

Top of the Batch: Interviews and the Match

Federico Echenique* Ruy González† Alistair Wilson‡ Leeat Yariv§¶

May, 2021

Abstract

Most doctors in the NRMP match with one of their most-preferred internship programs. However, surveys indicate doctors' preferences are similar, suggesting a puzzle: how can so many doctors match with their top choices when positions are scarce? We provide one possible explanation. We show that the patterns in the NRMP data may be an artifact of the interview process that precedes the match. Our study highlights the importance of understanding market interactions occurring before and after a matching clearinghouse. It casts doubts on analyses of clearinghouses that take reported preferences at face value.

Key words: NRMP, Deferred acceptance, Interviews, First-rank matches

JEL: C78, D47, J44

*Division of the Humanities and Social Sciences, Caltech.

†Caltech.

‡Department of Economics, University of Pittsburgh.

§Department of Economics, Princeton University, CEPR, and NBER.

¶Echenique gratefully acknowledges the support of NSF SES-1558757 and CNS-1518941. Yariv gratefully acknowledges the support of NSF grant SES-1629613. Detailed comments from SangMok Lee, Juan Pereyra, the coeditor, and four anonymous reviewers were very helpful in revising and improving the paper. A suggestion from SangMok Lee allowed us to generalize our main result.

1 Introduction

The National Resident Matching Program (NRMP) has matched millions of doctors to residency programs across the United States. In 2020 alone, 45,000 active applicants attempted to match with just over 37,000 positions. Match results reported by the NRMP for 2020 suggest comforting news for doctors: 46% of freshly-minted MDs from US schools were matched to their first-ranked choice, while 71% were matched to one of their top-three choices. The 2020 figures are by no means an aberration. The fraction of applicants matched to their first-ranked choice has been at least as high over the past two decades. We suggest these figures should not be taken at face value. In particular, we show that interactions outside of the main match—the interview process that precedes it—may be at least as important as the match itself.

Why should a large fraction of doctors matching to their first-ranked residency be surprising? The algorithm governing the NRMP match implements a stable matching over reported preferences. If applicants report similar preferences, only a few applicants can get their most-preferred option. For example, suppose 100 prospective residents are matched to 100 positions. Common preferences on both sides (an assortative market) yield an outcome in which only 1% are matched to their first-ranked program. As we show, even a small common component in doctors' preferences implies relatively few matches to top-ranked hospitals.

One explanation for the NRMP outcomes is that applicants' preferences are diametrically opposed, with only a handful ranking each position as their top outcome. This stands in the face of survey data and estimation results suggesting the importance of preference commonalities (Rees-Jones, 2018; Agarwal, 2015). Another explanation might be that preferences are similar across participants, but that each doctor and hospital consider only their top- k partners acceptable, as in Immorlica and Mahdian (2015). Matched participants would then have to receive a top- k outcome. As we demonstrate, this explanation too has shortcomings. First, it does not explain the relative prevalence of matches with first-ranked partners. Second, for k small enough to generate an effect, many applicants would be unmatched (see Arnosti, 2015; Beyhaghi and Tardos, 2018; Lee, 2016).

We propose another story. In the months preceding the centralized match, applicants submit their biographical and academic records, personal statements, and

letters of recommendation to hospitals.¹ From these applications, the hospitals select a group to interview. The process used to determine who interviews with whom is decentralized with two important features. First, applications and interviews are costly—both market sides have limited capacity. Second, hospitals and doctors submit rankings to the NRMP only for those *they interviewed with*.²

We assume that hospitals’ and applicants’ preferences are decomposable into common and idiosyncratic components. For hospitals, the common component may reflect doctors’ academic performance, test scores, and the strength of their letters (Agarwal, 2015). For doctors, it may reflect hospital rankings, quality of life in the local area, etc. In contrast, the idiosyncratic component reflects match-specific values. Using this model of preferences we consider the pre-match interview-selection process. Each candidate has a limit on the number of interviews they can attend, k , while each hospital has a maximum number of interview slots they can offer, k' . The decentralized interview process is then modeled as a stable many-to-many matching under the (k, k') capacity constraints. In the centralized matching stage, participants report rankings over interview partners only—their “interview-truncated” preference.

Truncation induced by interviewing necessarily narrows agents’ original preferences. However, unlike the truncation to the top- k , it is endogenous. Participants’ preferences are linked through the stability of interviews, so a large fraction of prospective doctors still end up matched—indeed, in our simulations, participants often end up with the same exact partner as in the centralized match with untruncated preferences. Moreover, the *reported* ranks for match outcomes are greatly inflated.

The presence of a common component in the preferences of prospective residents and hospitals is crucial for this conclusion. We show that with sufficient disagreement in doctors’ preferences, interviews may cause matched partners’ reported rank to *go down*, not up. While perfect agreement among doctors over hospital rankings will clearly lead to inflated rankings for matched programs under interview truncation, this obviously represents an extreme.³ Our main theoret-

¹See <https://www.nrmp.org/applying-interviewing-residency-programs/>

²The 2019 NRMP Applicant Survey (available from nrmp.org) provides median respondent data across four applicant-types and 21 medical-specialties. Of the 84 medians reported, 63 have perfectly coincident numbers for interviews attended and programs ranked, where 81 are ± 1 .

³A related idea appears in Beyhaghi and Tardos (2018), who show that interviews may increase the size of a match. See also Kadam (2015).

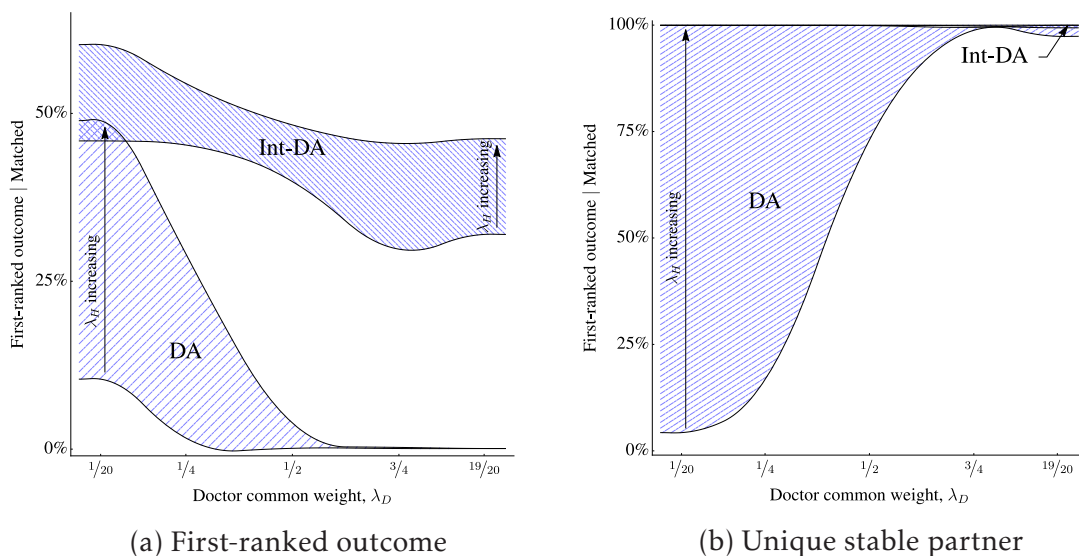


FIGURE 1: Simulated matched-resident outcomes (NRMP scale)

Note: Simulations under both DA and Int-DA algorithms are used to assess outcomes across many market sizes N . A composite response across 35,704 positions is then imputed across the speciality sub-markets (ranging in size from 22–9,127 positions) listed in the 2020 NRMP report (Table 13).

ical finding is that, in large markets, an *arbitrarily weak* common-component is sufficient for interviews to generate the pattern of high-reported ranks for match partners.

As our most-general result is asymptotic, we complement it with an array of simulations at more-moderate market sizes. Using the listed positions for the sub-markets that residents match within (anesthesiology through vascular surgery), we use our simulations to infer an aggregate-level outcome at the NRMP scale. We illustrate the results in Figure 1, serving both as a guide for the puzzle that motivates us, and the paper’s main result.

Figure 1 illustrates two measures for which data is available on NRMP outcomes: (a) The fraction of matched residents getting the hospital at the very top of their rank-order list; (b) Uniqueness, measured as the fraction of doctors with a unique stable match partner given the submitted rank-order lists. The figure illustrates simulated outcomes under both deferred acceptance with the full preferences (the region labeled *DA*) and interview-truncation (the region labeled *Int-DA*). In both figures we sweep across an array of possible preferences that modify the strength of the common component. On the horizontal axis we indicate the effect of increasing the common component weight within doctor’s preferences,

λ_D (with a residual weight $1 - \lambda_D$ on the idiosyncratic component). The width of each region then indicates the range in effect as we shift the common-component weight in hospital’s preferences, λ_H (where an arrow indicates the direction).

As mentioned, across two decades of annual NRMP matches, approximately one half of all matched residents obtain their first-ranked outcome.⁴ Figure 1(A) illustrates that, while the high fraction matched to their first-ranked outcome is possible under deferred acceptance, it requires a very particular form of preference: hospitals with a heavy weight on the common-value (λ_H close to one) and doctors with entirely idiosyncratic preferences (λ_D close to zero). For all other preferences, our simulations indicate the fraction of first-ranked matches is considerably lower. In particular, this statistic is close to zero whenever doctors place substantial weight on the common preference component.

While hospitals having a strong common component is consistent with NRMP survey data, the requirement that doctors’ preferences are almost completely idiosyncratic contradicts ample survey evidence.⁵ Simulation results under interview-truncation in the region labeled *Int-DA* illustrate our explanation. Our model of interviews leads to a near-opposite result to deferred acceptance: Except for a very particular and empirically unlikely preference—hospitals (doctors) with a small (large) common-value component—the NRMP-scale simulations lead to an approximately half of the residents obtaining their first-ranked outcome.

In Figure 1(B) we focus on another documented feature of the NRMP rankings: almost every doctor has a unique stable match partner. Using proprietary rank-order data from the NRMP, Roth and Peranson document small cores, where doctor-proposing and hospital-proposing deferred acceptance produce the same partner for 99.9% of the residents. Our simulation results in Figure 1(B) indicate the type of preferences that can generate this level of uniqueness within DA. The result that almost every participant has a unique stable match partner is a generic feature across all preference weights under the interview-truncated rankings (this is the close-to-degenerate region labeled at the very top of the figure). In contrast, this empirical feature is only possible under the full ranking when one of the two

⁴See Figure B.1 in Online Appendix B for the time-series.

⁵From the NRMP Director’s Survey, hospitals clearly place substantial weight on common features: test scores, recommendation letters, etc. However, the resident survey has common-value components (“reputation of program,” having an “academic medical center program,” as well as quality of the residents, faculty, and educational curriculum) cited at similar frequencies to potentially idiosyncratic ones (“perceived goodness of fit”) as the reasons for their ranking of programs.

sides' preferences are primarily driven by a common component.

The idea that doctors' reports in the residency match may not reflect true preferences is certainly present in other work. Hassidim, Marciano, Romm, and Shorrer (2017) survey evidence of misreports in the NRMP, suggesting four possible explanations: proposers' failure to identify the dominant strategy, mistrust in the mechanism, non-classical utility, and self-selection. The last of these is closest to ours. In this vein, Chen and Pereyra (2019) consider school-choice problems where students "self select" by only ranking schools they believe will plausibly admit them, showing evidence for this self-selection in Mexican high-school applications. While doctors and hospitals ranking only those they interview with is a manifestation of self-selection, our theoretical analysis offers a constructive process, shedding light on its underpinnings and impacts.⁶

Our results have important implications for the NRMP, and the matching literature more broadly. Doctors participating in the deferred-acceptance algorithm underlying the match have incentives to truthfully report preferences (Roth and Peranson, 1999). Traditionally, economists have viewed the NRMP as an ideal case-study in strategy-proof design. Our findings make clear that because reported preferences in the NRMP are filtered through the interview stage, they should be interpreted with caution. In particular, reported high-rank matches cannot be read literally, and any conclusions drawn about welfare using estimated preferences from the match itself are suspect. This message is particularly stark given that our paper ignores strategic effects at the interview stage.⁷ That said, our approach also indicates some constructive ways forward and suggests the potential importance of accounting for interactions preceding centralized clearinghouses.

⁶Lee and Schwarz (2017) also consider an interview process that precedes a centralized match. In their setting, workers are fully informed of their preferences, while firms view workers symmetrically at the outset and use costly interviews to infer their own preferences. In the NRMP context, Rees-Jones (2018) uses surveys to illustrate doctors' significant "misreporting" in the match. Rees-Jones and Skowronek (2018) use an online experiment with post-match residents, where 23% misrepresent their preferences in an incentivized NRMP-like matching task.

⁷See Beyhaghi, Saban, and Tardos (2017) for an analysis of some strategic implications of interviews and Manjunath and Morrill (2021) for the effects of changes in the interviewing costs on doctors' welfare.

2 The Model

Our model is a variant of the standard two-sided matching model (see, for example, Roth and Sotomayor, 1990), with an added interview stage.

2.1 Basic Definitions

A *market* is a triple (H, D, U) , where: H is a finite set of *hospitals*; D is a finite set of *doctors*; and $U = ((u_d)_{d \in D}, (u_h)_{h \in H})$ is a *utility function profile* (with $u_d : H \cup \{d\} \rightarrow \mathbf{R}$ and $u_h : D \cup \{h\} \rightarrow \mathbf{R}$ for each d and h).

A utility u_a induces an ordinal preference \succeq_a over the relevant set of alternatives, where we assume throughout that the resulting ordinal preferences are *strict*. The *rank-order* of b in u_a is one plus the number of b' with $u_a(b') > u_a(b)$ —a lower rank-order indicates a better ordinal outcome/higher ranking. In particular, agent a 's most-preferred match partner has rank-order 1. An agent b is *unacceptable* for a if $u_a(a) > u_a(b)$.

A *matching* is a function $\mu : H \cup D \rightarrow H \cup D$ with the properties that $\mu(h) \in D \cup \{h\}$, $\mu(d) \in H \cup \{d\}$, and $\mu(d) = h$ iff $\mu(h) = d$. A matching μ is *stable* for a market (H, D, U) if $u_a(\mu(a)) \geq u_a(a)$ for all $a \in D \cup H$, and there is no $(d, h) \in D \times H$ with $u_d(h) > u_d(\mu(d))$ and $u_h(d) > u_h(\mu(h))$.

A *many-to-many matching* is a function $\mu : H \cup D \rightarrow 2^{H \cup D}$ with the properties that $\mu(d) \subseteq H$, $\mu(h) \subseteq D$, and $h \in \mu(d)$ iff $d \in \mu(h)$. When an agent a is unassigned, we have $\mu(a) = \emptyset$. Given a pair of positive integers (k, k') , a many-to-many matching μ is *pairwise stable* for (k, k') if

- $|\mu(d)| \leq k$ and there is no $h \in \mu(d)$ with $u_d(h) < u_d(d)$;
- $|\mu(h)| \leq k'$ and there is no $d \in \mu(h)$ with $u_h(d) < u_h(h)$;
- There is no (h, d) such that $d \notin \mu(h)$ and any one of the following:
 - $u_d(h) > u_d(h')$ and $u_h(d) > u_h(d')$ for some $(h', d') \in \mu(d) \times \mu(h)$;
 - $u_d(h) > u_d(h')$, $u_h(d) > u_h(h)$, and $|\mu(h)| < k'$ for some $h' \in \mu(d)$;
 - $u_d(h) > u_d(d)$, $u_h(d) > u_h(d')$, and $|\mu(d)| < k$ for some $d' \in \mu(h)$.

2.2 Interview Schedules

In our model, doctors and hospitals first schedule interviews and then participate in the match.

An *interview schedule* is a many-to-many matching. Given a pair of integers (k, k') , a (k, k') -constrained *interview schedule* is a many-to-many matching μ with $|\mu(d)| \leq k$ and $|\mu(h)| \leq k'$ for all d and h . Each doctor can interview with at most k hospitals, and each hospital can interview at most k' doctors.

Given an interview schedule μ , agents' *interview-truncated* preferences are determined by setting $u_a(b) < u_a(a)$ for all $b \notin \mu(a)$. That is, interview-truncated preferences rank all interviewed agents as in the original preferences, and set all other agents as unacceptable.

The timing in our model is: (i) An interview schedule is determined as the doctor-optimal many-to-many (k, k') -stable matching;⁸ (ii) Doctors and hospitals report their interview-truncated preferences as inputs into doctor-proposing DA. The outcome is the doctor-optimal stable matching on the interview-truncated preferences. We term this two-step process **Int-DA**: the **I**nterview process followed by **D**eferred **A**cceptance.

A doctor-optimal interview schedule can be found algorithmically using the “T-algorithm” (see Blair, 1988; Fleiner, 2003; Echenique and Oviedo, 2006). We assume it is the result of a decentralized interview scheduling process. While modeling explicitly the interview process would certainly be desirable, our focus is on the tension between a “pure” application of DA, and one that is preceded by interviews. Assuming a stable outcome at the interview stage provides us with a simple, tractable model.⁹ A richer model might allow for some information to be transmitted at the interview stage. This could further inflate rank differences for two reasons.¹⁰ First, some doctors and hospitals may learn from interviews that they are not acceptable to one another, thereby limiting further the set of ranked

⁸Arguably, the doctor-optimal stable matching at the interview stage yields a smaller difference between reported and actual ranks than other selections of stable matchings.

⁹In one-to-one matching markets, experimental evidence suggests decentralized interactions yield stable outcomes at high rates, see Echenique and Yariv (2013). Melcher, Ashlagi, and Wapnir (2018) propose a stable-matching algorithm for internship interviews. For more on the theory of many-to-many matching, see Sotomayor (1999); Konishi and Ünver (2006).

¹⁰Indeed, while simulations where the idiosyncratic component is only revealed during interviews yield qualitatively similar results, quantitatively they do increase the effects (see Figure B.6 in Online Appendix B).

participants. Second, participants would not discover others whom they would, with full information, rank very highly but that they did not interview with.

We denote the final matching from Int-DA as μ^I . We compare μ^I to the matching obtained from the doctor-proposing DA algorithm using agents' original preferences, μ^{DA} .

3 The Impacts of Interviews

We start by illustrating that when doctors' preferences are misaligned, interviews alone cannot explain the findings in the data: Int-DA does *not* necessarily yield better-ranked partners in submitted preferences. We then show that some preference alignment, in the presence of interviews, leads to patterns similar to those in the data.

3.1 Preference Misalignment

To illustrate that interviews can have the opposite effect when doctors preferences are misaligned, consider a matching market with three doctors, $\{d_1, d_2, d_3\}$, and four hospitals, $\{h_1, h_2, h_3, h_4\}$.¹¹ Hospitals' preferences are common: they all prefer d_1 to d_2 , d_2 to d_3 , and d_3 to staying unmatched. Doctors rank all hospitals as acceptable, with preferences given by (first to last):

$$\begin{aligned} d_1: & h_1, h_3, h_2, h_4; \\ d_2: & h_2, h_3, h_1, h_4; \\ d_3: & h_3, h_1, h_4, h_2. \end{aligned}$$

Under DA, d_i matches to h_i . So the rank-order of d_3 's match is 1.

Suppose interview constraints are $k = k' = 2$. All doctors want to interview with h_3 , but only d_1 and d_2 are able to. The resulting interview schedule is: d_1 with h_1 and h_3 ; d_2 with h_2 and h_3 ; and d_3 with h_1 and h_4 .

Given the interview-truncated preferences, d_i matches with h_i for $i = 1, 2$, but d_3 is matched with h_4 . The Int-DA rank-order of d_3 's match is therefore 2, so the presence of interviews leads her to a strictly lower rank-order than under DA

¹¹It is straightforward to construct more intricate examples with equal numbers of doctors and hospitals.

without interviews. Furthermore, the outcome under Int-DA is unstable for the original preferences.

In this example, there is substantial disagreement between doctors' preferences—there are no pairwise comparisons of hospitals $\{h_1, h_2, h_3\}$ on which doctors agree. In contrast, our discussion of the NRMP data emphasized the role of common components in the doctors' and hospitals' preferences. We now show that some degree of agreement between doctors' ranking of hospitals rules out such examples, and interviews can explain observed high match ranks.

3.2 Aligned Preferences

We start with the extreme case where doctors' preferences are common.

Proposition 1. *Suppose $k = k'$ and that doctors' preferences are identical. For any doctor d , the rank-order of $\mu^I(d)$ in her interview-truncated preference is always weakly lower than the rank-order of $\mu^{DA}(d)$ in her actual preference \succeq_d .*

The proof appears in Online Appendix A. Intuitively, when doctors' preferences are common, only one of the doctors under DA is matched to the highest-ranked hospital, one to the second-highest, etc. In particular, $n - k$ doctors are matched to a hospital ranked below their top k . In contrast, interviews allow for presorting of doctors to hospitals they have a chance of matching with. Interviews also limit how low a matched hospital can be ranked in the reported preferences: it can never be lower than k .

The proposition assumes $k = k'$ mainly for expository reasons. In our main result below we allow the two capacities to differ.

3.3 Large Markets

We now demonstrate that perfect preference alignment is not necessary. As long as there is a common-value component in agents' preferences, *however small*, the message of our first result holds in large markets.

We expand the model to account for market size and randomly generated preferences. For each n , let $E_n = (D_n, H_n, U_n, k_n, k'_n)$ denote a market, where $D_n = \{d_1, \dots, d_n\}$, $H_n = \{h_1, \dots, h_n\}$, each utility function is randomly drawn with a common-value and idiosyncratic component, and k_n (k'_n) is the limit on doctors' (hospitals')

interviews. Specifically, suppose that

$$u_d^n(h) = \lambda_D c_h + (1 - \lambda_D) \eta_{d,h} \text{ and } u_h^n(d) = \lambda_H c_d + (1 - \lambda_H) \eta_{h,d},$$

for all $d \in D_n$ and $h \in H_n$, where $\lambda_D, \lambda_H \in (0, 1)$. Moreover, assume that $u_a^n(a) = 0$. The common-value components c_h and c_d are crucial for our results, but need not dominate doctors' utilities: $\lambda_D, \lambda_H > 0$ may be arbitrarily small.

Suppose that $c_h, c_d, \eta_{d,h}$, and $\eta_{h,d}$ are all drawn from absolutely continuous distributions with support on a convex subset of \mathbf{R}_+ . Let μ_n^I denote the matching resulting from the Int-DA process in the n -sized market E_n , and μ_n^{DA} the corresponding outcome of the doctor-proposing DA; these matchings are random and depend on realized utilities. The Int-DA procedure determines a matching μ_n^I by choosing a (k_n, k'_n) -constrained interview schedule $\hat{\mu}$ as the doctor-optimal many-to-many stable matching, followed by the doctor-proposing DA using the induced preferences.

Proposition 2. *Let $E_n = (D_n, H_n, U_n, k_n, k'_n)$ be a sequence of markets satisfying our assumptions. Fix an integer sequence $M_n \geq 1$, and let D'_n be the set of doctors $d \in D_n$ in market E_n for whom the rank of $\mu_n^I(d)$ in their interview-truncated preferences is M_n positions above the rank of $\mu_n^{DA}(d)$ in their true preference. If k_n and M_n are $o(n)$, then, for any $\theta > 0$,*

$$\lim_{n \rightarrow \infty} P\left(\frac{1}{n} |D'_n| \geq 1 - \theta\right) = 1.$$

The proof of Proposition 2 appears in the Appendix. The idea underlying it is simple. Consider DA, let $\varepsilon > 0$, and fix $M_n = 1$. By Lee (2016), when n is large, with high probability, the set $A_n(\varepsilon, (c, \eta))$ of doctors that are within ε of their “target” assortative utility in DA accounts for at least $1 - \theta/2$ of all doctors. Let $B(c_d, n)$ be the event that fewer than k_n hospitals provide doctor d a utility greater than d 's target utility. We denote by β_n the probability that a fraction of at least $\theta/2$ doctors have a “small” number (at most k_n) of hospitals above their target utility. We show that for n large enough, $\beta_n < \pi/2$, and by Lee (2016), $P\left(\frac{1}{n} |A_n(\varepsilon, (c, \eta))| \geq 1 - \theta/2\right) > 1 - \pi/2$. Thus, the event that $B(c_d, n)$ is false for a fraction $\geq 1 - \theta/2$ of doctors and $\frac{1}{n} |A_n(\varepsilon, (c, \varepsilon))| \geq 1 - \theta/2$, has probability $\geq (1 - \pi/2) + (1 - \pi/2) - 1 = 1 - \pi$. At the intersection of these conditions, for a fraction $\geq (1 - \theta/2) + (1 - \theta/2) - 1 = 1 - \theta$ of $d \in D_n$, we have that $B(c_d, n)$ is false and $d \in A_n(\varepsilon)$. Hence, for a fraction $\geq 1 - \theta$ of $d \in D_n$, there are more than k_n hospitals above their target utility, and they are

within ε of their target utilities.

Convergence rates for the large-market result in Proposition 2 are modest, with (poly-)logarithmic or polynomial growth in the relevant “approximation guarantees” θ and π . In words, the market size needed for Proposition 2 does not grow too quickly with the approximation guarantees. This message complements the simulations in Section 4, which assume (arguably) realistic market sizes, and can be formalized as follows (a proof is in Online Appendix A):

Proposition 3. *Let E_n and D'_n be as in the statement of Proposition 2. Fix $\theta, \delta \in (0, 1)$. Then $P\left(\frac{1}{n}|D'_n| \geq 1 - \theta\right) \geq 1 - \delta$ for $n = \Theta((\ln(1/\pi))^4)$ as $\pi \rightarrow 0$, and $n = \Theta((1/\theta)^4)$ as $\theta \rightarrow 0$.*

4 Simulations

Our theoretical findings raise three important questions. The first regards market size. Proposition 2 is asymptotic, and it is natural to consider whether interviews matter for smaller, more realistic, market sizes. The second question regards unmatched agents. One might worry that interview-truncated preferences give rise to large numbers of unmatched participants, beyond those observed in the NRMP. The final question regards stability. Ideally, the difference between outcomes under DA and the interview-truncated DA procedure would be small.

We address these questions using extensive numerical simulations across the preference parameters λ_D and λ_H . In total, we examine 275 different parameterizations, representing approximately 1 million simulated participants. We also consider a variety of robustness checks. Section B.2 in Online Appendix B provides a summary and more-detailed results. To compactly display our main results, Table 1 aggregates our findings across balanced markets (N doctors and N positions) corresponding to the many smaller sub-markets that compose the NRMP.¹² We examine three algorithmic solutions:¹³

¹²For each pair (λ_D, λ_H) and simulation proportion f_N , we fit a linear model for $\log(f_N/(1 - f_N))$ against $\log(N)$, where goodness of fit for this approach is shown in Figure B.4 in Online Appendix B. We predict the proportion $\hat{f}(N_j)$ for each sub-market of size N_j listed in Table 13 of the 2020 NRMP report (from 22 positions for Pediatrics/Medical Genetics to 9,127 for Internal Medicine). We then aggregate across these sub-markets to form a composite NRMP measure for f over the $N_T = 35,704$ positions as $\sum_j \frac{N_j}{N_T} \hat{f}(N_j)$.

¹³To assess robustness, we also conducted simulations examining: (i) imbalanced markets; (ii)

- Doctor- and hospital-proposing deferred acceptance (DA).
- The stable interview allocation with $k = k' = 5$ slots per position, followed by both doctor- and hospital-proposing deferred acceptance on the interview-truncated preferences (Int-DA).¹⁴
- Doctor- and hospital-proposing deferred acceptance on preferences truncated to the $k = k' = 5$ top-ranked options (Tr-DA).

Table 1 provides NRMP-scaled simulation outcomes for six different (λ_D, λ_H) pairs. Simulations with Tr-DA were added to distinguish the pure effect of truncation from the interview process our paper focuses on.¹⁵ The first panel in the table provides three characteristics of the match outcome: (i) the fraction of *unmatched* participants; (ii) the fraction of doctors matched to their *first-ranked program*; and (iii) the proportion of doctors matched to a *top-three-ranked program*.

Because our simulated markets have the same participant volume on each side, with all possible matches acceptable, the benchmark for DA with full preferences predicts no unmatched doctors. In contrast, the NRMP data indicates that 5.8% of US seniors are unmatched. The first result from our simulations in Table 1 illustrates that the two-stage Int-DA process leads to a similar unmatched rate as the NRMP. Doctors in our simulations are unmatched after the Int-DA process at a 5.5% rate. Moreover, this proportion does not change substantially with either market size or the common weight. In contrast, a direct truncation to the top-five participants on the other side leads to substantially more unmatched participants. Moreover, the unmatched rate grows sharply with increases to N and λ .

The next pair of results from the Int-DA simulations again match the NRMP data: a large fraction of doctors are matched to top-ranked hospitals. Looking to NRMP data from the past five years, 48% (73%) of US MD Seniors are matched to their first-ranked (top-three-ranked) program. The Int-DA simulations indicate

smaller/larger interview capacities; (iii) more extreme values for λ_D and λ_H ; (iv) an alternative interview selection method where hospitals only used the doctor's common components. Online Appendix B presents details on these simulations. The main effects of the interview stage are similar across these exercises.

¹⁴To mirror the theory section, we select the doctor-optimal stable interview schedule.

¹⁵In simulations where we allowed for information to be realized at the interview stage, hospitals ranked doctors purely on the common component. The idiosyncratic component is observed at the interviews, and used in the DA stage. Results are similar to the Int-DA procedure (see Figure B.6 in Online Appendix B) and suggest that simpler procedures that incorporate both sides preferences can also be used.

TABLE 1: Simulation Outcomes scaled to NRMP size

	$\lambda_H = 1/4$			$\lambda_H = 3/4$		
	$\lambda_D = 1/4$	$\lambda_D = 1/2$	$\lambda_D = 3/4$	$\lambda_D = 1/4$	$\lambda_D = 1/2$	$\lambda_D = 3/4$
Panel A: Matching outcomes						
	<i>Unmatched</i>			[DA: 0.0%, NRMP: 5.4%] [†]		
Int-DA	6.0%	6.4%	8.1%	8.2%	6.5%	5.6%
Tr-DA	26.1%	72.1%	96.0%	27.1%	71.7%	96.2%
	<i>First-ranked program</i>			[NRMP: 48.1%] [†]		
DA	2.5%	0.2%	0.1%	26.5%	3.0%	0.2%
Int-DA	43.5%	38.7%	32.5%	49.1%	43.7%	41.4%
Tr-DA	22.7%	4.1%	0.2%	31.2%	5.2%	0.2%
	<i>Top-three-ranked program</i>			[NRMP: 73.2%] [†]		
DA	7.4%	0.5%	0.3%	48.4%	8.1%	0.6%
Int-DA	81.6%	79.6%	75.0%	81.6%	81.5%	81.2%
Tr-DA	55.0%	15.1%	1.3%	59.1%	17.2%	1.3%
Panel B: Core size, similarity to DA, and stability						
	<i>Same partner under proposer change Matched</i> [NRMP: 99.9%] [‡]					
DA	41.9%	90.0%	99.5%	99.3%	99.2%	98.8%
Int-DA	99.9%	99.9%	100.0%	100.0%	99.9%	99.8%
	<i>Identical partner to DA Matched</i>					
Int-DA	73.8%	82.6%	79.8%	80.1%	74.6%	81.0%
	<i>Proportion blocking programs in Int-DA</i>					
Matched	0.1%	0.6%	1.0%	0.1%	0.7%	1.8%
Unmatched	9.1%	7.3%	8.3%	19.2%	24.6%	34.3%

Note: Table uses an estimated logit-model for each (λ_D, λ_H) pair to form an estimate at NRMP scale (~35,000 positions) using the sub-market sizes listed for the 2020 NRMP. See Online Appendix for tables at separate N †–Average for US MD Seniors in 2016–20. Source: *Results and Data: 2020 Main Residency Match*, Table 15, available from nrmp.org. ‡–Figure reported for main NRMP match in Roth and Peranson (1999). Smaller thoracic surgery market ($N \simeq 120$) has a 99.6 percent unique match for five reported years in 1991–96 (*ibid*, tables 1 and 3).

similarly-sized effects to the observed NRMP figures, at 40–50% (75–82%).¹⁶ In contrast, the pure DA algorithm on the full preference lists implies substantially lower rates of top-ranked outcomes—matches to the first-ranked program are only found at substantial rates when hospitals have strong common components, and doctors are almost fully idiosyncratic (as depicted in Figure 1).

In the second panel of Table 1 we turn to other observed match outcomes. These outcomes are not part of our explanation of reduced match ranks, but serve to evaluate the empirical relevance of our interviews model. The first outcome is motivated by Roth and Peranson’s (1999) finding that NRMP data exhibit small cores. Using NRMP ranking data from the 1990s, they examine the change in outcomes moving from the doctor- to the hospital-proposing DA. They find that 99.9% of doctors receive the same outcome, implying a unique stable partner. In the *same partner under proposer change* rows we mirror this exercise. Our DA simulations get close to the NRMP figure only with heavy weights on the common component. While most participants across each of the simulations do have a unique stable partner, the minority with multiple partners are at least an order of magnitude larger than in Roth and Peranson (1999). However, changing the proposing side over the interview-truncated rankings from Int-DA indicates much-closer effects to the NRMP field study.

Our simulations of the Int-DA procedure show that it can reproduce stylized results reflective of the observed NRMP figures—over unmatched rates, over the fraction of first-ranked outcomes, and over the small cores found in rank-order list data. Moreover, the Int-DA process does so generically, across market sizes and the common-preference weights.¹⁷

Given the fit with observed data regularities, a natural question regards the difference between outcomes under Int-DA and standard DA. The final set of results in Table 1 speaks to this question.

The *identical partner to DA* row directly contrasts the Int-DA and DA match outcomes. Our results illustrate that the large majority of matched doctors (~75–80%) in the Int-DA procedure are matched to the *exact same partner* they would

¹⁶The Int-DA fraction matched to their first-ranked program does increase slightly as we increase N (see Tables B.1–B.6 in Online Appendix B) and as we decrease λ_D /increase λ_H .

¹⁷We also examined imbalanced markets, with an excess of doctors. The Int-DA results for these simulations are qualitatively similar. However, we note that DA results allowing imbalance exhibit more frequent unique stable matchings Ashlagi, Kanoria, and Leshno (per 2017) (see Tables B.7 and Figure B.3 in Online Appendix B).

have matched with under DA with full preferences reports.¹⁸ While four of every five doctors are entirely unaffected by the interview process, one in five being affected is clearly far from negligible.¹⁹

In the last section of Table 1 we evaluate the effects on stability. For each doctor in our simulations we calculate the proportion of programs they form a blocking pair with. We report the average proportion, distinguishing between matched and unmatched doctors. Matched doctors exhibit some instability, despite both stages in the two-stage process being chosen to select stable outcomes. Averaging across parameterizations, a blocking pair is detected for matched doctors 0.7% of the time. Unsurprisingly, instabilities are more substantial for unmatched doctors. A randomly chosen hospital yields a blocking pair between 8% and 34% of the time for each unmatched doctor, depending on the parameterization.

5 Conclusion

Much of the matching literature has focused on the centralized clearinghouse governing the match of newly-minted doctors and residency positions. We illustrate the possibility that *decentralized* interactions preceding the match—namely, interviews—may dramatically impact ultimate outcomes.

For the NRMP, our results imply that empirical estimations based on preferences submitted to the clearinghouse should not be taken at face value. More broadly, beyond the NRMP, our paper suggests that interactions outside of the clearinghouse can have dramatic effects on outcomes.

Our model focuses on the role interviews can have purely in terms of selection. In our setting, there is no constructive role for information acquisition during the interviews (see, however, Figure B.6 in Online Appendix B, for one model variant allowing for information transmission during the interviews). Our algorithmic approach offers a base for future research to explore richer interactions preceding matching clearinghouses.

¹⁸In a series of robustness exercises we also examine balanced-market simulations with different values of k (see Tables B.8–B.11 and Figure B.5 in Online Appendix A), indicating that the DA and Int-DA outcomes become more similar as k increases. Nonetheless, we still find that, among matched doctors, ~75% get the same exact match as under DA even when $k = 2$.

¹⁹For those matched to distinct partners under Int-DA and DA, matches tend to be close in ranking terms. See Figures B.7 and B.7 in Online Appendix A for illustration.

References

- Agarwal, Nikhil (2015), “An empirical model of the medical match.” *The American Economic Review*, 105, 1939–78.
- Arnosti, Nick (2015), “Short lists in centralized clearinghouses.” In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, 751, Association for Computing Machinery, New York, NY, USA, URL <https://doi.org/10.1145/2764468.2764533>.
- Ashlagi, Itai, Yash Kanoria, and Jacob D Leshno (2017), “Unbalanced random matching markets: The stark effect of competition.” *Journal of Political Economy*, 125, 69–98.
- Beyhaghi, Hedyeh, Daniela Saban, and Eva Tardos (2017), “Effect of selfish choices in deferred acceptance with short lists.” MATCH-UP.
- Beyhaghi, Hedyeh and Eva Tardos (2018), “Two-sided matching with limited interviews.” Mimeo, Cornell University.
- Blair, Charles (1988), “The lattice structure of the set of stable matchings with multiple partners.” *Mathematics of operations research*, 13, 619–628.
- Chen, Li and Juan Sebastián Pereyra (2019), “Self-selection in school choice.” *Games and Economic Behavior*, 117, 59 – 81, URL <http://www.sciencedirect.com/science/article/pii/S0899825619300843>.
- Echenique, Federico and Jorge Oviedo (2006), “A theory of stability in many-to-many matching markets.” *Theoretical Economics*, 1, 233–273.
- Echenique, Federico and Leeat Yariv (2013), “An experimental study of decentralized matching.” California Institute of Technology working paper.
- Fleiner, Tamás (2003), “A fixed-point approach to stable matchings and some applications.” *Mathematics of Operations Research*, 28, 103–126.
- Hassidim, Avinatan, Déborah Marciano, Assaf Romm, and Ran I. Shorrer (2017), “The mechanism is truthful, why aren’t you?” *American Economic Review*, 107, 220–24, URL <https://www.aeaweb.org/articles?id=10.1257/aer.p20171027>.

- Immorlica, Nicole and Mohammad Mahdian (2015), “Incentives in large random two-sided markets.” *ACM Trans. Econ. Comput.*, 3, URL <https://doi.org/10.1145/2656202>.
- Kadam, Sangram V (2015), “Interviewing in matching markets.” Mimeo Harvard University.
- Konishi, Hideo and M Utku Ünver (2006), “Credible group stability in many-to-many matching problems.” *Journal of Economic Theory*, 129, 57–80.
- Lee, Robin S. and Michael Schwarz (2017), “Interviewing in two-sided matching markets.” *The RAND Journal of Economics*, 48, 835–855, URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12193>.
- Lee, SangMok (2016), “Incentive compatibility of large centralized matching markets.” *The Review of Economic Studies*, 84, 444–463.
- Manjunath, Vikram and Thayer Morrill (2021), “Interview hoarding.” ArXiv:2102.06440.
- Melcher, Marc L., Itai Ashlagi, and Irene Wapnir (2018), “Matching for Fellowship Interviews.” *JAMA*, 320, 1639–1640, URL <https://doi.org/10.1001/jama.2018.13080>.
- Rees-Jones, Alex (2018), “Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match.” *Games and Economic Behavior*, 108, 317–330.
- Rees-Jones, Alex and Samuel Skowronek (2018), “An experimental investigation of preference misrepresentation in the residency match.” *Proceedings of the National Academy of Sciences*, 115, 11471–11476, URL <https://www.pnas.org/content/115/45/11471>.
- Roth, Alvin E and Elliott Peranson (1999), “The redesign of the matching market for american physicians: Some engineering aspects of economic design.” *American economic review*, 89, 748–780.
- Roth, Alvin E and Marilda Sotomayor (1990), *Two-sided Matching: A Study in Game-theoretic Modeling and Analysis*. Econometric Society Monographs volume 18, Cambridge University Press.

Sotomayor, Marilda (1999), “Three remarks on the many-to-many stable matching problem.” *Mathematical social sciences*, 38, 55–70.

A Proof of Proposition 2

Let $k_n'' = k_n + M_n$ and note that k_n'' is $o(n)$.

With some notational abuse, we drop λ^D and $1 - \lambda^D$, and write c_d for $\lambda^D c_d$, $\eta_{d,h}$ for $(1 - \lambda^D)\eta_{d,h}$, etc. This re-scaling implies that utilities are sums of the common and private value components: $u_d^n(h) = c_h + \eta_{d,h}$, and $u_h^n(d) = c_d + \eta_{h,d}$. The probability distributions are re-scaled correspondingly, but remain absolutely continuous. Without loss of generality, we assume distributions have compact support (otherwise, choose a compact set that accumulates large enough probability). Moreover, we take the support to be $[0, 1]$.

Let $D = \cup_n D_n$ and $H = \cup_n H_n$. Consider tuples (c, η) , with $c = (c_a)_{a \in H \cup D}$ and

$$\eta = ((\eta_{a,b})_{(a,b) \in H \times D}, (\eta_{a,b})_{(a,b) \in D \times H}).$$

The tuples (c, d) are endowed with the product probability measure from the i.i.d. distributions described.

Let G denote the cumulative distribution function corresponding to c_d . Fix $\theta, \pi \in (0, 1)$. Choose c^* and $\varepsilon, \delta \in (0, 1)$ such that $1 - G(c^*) + \delta < \theta/4$ while $0 < P(c_h + \eta_{d,h} > c^* + 1 + \varepsilon)$. These choices are possible due to the absolute continuity of the distributions of c_d , c_h , and $\eta_{d,h}$. Write $p(c^*)$ for $P(c_h + \eta_{d,h} > c^* + 1 + \varepsilon)$.

If agents match assortatively based on the common component, a doctor d should be able to find a hospital h for which it has idiosyncratic utility close to 1, and this hospital should provide d with (approximately) the same utility $c_d + 1$ as it receives from matching with d . Think of $c_d + 1$ as d 's “target utility.” Let

$$A_n(\varepsilon, (c, \eta)) = \{d \in D_n : c_d + 1 - \varepsilon < u_d(\mu_n^{DA}(d)) < c_d + 1 + \varepsilon\}$$

be the doctors for which this target is achieved (in DA), up to ε . We prove that, when n is large enough, with large probability, a fraction of at least $1 - \theta/2$ doctors are in $A_n(\varepsilon, (c, \eta))$.

Consider the number of hospitals ranked above a doctor's adjusted target util-

ity $c_d + 1 + \varepsilon$. Let

$$B(c_d, n) = \left\{ \left| \{h \in H_n : c_h + \eta_{d,h} > c_d + 1 + \varepsilon\} \right| \leq k_n'' \right\}$$

be the event that fewer than k_n'' hospitals provide d a utility greater than d 's target utility. We denote by β_n the probability that a fraction of at least $\theta/2$ doctors have a “small” number, at most k_n'' , of hospitals above their target utility.

We prove that for n large enough, $\beta_n < \pi/2$ and $P\left(\frac{1}{n} |A_n(\varepsilon, (c, \eta))| \geq 1 - \theta/2\right) > 1 - \pi/2$. Thus, the event that $B(c_d, n)$ is false for a fraction $\geq 1 - \theta/2$ of doctors *and* the event $\left(\frac{1}{n} |A_n(\varepsilon, (c, \varepsilon))| \geq 1 - \theta/2\right)$ holds, has probability $\geq (1 - \pi/2) + (1 - \pi/2) - 1 = 1 - \pi$. At the intersection of these events, it holds for a fraction $\geq (1 - \theta/2) + (1 - \theta/2) - 1 = 1 - \theta$ of $d \in D_n$ that $B(c_d, n)$ is false and $d \in A_n(\varepsilon)$. Hence, for a fraction $\geq 1 - \theta$ of $d \in D_n$ there are more than k_n'' hospitals above their adjusted target utility, and they are within ε of their target utilities, hence below the adjusted target. The rank-order of any partner in μ^I is at most k_n . Since $M_n = k_n'' - k_n$, these statements prove the proposition.

To complete the proof we present the required calculations.

If $c_d \leq c^\star$, then

$$\begin{aligned} P(B(c_d, n)) &= P\left(\sum_{h \in H_n} \mathbf{1}_{c_h + \eta_{d,h} > c_d + 1 + \varepsilon} \leq k_n''\right) \\ &\leq P\left(\frac{1}{n} \sum_{h \in H_n} \mathbf{1}_{c_h + \eta_{d,h} > c^\star + 1 + \varepsilon} \leq p(c^\star) - \left(p(c^\star) - \frac{k_n''}{n}\right)\right) \\ &\leq \exp\left(-2\left(p(c^\star) - \frac{k_n''}{n}\right)^2 n\right) \end{aligned} \tag{1}$$

by Hoeffding's inequality (observe that, eventually, $p(c^\star) - \frac{k_n''}{n} > 0$).

Let

$$\begin{aligned} \beta_n &= P(|\{d \in D_n : B(c_d, n)\}| > n\theta/2) \\ &\leq P\left(\underbrace{|\{d \in D_n : B(d, n) \text{ and } c_d \leq c^\star\}|}_{Z_n} + \underbrace{|\{d \in D_n : c_d > c^\star\}|}_{Y_n} > n\theta/2\right) \\ &\leq P\left(\frac{1}{n} Z_n + 1 - G(c^\star) + \delta > \theta/2\right) + P\left(\frac{1}{n} Y_n > 1 - G(c^\star) + \delta\right). \end{aligned}$$

The first inequality follows by counting all d with $c_d > c^\star$ as if $B(d, n)$ were true, so the random variable Y_n counts all $d \in D_n$ with $c_d > c^\star$ as if they were in $B(c_d, n)$.

The second inequality is a truncation exercise, partitioning the probability space into two events. The first is $\frac{1}{n}Y_n \leq 1 - G(c^\star) + \delta$ and the second $\frac{1}{n}Y_n > 1 - G(c^\star) + \delta$. Under the second event, $\frac{1}{n}Z_n + \frac{1}{n}Y_n > \theta/2$ as $1 - G(c^\star) + \delta > \theta/2$. Under the first event, the inequality is obtained by raising $\frac{1}{n}Y_n$ to $1 - G(c^\star) + \delta$.

Applying Hoeffding's inequality again,

$$P\left(\frac{1}{n}Y_n > 1 - G(c^\star) + \delta\right) \leq \exp(-2\delta^2 n). \quad (2)$$

Now,

$$\begin{aligned} P(Z_n > n(\theta/2 - [1 - G(c^\star) + \delta])) &\leq P(\cup_{d \in D_n} B(d, n) | c_d = c^\star) \\ &\leq \sum_{d \in D_n} P(B(d, n) | c_d = c^\star) \\ &\leq n \exp(-2(p(c^\star) - \frac{k''_n}{n})^2 n), \end{aligned} \quad (3)$$

where the first inequality follows as $n(\theta/2 - (1 - G(c^\star) + \delta)) \geq 1$, and the probability of $B(d, n)$ is maximized when $c_d = c^\star$.

Choose n such that

$$n(\theta/2 - [1 - G(c^\star) + \delta]) > 1, \quad (4)$$

$$\exp(-2\delta^2 n) < \pi/4, \quad (5)$$

$$n \exp(-2(p(c^\star) - \frac{k''_n}{n})^2 n) < \pi/4, \quad (6)$$

$$\text{and } P\left(\frac{1}{n} |A_n(\varepsilon, (c, \eta))| \geq 1 - \theta/2\right) > 1 - \pi/2. \quad (7)$$

Observe that (4) is possible as $\theta/2 - [1 - G(c^\star) + \delta] > 0$. Inequality (6) requires k''_n be $o(n)$, which holds by hypothesis, and our choice of c^\star to ensure that $p(c^\star) - k''_n/n > 0$ is eventually bounded away from zero. Inequality (7) is possible by Theorem 1 of Lee (2016).

By (2),(3),(5), and (6), we obtain that

$$\beta_n \leq n \exp(-2(p(c^\star) - \frac{k''_n}{n})^2 n) + \exp(-2\delta^2 n) < \pi/2. \quad (8)$$

Statements (7) and (8) provide the required bounds.