Inefficiencies in Residency Matching Associated with Gale–Shapley Algorithms

Yue Wu, PhD¹ Parisa Taravati, MD¹ Ryan T. Yanagihara, MD¹ Courtney E. Francis, MD¹
Marian Blazes, MD¹ Cecilia S. Lee, MD, MS¹ Aaron Y. Lee, MD, MSCI^{1,*} Russell N. Van Gelder, MD, PhD^{1,*}

¹ Department of Ophthalmology, University of Washington, Seattle, Washington

Address for correspondence Aaron Y. Lee, MD, MSCI, Department of Ophthalmology, University of Washington, Box 359608, 325 Ninth Avenue, Seattle, WA 98104 (e-mail: leey@uw.edu).

J Acad Ophthalmol 2021;13:e175–e182.

Abstract

Objective This study aimed to investigate emerging trends and increasing costs in the National Residency Matching Program (NRMP) and San Francisco Residency and Fellowship Match Services (SF Match) associated with the current applicant/program Gale–Shapley-type matching algorithms.

Design A longitudinal observational study of behavioral trends in national residency matching systems with modeling of match results with alternative parameters.

Patients and Methods We analyzed publicly available data from the SF Match and NRMP websites from 1985 to 2020 for trends in the total number of applicants and available positions, as well the average number of applications and interviews per applicant for multiple specialties. To understand these trends and the algorithms' effect on the residency programs and applicants, we analyzed anonymized rank list and match data for ophthalmology from the SF Match between 2011 and 2019. Match results using current match parameters, as well as under conditions in which applicant and/or program rank lists were truncated with finalized rank lists, were analyzed.

Results Both the number of applications and length of programs' rank lists have increased steadily throughout residency programs, particularly those with competitive specialties. Capping student rank lists at seven programs, or less than 80% of the average 8.9 programs currently ranked, results in a 0.71% decrease in the total number of positions filled. Similarly, capping program rank lists at seven applicants per spot, or less than 60% of the average 11.5 applicants ranked per spot, results in a 5% decrease in the total number of positions filled.

Conclusion While the number of ophthalmology positions in the United States has increased only modestly, the number of applications under consideration has increased substantially over the past two decades. The current study suggests that both programs and applicants rank more choices than are required for a nearly complete and stable match, creating excess cost and work for both applicants and programs. "Stable-marriage" type algorithms induce applicants and programs to rank as many counterparties as possible to maximize individual chances of optimizing the match.

Keywords

- ▶ resident matching
- ▶ stable marriage
- ▶ Gale–Shapley
- ▶ Nash equilibrium

* These authors contributed equally and should be considered co-senior authors.

received
January 4, 2021
accepted after revision
June 22, 2021

DOI <https://doi.org/10.1055/s-0041-1735951>.
ISSN 2475-4757.

© 2021. The Author(s).

This is an open access article published by Thieme under the terms of the Creative Commons Attribution-NonDerivative-NonCommercial-License, permitting copying and reproduction so long as the original work is given appropriate credit. Contents may not be used for commercial purposes, or adapted, remixed, transformed or built upon. (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Thieme Medical Publishers, Inc., 333 Seventh Avenue, 18th Floor, New York, NY 10001, USA

The San Francisco Residency and Fellowship Match Services (SF Match)¹ and the National Residency Match Program (NRMP)² are two national matching systems used to place physicians into residency and fellowship training programs. Both use versions of the Gale–Shapley algorithm³ to pair applicants with programs in a binding system that has been used for over 50 years. While there has been no significant alteration in these systems in the past half century, the landscape of the application process has changed substantially, mainly due to a growing number of applicants and participating programs each year while the number of applicants per position remains relatively stable. The year 2020 has the largest match numbers to date, with 40,084 applicants applying to 37,256 positions.⁴

In the 1920s, when the first residency programs were introduced as optional postgraduate training, only a few medical graduates participated. This inadequate supply of interns led to fierce competition among the programs, which manifested as a race between programs to secure binding commitments from potential graduates as early as possible.⁵ This resulted in medical students receiving internship offers up to 2 years before graduation.⁶ To avoid this race between programs, the National Interassociation Committee on Internships (NICI) was formed in 1950 to examine existing matching plans and performed a trial for a centralized match system. In October 1951, 79 medical schools formed the National Student Internship Committee and adopted a modified Boston Pool Plan⁷ nationally based on the recommendations of the NICI. The National Internship Matching Program (NIMP, now the NRMP) was incorporated in 1953 to manage and administer the matching process, and has continued to do so for most medical residency programs and many fellowship programs.⁸ The SF Match oversees ophthalmology and plastic surgery residency programs, as well as multiple specialty fellowship programs.

The residency matching dilemma can be described as a stable marriage problem, with the applicants as one side of the “marriage” and the residency programs as the other. A marriage or match is stable when there is no applicant matched to program A while preferring program B, when program B also prefers this applicant over at least one other candidate that is currently matched with program B. Gale and Shapley proved that for an equal number of participants on each side who have each ranked every potential partner, stable matches for all participants exist³ and their eponymous algorithm finds a solution. (The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel for 2012 was awarded to Lloyd S. Shapley for this work.) The resident matching algorithms used by NRMP and SF Match appear mathematically equivalent to the Gale–Shapley algorithm. The current NRMP match algorithm was implemented in 1995.

The Gale–Shapley algorithm takes the rank order lists (a ranked list of choices) from each of the participants on both sides along with a predetermined proposing side (either the programs or the applicants in this case). For example, if the applicants are the proposing side, the algorithm first selects an applicant at random from the pool of applicants. This

applicant will first propose to its most preferred program. If that program has open positions and has ranked that applicant, then a tentative match is formed between them and the algorithm picks another applicant to start proposing to his or her most preferred program. If the program's positions are already filled, the algorithm checks if the program would prefer the new proposing applicant over one of their currently matched applicants. If the program prefers the new proposing applicant, then the program's match with its least preferred previously matched candidate is annulled, and that candidate is added back to the applicant pool. The algorithm continues until all program positions have been filled.

Gale–Shapley requires choosing one proposing side: applicants versus programs. In addition, in the original formulation, both parties must rank all possible matches. The algorithm then works to achieve a stable match and fill every available position. While the algorithm favors the proposing side,³ Roth and Peranson showed that in the case of NRMP, the algorithm produces similar match results whether applicants or programs propose.⁹

The Gale–Shapley algorithm assumption of submitting full rank lists has practical implications, since this is not true in practice. Applicants cannot directly rank all programs because they must first apply to programs for interviews. While applicants can rank programs that did not interview them, programs will generally only rank applicants whom they have interviewed. However, under the current applicant-proposing version of Gale–Shapley, applicants cannot do worse by ranking and being ranked by more programs. This implies that applicants should apply to as many programs as resources allow in the hopes of being invited for more interviews and then being ranked more often. Similarly, programs likely feel induced to interview and rank as many applicants as possible to increase the likelihood of matching all positions in their program.

In the current study, we examine recent trends in the number of applicants and available positions, as well the average number of applications and interviews per applicant for ophthalmology and multiple NRMP specialties. To determine whether these numbers are insufficient, optimal, or excessive, we simulated matches using the Gale–Shapley algorithm, comparing present conditions with simulated matches in which the number of positions that applicants or programs could rank are limited.

Patients and Methods

This retrospective study was conducted in accordance with the Declaration of Helsinki. The study was exempted from approval by the institutional review board of the University of Washington, Seattle, WA. Publicly available historical data were collected from the NRMP and SF Match websites, and from archived versions of the websites using the Wayback Machine.¹⁰ Data were collected for matches between 1985 and 2020. In addition, fully anonymized rank lists and match data for ophthalmology applicants and programs were obtained from the SF Match for the years 2011 to 2019 with approval from the Board of Trustees of the Association

of University Professors of Ophthalmology, who oversees SF Match.

Longitudinal Trend Analysis

Using historical match statistics from the NRMP and SF Match, the total number of applicants and positions and the average number of applications and interviews over time were obtained for the following specialties: dermatology, otolaryngology, internal medicine, orthopedic surgery, plastic surgery, diagnostic radiology, radiation oncology, and ophthalmology.

To evaluate the trends in the ranking behaviors of the residency programs, we modeled the length of program rank lists (taken as a proxy and a lower bound for the number of interviews), and the number of available positions over time in years using a multivariable ordinary least squares model. The regression model was fitted by using the anonymized SF match rank list data.

We performed a cost and risk analysis of the ophthalmology residency match for applicants to determine the economics behind residency matching. Cost estimates were based on a financial analysis study of the ophthalmology residency match program.¹¹

Capping Analyses and Truncation Analysis

We investigated the extent of the universal excessive ranking that occurs, defined as ranking more programs/applicants than necessary to ensure a match and filling all available spots. We capped the length of the finalized rank lists of applicants using anonymized, actual SF Match rank lists as the basis for our experiment. We applied progressively more capping restrictions to limit the maximum number of entries on the rank lists. Next, to cap the programs, the number of applicants per spot was increasingly restricted to account for programs of different sizes. As a final analysis, we capped both applicant and program rank lists. The Gale–Shapley algorithm was then applied to the modified rank lists. The percent of all available ophthalmology positions filled was computed for each capping level.

To understand the pressures behind the universal excessive ranking behavior, we performed individual truncation experiments where the rank list of each applicant or program was successively truncated while all other rank lists were unchanged and the Gale–Shapley algorithm was rerun. For applicants, the change in rank status in going from matched to unmatched was measured. For programs, the percentage of spots filled was measured as a function of rank list truncation.

Results

Applicants and Ranked Applicants are Steadily Increasing Despite Minimal Changes in Number of Applicants per Position

The number of applicants relative to the number of available residency positions in ophthalmology has been steady at approximately 1.40 applicants per available position every year (95% confidence interval: 1.28–1.54) since 2000

Fig. 1 Longitudinal trends in ophthalmology match. (A) Total number of matched and unmatched applicants by year for SF Match. (B) Average number of applications and interviews by year for SF Match. Data for the average number of interviews were not available for the years in the gray box. (C) Comparison of the number of applications as a percentage of all programs in 2019 for ophthalmology and National Resident Matching Program specialties internal medicine, radiology, orthopedic surgery, otolaryngology (ear, nose, and throat), dermatology, radiation oncology, and plastic surgery. SF, San Francisco Residency and Fellowship Match Services.

(**Fig. 1A**). In contrast, the average number of applications submitted and the average number of interviews per applicant have been rising continuously in ophthalmology. The average number of applications per applicant submitted annually between 1985 and 2020 was increased from 24 to 77. Linear regression of these data since 2000 indicates an annual increase in applications of 2.07/year/applicant (applications = $2.07 + 32.44t$, $r^2 = 0.98$, $p = 5.7e-17$). Although data were not available for the time period of 2000 to 2010, looking at a longer timescale, the average number of interviews per applicant was increased 56% (5.7–8.9) between 1985 and 2020 (**Fig. 1B**), although this number appears to have stabilized over the past 5 years.

Similar trends were found for NRMP-matched specialties (**Supplementary Fig. S1** [available in the online version]). The median number of applications per candidate has increased from 27.6 to 39 (41.3%) between 2008 and 2019 across all NRMP specialties. For seven selective NRMP specialties, the median number of applications has increased 38.5%. Similarly, the median number of interviews has increased by 19.3 and 5.2% over this time across all NRMP

specialties and for the seven selected specialties, respectively (► **Supplementary Fig. S1** [available in the online version]). For comparison over the same time frame (2008–2019), the average number of applications in ophthalmology increased 56.3% (48–75). In addition, between 2011 and 2019, where the data were available, the average number of interviews in ophthalmology was increased by 7.5% (8.24–8.86).

The median number of applications as a ratio of all programs for seven selective NRMP-matched specialties and ophthalmology is shown in ► **Fig. 1C**. For specialties, such as radiation oncology and plastic surgery, the applicants typically apply to over 90% of all programs. In ophthalmology, applicants apply to 65% (75/116) of all programs on average.

In 2019, the average ophthalmology applicant ranked 8.86 ± 5.53 programs, and the average program ranked 11.54 ± 4.26 applicants per available position. On the program side, the length of the rank lists has increased over time. By linear regression modeling, programs ranked 8.34 candidates per open spot in 2011 and have been ranking 1.83 (95% CI: 1.53–2.14) more candidates per available position every subsequent year.

Moderate Capping of Match List Length for Applicants and Programs have Minimal Effect on Overall Match Success

It is possible that the current match list lengths have increased over time to ensure a complete match (i.e., to fill nearly every available position). To estimate the impact of shortened rank lists on overall match success for both applicants and programs, we re-simulated the match for each year from 2011 to 2019 while capping the maximum number of entries on either applicant or program rank lists by progressive degrees. The total percentage of positions filled for applicant, program, and both combined after rank list capping are shown in ► **Fig. 2A–C**, respectively. When we compared the number of applicants who matched without rank list limitations versus the number matched under successively shorter capped maximum rank list lengths, only capping applicant rank lists below three positions resulted in more than 5% change in this number (► **Fig. 2A**). Notably, capping program rank lists up to the limit of analysis of 15 positions did not affect the overall success of the match.

Individual Ranking Behavior of Applicants and Programs through Truncation Experiments

From these data, it appears that both programs and applicants are ranking more counterparts than are necessary for a stable match. To understand the pressures behind the over-ranking, we performed unilateral and individual truncation experiments, where we systematically removed the last entry of each applicant's rank list while not changing any other applicant or program rank lists and re-ran Gale–Shapley to observe changes in their match outcome. We then repeated this until each applicant's rank list was reduced to a single entry (► **Fig. 3**). We found that applicants who ranked up to 10 programs had a change in their match

outcome even with the removal of a single entry from their rank lists, and applicants who ranked 13 or more could remove the bottom two entries on their rank lists without a change in the match outcome. These results show that, at an individual level, applicants benefit from submitting long match lists.

We also analyzed the proportion of applicants matching after truncating the number of applicants ranked for each program while holding all other programs' and all applicant rank lists unchanged, stratified by the program size (► **Fig. 4**). A stepwise decrease was noted in the number of applicants ranked per spot with respect to the program size (► **Fig. 4A**). Smaller programs rank more applicants per spot compared with larger programs; a program with four positions available ranks a median of 11.6 applicants per spot, while a program with eight available positions ranks a median of 7.3 applicants per spot. When the rank lists of individual programs were truncated, a negative effect was seen for smaller programs earlier than larger programs (► **Fig. 4B**). For instance, two-person programs would lower their fill rate to 90% by truncating just three to four entries in their rank lists, while four-person programs can truncate 19 ranks to reach a similar rate. Thus, individual programs, particularly smaller programs, benefit from increasing the length of their match list.

Discussion

The results of the current analysis demonstrate that for both ophthalmology and other “competitive” specialties (i.e., where applicants significantly outnumber positions), (1) the number of applications per applicant and number of interviews per program have increased substantially over the past 20 years; (2) the current numbers of ophthalmology applications and interviews are in excess of those necessary to ensure a near-complete match; and (3) individual truncation of match list length by either applicant or program negatively impacts the likelihood of a successful match for the individual.

The increases seen in number of applications and interviews are driven by the Nash equilibrium.¹² A Nash equilibrium is when no player of a game can improve their payoff by changing their strategy while all the other players keep their strategies unchanged. In this “game,” the players are the applicants, who have realized that they will do no worse by applying to more programs because under the current applicant-proposing Gale–Shapley stable marriage residency algorithm, applicants will always match to their most preferred program if that program also prefers them over other candidates. Consequently, if applicants reduce their rank lists by even a single entry, they run the risk of becoming unmatched (► **Fig. 3**). The only deterrent to applying to more programs for applicants is increased cost.^{13–15} The average cost of submitting applications in ophthalmology has risen \$805 in less than 10 years from \$930 in 2011 to \$1,735 in 2019. It can rise another $\$35 \times (116 - 75) = \$1,435$ if the average applicant applies to all programs, the saturation point. The extra \$1,435 is 0.39% of the average

Fig. 2 Relationship between the percentage of total ophthalmology residency positions filled before capping at different maximum rank list lengths for (A) applicants, (B) programs, and (C) both applicants and programs. For each experiment, the Gale–Shapley algorithm was rerun to simulate the match.

ophthalmologist's annual income of \$366,000 in 2019, or 0.013% annualized over a 30-year career.¹⁶ Thus, the financial burdens of application (and indeed, interviewing as well) although substantial for the student are trivial with respect to the cost of not matching.

The Nash equilibrium of Gale–Shapley not only increases costs for applicants, but also for programs. In 2011, 621 ophthalmology applicants applied to an average of 53 programs; while in 2019, 648 applicants applied to an average of 75 programs each. During this time, the number of programs was increased from 113 to 116. Using an estimated 5-minute initial review time per application,¹⁷ a program director would spend on average 10 additional hours reviewing applications ($5 \times 648 \times 75/116/60 = 35.1$ hours in 2019 vs. $5 \times 621 \times 53/113/60 = 25.3$ hours in 2011). The application

review process will likely become more time consuming in the future, as the USMLE Step 1 exam, cited by 94% of program directors as an important factor in extending interview offers, will become pass/fail in 2022.¹⁸ Programs also experience substantial financial burden to interview candidates.^{15,19} Ophthalmology programs spend approximately \$3,736 per interviewed candidate when application screening time and lost clinical revenue for interview time are accounted for.¹¹ Ophthalmology interview costs are already large, as an average of 8.34 candidates are interviewed per available position, but costs will rise further with the average increasing by 1.83 additional candidates per position every year. This burden is particularly challenging for smaller programs, which must rank more applicants per position to ensure filling (►Fig. 4).

Fig. 3 Effect of the individual truncation of applicant rank list. Applicant outcomes (% matched) grouped by the number of programs the applicants ranked (rank list lengths are shown in gray), where each applicant's rank list is truncated by different amounts while no other applicant and program rank lists are modified.

The Nash equilibrium challenge could be addressed by an agreement among all programs and applicants to cap the number of interviews per available slot at 8 (→Fig. 2, bottom), and for applicants to rank only those programs at which they interviewed (i.e., no more than eight). For example, in 2019, adopting this policy would have resulted in a 29.6% reduction in total ranked positions by programs, and a 29.1% decrease in the number of programs ranked by applicants. Despite the reductions, approximately 95% of candidates would have still matched with a stable-marriage result. Overall, such an agreement would have reduced the number of total interviews system-wide in 2019 from 5,856 to 4,190 (28.5%). Given per interview cost of \$404 for candidates and \$3,736 for programs, this would have resulted in a net savings of \$6,897,240 for the system at a cost of 22 candidates needing to enter the scramble to secure a position. However, such a change would not mitigate screening of initial applications and might represent a restraint of choice for applicants. A “cascaded match” in which applicants first “match” to interviews (perhaps by remotely conducting preliminary interviews), with those interviews limited to eight programs,

might be a reasonable approach for implementing such a system.

Several limitations of the present study should be noted. This was an observational longitudinal study, and only certain specialties were examined. The consequences of Gale–Shapley algorithms might differ for specialties where available positions are in excess of qualified candidates, for example. However, the trend toward increased applications was remarkably similar across multiple competitive specialties. In addition, the analysis of applicant and program behavior based on rank list, and match information was limited to ophthalmology due to the availability of data. Nevertheless, we believe that the same concerns we highlight would equally apply to other specialties since the same match algorithm is used by NMRP. We also did not have access to which programs each applicant actually interviewed at and, as such, do not know the applicant-program pairs in which both sides decided not to rank each other. Finally, the study does not examine the effect of the strict ordinal rankings used in the Gale–Shapley algorithm on participant behavior. Ordinal rankings cannot express

Fig. 4 Effect of individual truncation of program rank list. (A) The number of applicants ranked per spot was grouped by program size. Smaller programs rank more applicants than larger programs. (B) Total percentage of positions matched grouped by the number of program spots, where every program truncated its rank list while the rank lists of all other programs and all applicants remain unchanged, with the match simulated under Gale–Shapley.

relative preference and may not be representative of either applicant or program preferences.²⁰ It is possible that weighted match list ranking (in which candidates and programs can “weight” their preferences in a nonlinear fashion) might show different behavior under truncation. Our capping and truncation analysis assumes that the final ranking behavior of the applicants and programs would not change after having been interviewed.

These challenges of the Nash equilibrium driving applicants toward applying to all programs nationally and inducing programs to interview increasing numbers of applicants are a direct consequence of the current structure of Gale–Shapley based match systems. Our analysis of rank list truncation demonstrates that rank lists are currently excessively long for a successful match. However, mandatory capping of rank list length is likely not feasible as it would be viewed as constraining choice. An alternative to capping the number of interviews would be to utilize non-Gale–Shapley algorithms, which might have different Nash equilibrium behavior. For instance, providing a budget of rank weightings to programs and candidates (as opposed to the current ordinal ranking) might intrinsically reduce application numbers while improving satisfaction with the match by allowing candidates to better express their preferences. Ideally, an improved algorithm would optimize the preferences of the entire match and better incorporate the relative preferences of participants, while achieving major cost savings for all participants.

Note

A.Y.L. has received honoraria from Topcon, Genentech, and Verana Health. A.Y.L. works with the U.S. Food and Drug Administration.

Funding

A.Y.E. reports grants from National Eye Institute, Novartis, Regeneron Santen, Carl Zeiss Meditec, Microsoft, and NVIDIA. He reported personal fees from U.S. Food and Drug Administration, Genentech, Topcon, and Verana Health outside the submitted work. was supported by the Research to Prevent Blindness, University of Washington CoMotion Innovation Fund (NEI/NIH K23EY029246), and supported in part by the Machado-Joseph disease Research Fund. C.S.L. reported grants from National Institute on Aging outside the submitted work.

Conflict of Interest

None declared.

Acknowledgments

The authors would like to thank Tim Losch and Dennis Thomatos from SF Match for his assistance in preparing, packaging, and sharing the SF Match data for this manuscript. R.N.V.G. and colleagues have applied for a patent for novel algorithms for matching. Y.U. has a patent optimization framework for two-sided markets pending.

References

- 1 SF Match homepage. SF Match. Accessed 2021 at: <https://sfmatch.org/>
- 2 2020Main Residency Match®. NRMP. Accessed 2021 at: www.nrmp.org
- 3 Gale D, Shapley LS. College admissions and the stability of marriage. *Am Math Mon* 1962;69(01):9–15
- 4 Main Residency Match Data and Reports. Accessed 2021 at: <https://www.nrmp.org/main-residency-match-data/>

- 5 Roth AE. The evolution of the labor market for medical interns and residents: a case study in game theory. *J Polit Econ* 1984;92(06):991–1016
- 6 Turner J. Intern selection: wanted, an orderly plan. *Acad Med* 1945;20(01):26
- 7 Mullin FJ, Stalnaker JM. The matching plan for internship appointment. *J Med Educ* 1951;26(05):341–346
- 8 Roth AE. The origins, history, and design of the resident match. *JAMA* 2003;289(07):909–912
- 9 Roth AE, Peranson E. The redesign of the matching market for american physicians: some engineering aspects of economic design. *Am Econ Rev* 1999;89(04):748–780
- 10 Internet archive: digital library of free & borrowable books, movies, music & wayback machine. Accessed 2021 at: <https://archive.org/>
- 11 Moore DB. Not a cheap investment: estimating the cost of the 2017 to 2018 ophthalmology residency match to the applicant and program. *J Acad Ophthalmol* 2018;10(01):e158–e162
- 12 Nash J. Non-cooperative games. *Ann Math* 1951;54(02):286–295
- 13 Nikonow TN, Lyon TD, Jackman SV, Averch TD. Survey of applicant experience and cost in the urology match: opportunities for reform. *J Urol* 2015;194(04):1063–1067
- 14 Polacco MA, Lally J, Walls A, Harrold LR, Malekzadeh S, Chen EY. Digging into debt: the financial burden associated with the otolaryngology match. *Otolaryngol Head Neck Surg* 2017;156(06):1091–1096
- 15 Van Dermark JT, Wald DA, Corker JR, Reid DG. Financial implications of the emergency medicine interview process. *AEM Educ Train* 2017;1(01):60–69
- 16 Kane L. Medscape physician compensation report 2019. Medscape Published April 10, 2019. Accessed 2021 at: <https://www.medscape.com/slideshow/2019-compensation-overview-6011286>
- 17 Weissbart SJ, Hall SJ, Fultz BR, Stock JA. The urology match as a prisoner's dilemma: a game theory perspective. *Urology* 2013;82(04):791–797
- 18 The United States Medical Licensing Examination. Accessed 2021 at: <https://www.usmle.org/incus/>
- 19 Susarla SM, Swanson EW, Slezak S, Lifchez SD, Redett RJ. The perception and costs of the interview process for plastic surgery residency programs: can the process be streamlined? *Plast Reconstr Surg* 2017;139(01):302e–309e
- 20 Benjamin DJ, Heffetz O, Kimball MS, Rees-Jones A. Can marginal rates of substitution be inferred from happiness data? Evidence from residency choices. *Am Econ Rev* 2014;104(11):3498–3528