

An Empirical Model of the Medical Match[†]

By NIKHIL AGARWAL*

This paper develops a framework for estimating preferences in a many-to-one matching market using only observed matches. I use pairwise stability and a vertical preference restriction on one side to identify preferences on both sides of the market. Counterfactual simulations are used to analyze the antitrust allegation that the centralized medical residency match is responsible for salary depression. Due to residents' willingness to pay for desirable programs and capacity constraints, salaries in any competitive equilibrium would remain, on average, at least \$23,000 below the marginal product of labor. Therefore, the match is not the likely cause of low salaries. (JEL C78, I11, J31, J44, K21, L44)

Each year, the placement of about 25,000 medical residents and fellows is determined via a centralized clearinghouse known as National Residency Matching Program (NRMP) or “the match.” During the match, a stable matching algorithm uses rank order lists reported by residents and programs to assign applicants to positions. Agents on both sides of the market are heterogeneous but salaries paid by residency programs are not individually negotiated with residents. Therefore, preferences of residents and programs, rather than prices alone, determine equilibrium assignments. The medical match is iconic in the stable matching literature, but with few exceptions this literature has been primarily theoretical.

This paper estimates the preference distribution in the market for family medicine residents in the United States, and empirically analyzes the antitrust allegation that the centralized market structure is responsible for the low salaries paid to residents. The plaintiffs in a 2002 lawsuit argued that the match limited the bargaining power of the residents because salaries are set before ranks are submitted. They reasoned that a “traditional market” would allow residents to use multiple offers and salary bargaining to make programs bid for their labor. Using a perfect competition

*Department of Economics, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02142, and NBER (e-mail: agarwaln@mit.edu). I am grateful to my advisors Susan Athey, Ariel Pakes, Parag Pathak, and Al Roth for their constant support and guidance. I thank Atila Abdulkadiroglu, Raj Chetty, Rebecca Diamond, William Diamond, Adam Guren, Guido Imbens, Joel Katz, Larry Katz, Greg Lewis, Julie Mortimer, seminar participants at numerous universities, and four anonymous referees for helpful discussions, suggestions, and comments. Data acquisition support from LEAP and the Kuznets Award, financial support from the NBER Nonprofit Fellowship, and the generous hospitality of the Cowles Foundation is gratefully acknowledged. Data on matches from the Graduate Medical Education Database, 2012, American Medical Association, Chicago, IL. The author obtained IRB approval. The author declares that he has no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <http://dx.doi.org/10.1257/aer.20131006> to visit the article page for additional materials and author disclosure statement.

model as the alternative, they argued that the large salary gap between residents and nurse practitioners or physician assistants is a symptom of competitive restraints imposed by centralization. Although the lawsuit was dismissed due to a legislated Congressional exception, it sparked an academic debate on whether inflexibility results in low salaries (Bulow and Levin 2006; Kojima 2007). Observational studies of medical fellowship markets do not find an association between low salaries and the presence of a centralized match (Niederle and Roth 2003, 2009). While these studies strongly suggest that the match is not the primary cause of low salaries in this market, they do not explain why salaries in decentralized markets may remain lower than the perfect competition salary benchmark suggested by the plaintiffs. I use a stylized theoretical model to show that residents' willingness to pay for training at desirable programs and capacity constraints result in an "implicit tuition" that depresses salaries in a decentralized market.

To quantify the implicit tuition, the paper develops a framework for estimating preferences (market primitives) of both sides of a many-to-one two-sided matching market using data only on final matches from a large market. These primitives are important inputs into the counterfactual analysis of government interventions or outcomes under alternative market designs. However, direct data on these market primitives is frequently not available due to confidentiality concerns as in the case of the NRMP, or because they are not collected in the first place. When only data on final matches are available, it is not immediately clear how to use these data to estimate the distribution of preferences. The approach may therefore be useful for studying other matching markets as well.

The empirical techniques in this paper require data from a large many-to-one two-sided matching market with low frictions. The primary assumptions are that one side of the market has homogeneous preferences (while the other side may have heterogeneous preferences) and observed matches are described by a pairwise stable equilibrium. According to this equilibrium concept, no two agents on opposite sides of the market prefer each other over their match partners at predetermined salary levels. In this empirical context, the equilibrium assumption can be justified by properties of the medical match (Roth and Peranson 1999). Following the discrete choice literature, I model the preferences of each side of the market over the other as a function of characteristics of residents and programs, some of which are not known to the econometrician. I use the pure characteristics model of Berry and Pakes (2007) for the preferences of residents for programs. This model allows for substantial heterogeneity in preferences. However, a similarly flexible model for a program's preferences for residents raises identification issues and other methodological difficulties. In the medical residency market, anecdotal evidence suggests that residents are largely vertically differentiated in skill because medical school quality, academic record, and clinical performance are the main determinants of a resident's desirability to a program. To account for the primary unobservables (academic record and recommendation letters) and difficulties with generalizing the model, I restrict programs' preferences for residents to be homogeneous while allowing for an unobservable determinant of resident skill. The assumption implies the existence of a unique pairwise stable match and a computationally tractable simulation algorithm. While restrictive, vertical preferences may be a good modeling approximation in some other contexts. For example, colleges in many countries use a single aggregate high school exam score

to determine admission eligibility. Workers differentiated by an ordinal human capital index may also reasonably approximate firm preferences in some other narrowly defined, high-skilled labor markets.

The empirical strategy must confront the fact that “choice sets” of agents in the market are not observed because they depend on the preferences of other agents in the market. Instead of a standard revealed preference approach, I identify the model using the observed assortativity between resident and program characteristics (which I will refer to as “sorting patterns”) and information only available in an environment with many-to-one matching. Agarwal and Diamond (2014) formally studies nonparametric identification of a model with homogeneous preferences on both sides and shows that it is essential to use information in many-to-one matches. Intuitively, the degree of assortativity between, say, medical school prestige and hospital size increases with the importance of these characteristics in agent preferences. For this reason, a high weight on medical school prestige and a low weight on hospital size results in a similar degree of assortative matching as a high weight on hospital size and low weight on medical school prestige. Fortunately, data from many-to-one matches has additional information that assists in identification. In a pairwise stable match, all residents at a given program must have similar human capital. Otherwise, the program can likely replace the least skilled resident with a better resident. Because the variation in human capital within a program is low, the within-program variation in medical school prestige decreases with the correlation of human capital with medical school prestige. Note that it is only possible to calculate the within-program variation in a resident characteristic if many residents are matched to the same program. Finally, to learn about heterogeneity in preferences, I use observable characteristics of one side of the market that are excluded from the preferences of the other side. These exclusion restrictions shift the preferences of, say, residents, without affecting the preferences of programs, thereby allowing sorting on excluded characteristics to be interpreted in terms of preferences.

I estimate the model using a simulated minimum distance estimator (McFadden 1989; Pakes and Pollard 1989; Gouriéroux and Monfort 1997) that matches moments in the data to those predicted by the model, and data from the market for family medicine residents between 2003 and 2010. Approximately 430 programs and 3,000 medical residents participate in this market each year. Moments used in estimation include summaries of the sorting patterns observed in the data and the within-program variation in observable characteristics of the residents.

Since I will be estimating the effect of salaries on resident choices, I show how to correct for potential endogeneity between salaries and unobserved program characteristics. The technique is based on a control function approach and relies on the availability of an instrument that is excludable from the preferences of the residents (see Heckman and Robb 1985; Blundell and Powell 2003; Imbens and Newey 2009). This approach can be useful in other applications in education or labor markets where endogeneity may arise due correlation of financial aid or salaries with unobserved characteristics. For this setting, I construct an instrument using Medicare’s reimbursement rates to competitor residency programs, which are based on regulations enacted in 1985. The results from the instrumented version of the model are imprecise but indicate that salaries are likely positively correlated with unobserved program quality.

I assess the fit of the model, both in-sample and out-of-sample. The out-of-sample fit uses the most recent match results, taken from the 2011–2012 market. These data were not accessed until estimates were obtained. The observed sorting patterns for resident groups mimic those predicted by the model, both in-sample and out-of sample, suggesting that the model is appropriate for counterfactuals. Further, Monte Carlo simulations suggest that parameter estimates important for the counterfactuals in this paper are not very sensitive to moderate levels of misspecification in preferences and the equilibrium assumptions.

I use these estimates to study the antitrust allegation against the medical match. In the lawsuit, the plaintiffs used a perfect competition model to argue that residents' salaries are lower than those paid to substitute health professionals because the match eliminates salary bargaining. This reasoning does not account for the effects of the limited supply of each type of program and resident. Capacity constraints at desirable residency programs due to accreditation requirements may lower salaries at these programs. I use a stylized model to show that when residents value program quality, salaries in every competitive equilibrium are below the benchmark level suggested by the plaintiffs. The markdown is due to an implicit tuition arising from residents' willingness to pay for training at a program, and is in addition to any costs of training passed through to the residents. I estimate an average implicit tuition of at least \$23,000, with larger implicit tuitions at more desirable programs. Although imprecisely estimated, models using salary instruments estimate an average implicit tuition that is much higher: \$43,000 or more. The results weigh against the plaintiffs' claim that in the absence of the match, salaries paid to residents would be equal to the marginal product of their labor, close to salaries of physician assistants and nurse practitioners. At a median salary of \$86,000, physician assistants earn approximately \$40,000 more than medical residents. The upper-end of the estimated implicit tuition can explain this difference. The results imply that the implicit tuition would result in low salaries even in the absence of a match.

Related Literature.—The empirical methods in this paper contribute to the recent literature on estimating preference models using data from observed matches and pairwise stability in decentralized markets (for a survey, see Fox 2009). The majority of papers focus on estimating a single aggregate surplus that is divided between match partners. Chiappori, Salanié, and Weiss (2015), and Galichon and Salanié (2015), among others, build on the seminal work of Choo and Siow (2006) for studying transferable utility models of the marriage market in which agents have unobserved and heterogeneous preferences for partners of types known to the econometrician. Fox (2010) proposes a different approach for estimation, also for the transferable utility case, with applications in Fox and Bajari (2013), Akkus, Cookson, and Hortacsu (2015), among others. This approach is based on assuming that the structural unobservables are such that the observed matches are more likely than alternatives to maximize the systematic component of the total surplus. Sorensen (2007) is an example that estimates a single surplus function, but in a non-transferable utility model. Another set of papers measures benefits of mergers using related cooperative solution concepts (Gordon and Knight 2009; Uetake and Watanabe 2012; Weese forthcoming).

This paper estimates the distribution of preferences of each of the two sides of the market in a non-transferable utility model, with salary as an (potentially endogenous) additional characteristic that is valued by residents. Separately estimating the resident preference distribution is necessary for calculating the implicit tuition. From a methodological perspective, a primary distinction with the literature on transferable utility matching is in the treatment of transfers in the model. While this paper also allows for utility to be transferable through salaries, it models endogeneity through econometric correlation with unobserved characteristics rather than an equilibrium outcome simultaneously determined with the match. This approach may be better suited for environments where transfers are determined outside the matching process as in the medical residency market, public sector labor markets with salary schedules, or some higher education markets.¹

Similar models, but without endogenous transfers, are estimated by Logan, Hoff, and Newton (2008) and Boyd et al. (2013) in decentralized markets with the goal of measuring preferences for various characteristics. Logan, Hoff, and Newton (2008) proposes a Bayesian method for estimating preference distributions in a marriage market with no monetary transfers. Boyd et al. (2013) uses the method of simulated moments to estimate the preference distribution of teachers for schools and of schools for teachers, arguing identification of the model by analogy to the discrete choice literature. An important methodological difference is that these previous papers use only aggregate sorting patterns in the data in their empirical approach. Two recent papers formally show that information available in one-to-one matches or aggregate sorting patterns are not sufficient to identify the preference distribution on both sides of the market. Agarwal and Diamond (2014) prove that even under a very restrictive model with no preference heterogeneity on either side of the market, sorting patterns alone cannot fully identify the preference parameters of the model. Under parametric assumptions on the preference distribution but in a model that allows for heterogeneity, Menzel (forthcoming) shows that only the sum of the individual surplus is identified from data on one-to-one matches. Non-identification of this type can yield unreliable predictions for the counterfactual studied in this paper. To solve this problem, this paper leverages information made available through many-to-one matches, in addition to sorting patterns, for identifying two distributions of preferences. Agarwal and Diamond (2014) formally show nonparametric identification of their model using data from two-to-one matches. This paper intuitively discusses the importance of using information in many-to-one matches, and discusses exclusion restrictions useful for identifying a model with heterogeneous preferences.

The results on salary depression may also be of independent interest for its analysis of labor markets with compensating differentials, especially those with on-the-job training. Previous theoretical work on markets with on-the-job training has used perfect competition models to show that salaries are reduced by the marginal cost of training (Rosen 1972; Becker 1975). Counterfactuals in this paper using the competitive equilibrium model compute an implicit tuition, a markdown due to the value of training that is in addition to costs of training passed through to the resident.

¹For instance, the approach allows for college financial aid that is determined by income and demographic criteria as long as awards are not made in consideration of offers made to the student by other colleges.

Overview.—Section I describes the market for family medicine residents and the sorting patterns observed in the data. Sections II–VI present the empirical framework used to analyze this market, the identification strategy, the method for correcting potential endogeneity in salaries, the estimation approach, and parameter estimates, respectively. These sections omit details relevant exclusively to the application, which is discussed in Section VII. All technical details are relegated to online appendices which follow the organization of the paper.

To better understand the framework in the context of the application, I briefly outline the approach in Section VII. The theoretical analysis uses a stylized model in which residents and programs are vertically differentiated to show that residents' willingness to pay for desirable programs results in a markdown in salaries (the implicit tuition) in any competitive equilibrium (core allocations in Shapley and Shubik 1971). I show that the implicit tuition can be computed even when residents have heterogeneous preferences using estimates only of the distribution of residents' willingness to pay for different programs. This calculation requires a production function in which residents are vertically differentiated from a program's perspective. This restriction circumvents the need to monetize the value of a resident to a program in dollars.

I. Market Description and Data

The data on the family medicine residency market are taken from the 2003–2004 to 2010–2011 waves of the National Graduate Medical Education census (GME census) which provides characteristics of residents linked with information about the program at which they are training.² Family medicine is the second largest specialty (after internal medicine), constituting about one-eighth of all residents in the match.

I focus on five major types of program characteristics: the prestige/quality of the program as measured by National Institutes of Health (NIH) funding secured by a program's major and minor medical school affiliates; the size of the primary clinical hospital as measured by the number of beds; the Medicare case mix index as a measure of the diagnostic mix a resident is exposed to; location characteristics such as the median rent in the county a program is located in and the Medicare wage index as a measure of local health care labor costs; and the program type indicating the community, university, and/or rural setting of a program.

Panel A of Table 1 summarizes the characteristics of programs in the market. The market has approximately 430 programs, each offering approximately 8 first-year positions. Although not reported in the table, in all but 1.5 percent of the 3,441 cases across program-years, a program has more than one match. Except for program type (community- and/or university-based), there is little annual variation in the composition of programs in the market. Salaries paid to residents track inflation with

²I consider all non-military programs participating in the match, accredited by the Accreditation Council of Graduate Medical Education (ACGME) and not located in Puerto Rico. I restrict attention to residents matched with these programs. Detailed description of data sources and variable construction are in online Appendix A.

TABLE 1—PROGRAM AND RESIDENT CHARACTERISTICS

	2010–2011		2002–2003 to 2010–2011	
	Mean	SD	Mean	SD
<i>Panel A. Programs</i>				
Number of programs	428		3,441	
First year salary (2010 dollars)	\$47,331	\$2,953	\$46,394	\$3,239
NIH funding (major affiliates, millions \$)	88.37	92.26	88.35	91.62
NIH funding (minor affiliates, millions \$)	89.90	79.85	99.04	92.46
Beds (primary institution)	421.54	284.15	418.41	273.17
Medicare case mix index (primary institution)	1.61	0.23	1.57	0.22
Medicare wage index (primary institution)	1.00	0.14	1.01	0.14
Community-based program	0.25	0.43	0.33	0.47
Community-university program	0.62	0.49	0.54	0.50
University-based program	0.13	0.34	0.12	0.33
Rural program	0.15	0.35	0.14	0.35
Number of positions	7.70	2.83	7.57	2.77
Number of matches	7.36	2.93	7.01	2.92
Number of interviews	63.38	31.10	55.56	30.17
<i>Panel B. Residents</i>				
Number of residents	3,148		24,115	
Allopathic/MD graduate	0.45	0.50	0.45	0.50
Osteopathic/DO graduate	0.15	0.36	0.14	0.34
Foreign medical graduate	0.39	0.49	0.41	0.49
NIH funding (MD graduates, millions \$)	83.26	82.42	84.08	83.96
Median MCAT score (MD graduates)	31.24	2.25	31.31	2.20
Female	0.56	0.50	0.55	0.50
US-born foreign graduate	0.12	0.33	0.09	0.29
Rural-born resident	0.11	0.31	0.10	0.30

Notes: Variable construction detailed in the online Appendix. Interviews and Medicare fields reported conditional on nonmissing data (less than 2 percent missing). NIH funding only for programs with funded affiliates. About 35 percent of the programs have no NIH funded major affiliates, while about 46 percent have no minor affiliates. About 8 percent of programs have no NIH funded medical school affiliates. A resident is considered rural born if her city of birth is not in an MSA (missing for about 7.3 percent of residents). Residents with missing country of birth (14.6 percent) are treated as not born in the United States.

a mean of about \$47,000 (in 2010 dollars) and a standard deviation of about \$3,000.³ Although not reported in the table, the salaries range from \$31,000–\$65,000.

The data contain information on a resident's medical degree type, characteristics of graduating medical school and city of birth. Panel B in Table 1 describes the characteristics of residents matching with family medicine programs. The composition of this side of the market has also been stable over this sample period with only minor annual changes. A little less than one-half of the residents in family medicine are graduates of MD-granting medical schools in the US. A large fraction, about 40 percent, of residents obtained medical degrees from non-US schools while the

³Resident salaries after the first year are highly correlated with the first year salary with a coefficient that is close to 1 and a R^2 of 0.8 or higher.

rest have US osteopathic (DO) degrees.⁴ One in ten US-born medical residents are born in rural counties.

A. *The Match*

A prospective medical resident begins her search for a position by gathering information about the academic curriculum and terms of employment (including salaries) at programs from various sources, such as online directories and official publications. The publications list the number of positions being offered by the program as well, which is regulated by the accreditation body (ACGME). Subsequently, she electronically submits applications to several residency programs which then select a subset of applicants to interview. On average, approximately eight residents are interviewed per position (panel A of Table 1). Anecdotal evidence suggests that during or after interviews, informal communication channels actively operate allowing agents on both sides of the market to gather more information about preferences. Finally, residency programs and applicants submit lists stating their preferences for their match partners. Programs do not individually negotiate salaries with residents during this process. The algorithm described in Roth and Peranson (1999) uses these rank order lists to determine the final match. The terms of participating in the match commit both the applicant and the program to honor this assignment. The algorithm itself substantially reduces incentives for residents and programs to rematch by producing a match in which no applicant and program pair could have ranked each other higher than their assignments. I refer the reader to Roth (1984); Roth and Xing (1994); and Roth and Peranson (1999) for a historical perspective on the evolution of this market.

A few positions are filled before the match begins and some positions not filled after the main match are offered in the “scramble.” During the scramble, residents and programs are informed if they were not matched in the main process and can use a list of unmatched agents to contract with each other.⁵

B. *Descriptive Evidence on Sorting*

The empirical strategy uses sorting patterns between resident and program characteristics observed in the data and features of the many-to-one matching structure to infer preference parameters. I defer a discussion of summary statistics based on many-to-one matches to Section IIIB.

There is a significant degree of positive assortative matching between measures of a resident’s medical school quality and that of a program’s medical school and hospital affiliates. Table 2 presents regressions of a resident’s characteristic on the characteristics of programs with which she is matched. Programs that are associated with better NIH-funded medical schools tend to match with residents from

⁴Doctors of allopathic and osteopathic medicine have similar licenses to write prescriptions and treat patients but are trained in differing treatment approaches.

⁵A total of 142 positions in family medicine (approximately 5 percent) were filled through a different supplementary process in 2012. The scramble was likely of a similar size in the earlier years. See Signer (2012).

TABLE 2—SORTING BETWEEN RESIDENTS AND PROGRAMS

	log NIH fund (MD) (1)	Median MCAT (MD) (2)	MD degree (3)	DO degree (4)
log NIH fund (major)	0.3724*** (0.0119)	0.0154*** (0.0007)	0.0462*** (0.0032)	0.0025 (0.0022)
log NIH fund (minor)	0.1498*** (0.0137)	0.0084*** (0.0008)	0.0208*** (0.0040)	0.0048* (0.0028)
log number beds	-0.0972*** (0.0221)	-0.0021 (0.0014)	-0.0104 (0.0064)	-0.0098** (0.0045)
Rural program	-0.0687 (0.0437)	-0.0040 (0.0027)	-0.0010 (0.0117)	0.0138* (0.0082)
log case-mix index	0.1894** (0.0940)	0.0136** (0.0058)	0.4670*** (0.0255)	0.0574*** (0.0179)
log first-year salary	0.0126 (0.1717)	0.0590*** (0.0106)	0.3001*** (0.0467)	0.0969*** (0.0327)
log rent	0.4612*** (0.0600)	0.0727*** (0.0037)	0.1811*** (0.0168)	-0.0012 (0.0118)
Observations	10,842	10,872	23,984	23,984
R ²	0.1318	0.1282	0.0381	0.0079

Notes: Linear regression of resident characteristic on matched program characteristics. Column 1 restricts to the residents from medical schools with positive NIH funding. Column 2 restricts to the residents institutions reporting MCAT scores in the medical school admission requirements in 2010–2011. All specifications include indicators for programs with no NIH funding at major affiliates, no NIH funding at minor affiliates, and a missing Medicare ID for the primary institution. Standard errors in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

better medical schools as well, whether the quality of a resident's medical school is measured by NIH funding, MCAT scores of matriculants, or the degree type conferred by the school. This observation also holds true for programs at hospitals with a higher Medicare case mix index. Rent is positively associated with resident quality, potentially because cities with high rent may also be the ones that are more desirable to train or live in.

To highlight the geographical sorting observed in the data, Table 3 regresses characteristics of a resident's matched program on her own characteristics and indicators of whether the program is in her state of birth or medical school state. Residents that match with programs in the same state as their medical school tend to match with less prestigious programs, as measured by the NIH funds of a program's affiliates. These residents also match with programs that are at larger hospitals and have lower case mix indices. Column 5 shows that rural-born residents are about 7 percentage points more likely to place at rural programs than their urban-born counterparts.

Since these patterns arise from the mutual choices of residents and programs, estimates from these regressions are not readily interpretable in terms of the preference parameters of either side of the market. The next section develops a model of the market that uses these patterns and other features of the data to estimate preference parameters.

TABLE 3—GEOGRAPHICAL SORTING BETWEEN RESIDENTS AND PROGRAMS

	log NIH fund (major) (1)	log NIH fund (minor) (2)	log number beds (3)	log case mix index (4)	Rural program (5)
log NIH fund (MD)	0.4058*** (0.0124)	0.1555*** (0.0116)	-0.0213*** (0.0046)	-0.0002 (0.0011)	-0.0110*** (0.0023)
log median MCAT (MD)	0.6953*** (0.1009)	0.4704*** (0.0914)	0.0830** (0.0364)	0.0023 (0.0091)	-0.0877*** (0.0184)
US born (for)	-0.0711* (0.0374)	-0.1032*** (0.0366)	-0.0025 (0.0143)	0.0186*** (0.0036)	0.0141* (0.0072)
Match in medical school state	-0.4463*** (0.0322)	-0.2646*** (0.0303)	0.0468*** (0.0121)	-0.0057* (0.0030)	0.0111* (0.0061)
Match in birth state	-0.0038 (0.0285)	0.0197 (0.0264)	-0.0376*** (0.0105)	-0.0075*** (0.0026)	-0.0115** (0.0053)
Rural-born resident					0.0714*** (0.0066)
Observations	15,394	13,099	24,115	23,652	24,115
R^2	0.1211	0.0299	0.0052	0.0167	0.0101

Notes: Linear regression of program characteristics on characteristics of matched residents. Column 1 restricts to programs with major affiliates that have positive NIH funding. Column 2 restricts to programs with a minor affiliate with nonzero NIH funding. Column 4 excludes programs with missing Medicare ID. All specifications include medical school type dummies and a dummy for residents graduating from MD medical schools without NIH funding. Column 5 includes a dummy for nonreliable city of birth information for US-born residents. Standard errors in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

II. A Framework for Analyzing Matching Markets

This section presents the empirical framework, treating salaries as exogenous. I demonstrate how an instrument can be used to correct for correlation between salaries and unobserved program characteristics in Section IV.

A. Pairwise Stability

I assume that the observed matches are pairwise stable with respect to the true preferences of the agents, represented by \succeq_k for a program or resident indexed by k . Each market, indexed by t , is composed of N_t residents, $i \in \mathcal{N}_t$, and J_t programs, $j \in \mathcal{J}_t$. The data consist of the number of positions offered by program j in each period, denoted c_{jt} , and a match, given by the function $\mu_t : \mathcal{N}_t \rightarrow \mathcal{J}_t \cup \{0\}$, where 0 denotes being unmatched. Let $\mu_t^{-1}(j)$ denote the set of residents matched with program j .

A pairwise stable match satisfies two properties for all agents i and j participating in market t :

(i) Individual Rationality

- For residents: $\mu_t(i) \succeq_i 0$.
- For programs: $|\mu_t^{-1}(j)| \leq c_{jt}$ and $\mu_t^{-1}(j) \succeq_j \mu_t^{-1}(j) \setminus \{i\}$ for all $i \in \mu_t^{-1}(j)$.

(ii) No Blocking: if $j \succ_i \mu_r(i)$ then

- For all $i' \in \mu_r^{-1}(j)$, $\mu_r^{-1}(j) \succeq_j (\mu_r^{-1}(j) \setminus \{i'\}) \cup \{i\}$.
- Further, if $|\mu_r^{-1}(j)| < c_j$, then $\mu_r^{-1}(j) \succeq_j \mu_r^{-1}(j) \cup \{i\}$.

A pairwise stable need not exist in general or there may be multiple pairwise stable matches. The preference model described in the subsequent sections guarantees the existence and uniqueness of a pairwise stable match.

Individual rationality, also known as acceptability, implies that no program or resident would prefer to unilaterally break a match contract. Because I do not observe data on unmatched residents, I assume that no program prefers keeping a position empty to filling it with a resident in the sample, and that all residents prefer being matched to being unmatched. Almost all US graduates applying to family medicine residencies as their primary choice are successful in matching to a family medicine program, and the number of unfilled positions in residency programs in this specialty is under 10 percent.⁶ The primary limitation of this assumption is the inability to consider preferences for agents not matched in the family medicine market, such as programs or residents in other specialties.

Under the no-blocking condition, no resident prefers a program (to her current match) that would prefer hiring that resident in place of a currently matched resident if the program has exhausted its capacity. If the program a resident prefers has unfilled positions, the program would not prefer filling a position with that resident.

Theoretical properties of the mechanism used by the NRMP guarantee that the final match is pairwise stable with respect to submitted rank order lists, but not necessarily with respect to true preferences. Strategic ranking and interviewing, especially in the presence of incomplete information, is likely the primary threat to using pairwise stability in this market.⁷ The large number of interviews per position suggests that this may not be of concern in this market, however, it may be implausible in some decentralized markets. Online Appendix B1 presents Monte Carlo evidence that suggests that the estimation procedure is not particularly sensitive to moderate deviations from pairwise stability.

This equilibrium concept also implicitly assumes that agents' preferences over matches is determined only by their match, not by the match of other agents. This restriction rules out the explicit consideration of couples that participate in the match by listing joint preferences.⁸ According to data reports from the NRMP, in recent years, only about 1,600 out of 30,000 individuals participated in the main residency

⁶While residents may apply to many specialties in principle, data from the NRMP suggests that a typical applicant applies to only one or two specialties. Greater than 95 percent of MD graduates interested in family medicine as their first choice, however, apply only to family medicine programs. Upward of 97 percent of residents who list a family medicine program as their first choice match to a family medicine program in the main match (See "Charting Outcomes in the Match," NRMP 2006, 2007, 2009, 2011).

⁷The data and approach do not make a distinction for positions offered outside the match or during the scramble. The no-blocking condition should be a reasonable approximation for the positions filled before the match as it is not incentive compatible for the agents to agree to such arrangements if either side expects a better outcome after the match. For the small number of residents and programs matched during the scramble, note that the residents (programs) should not form blocking pairs with the set of programs (residents) that they ranked in the main round.

⁸Couples can pose a threat to the existence of stable matches (Roth 1984) although results in Kojima, Pathak, and Roth (2013) suggest that stable matches exist in large markets if the fraction of couples is small.

match as part of a couple. I model all agents as single agents because data from the GME census does not identify an individual as part of a couple.

B. Preferences of the Residents

Following the discrete choice literature, I model the latent indirect utility representing residents' preferences \succeq_i as a function $U(\mathbf{z}_{jt}, \xi_{jt}, w_{jt}, \beta_i; \theta)$ of observed program traits \mathbf{z}_{jt} , the program's salary offer w_{jt} , unobserved trait ξ_{jt} , and taste parameters β_i .

I use the pure characteristics demand model of Berry and Pakes (2007) for this indirect utility:

$$(1) \quad u_{ijt} = \mathbf{z}_{jt} \beta_i^z + w_{jt} \beta_i^w + \xi_{jt}.$$

In models that do not use a salary instrument, I assume that the unobserved trait ξ_{jt} has a standard normal distribution that is independent of the other variables. I normalize the mean utility to zero for $(\mathbf{z}, w) = 0$. The scale and location normalizations are without loss of generality. The independence of ξ_{jt} from w_{jt} is relaxed in the model correcting for potential endogeneity in salaries. Depending on the flexibility desired, β_i can be modeled as a constant, a function of observable characteristics x_i of a resident and/or of unobserved taste determinants η_i :

$$(2) \quad \beta_i^z = \mathbf{x}_i \mathbf{\Pi} + \eta_i.$$

The taste parameters η_i are drawn from a mean-zero normal distribution with a variance that is estimated. The richest specification used in this paper allows for heterogeneity via normally distributed random coefficients for NIH funding at major affiliates, beds, and case mix index. This specification also allows for preference heterogeneity for rural programs based on rural or urban birth location of the resident and heterogeneity in preference for programs in the resident's birth state or medical school state through interaction of x_i and z_{jt} . These terms are included to account for the geographic sorting observed in the market. I do not include random coefficients on salaries because of the limited variation in the data, thereby restricting $\beta_i^w = \beta^w$ for all i .

The pure characteristics model is micro-founded on residents having tastes for a finite set of program attributes while allowing for both unobserved heterogeneity in tastes and unobserved characteristics. It omits a commonly used additive ϵ_{ijt} term that is independent and identically distributed across residents, programs, and markets. I do not include both ϵ_{ijt} and ξ_{jt} because it is unclear how their distributions (variances) can be separately identified since aggregate demand used to invert for product-level unobservables in consumer models (Berry 1994; Berry, Gandhi, and Haile 2013) is not observed in our dataset. Given a choice between including ϵ_{ijt} and ξ_{jt} , I prefer an approach that includes ξ_{jt} for several reasons. First, product-level unobservables ξ_{jt} are important for building in endogeneity of salaries into the model (Berry and Haile 2014). Second, a common motivation for including ϵ_{ijt} has been that it ensures that no choice is dominated for all agents which is appealing in consumer markets since dominated choices are likely to exit. This motivation does not

apply to matching markets since dominated choices can prevail in matching markets due to capacity constraints faced by more desirable agents. Third, discrete choice models employing ϵ_{ijt} implicitly assume tastes for programs through a characteristic space that increases in dimension with the number of programs and residents (Berry and Pakes 2007), which are large in this empirical setting.⁹

C. Preferences of the Programs

Since the value of a team of residents to a program is not observed, I model residency program preferences through a latent variable. A very rich specification creates two extreme problems. On the one hand, a pairwise stable match need not exist if a program's (ordinal) preference for a given resident depends crucially on the other residents it hires. On the other hand, the number of stable matches can be exponentially large in the number of agents when programs have heterogeneous preferences.¹⁰ Additional hurdles include identifying a rich specification from observed matches alone, and computational difficulties in simulating outcomes using a dataset from a large market.

I use a restrictive specification for program preferences that solves the problems above while accounting for the primary resident characteristics that are valued by programs but are not observed in the dataset. My conversations with residency program and medical school administrators suggest that programs assess resident skill using test scores in medical exams, clinical performance, and the strength of recommendation letters. While fit in the program is considered, they broadly agree on resident qualifications and refer to a "pecking order" for residency slots. Therefore, I model a program's preference for a resident using an ordinal human capital index $H(x_i, \epsilon_i)$ that is a function of observable characteristics x_i of a resident and an unobservable determinant ϵ_i . I use the parametric form

$$(3) \quad h_i = \mathbf{x}_i \alpha + \epsilon_i,$$

where ϵ_i is normally distributed with a variance that depends on the type of medical school a resident graduated from. For graduates of allopathic (MD) medical schools, \mathbf{x}_i includes the log NIH funding and median MCAT scores of the resident's medical school. Characteristics also include the medical school type for residents, i.e., whether a resident earned an osteopathic degree (DO) or graduated from a foreign medical school. I also include an indicator for whether a resident that graduated from a foreign medical school was born in the United States. Without loss of generality, the variance of ϵ_i for residents with MD degrees is normalized to 1 and the mean of h at $\mathbf{x} = 0$ is normalized to zero.

⁹This fact also generates computational demands for a simulation-based estimation since memory requirements grow with the product of the number of residents and programs. This difficulty can be avoided in demand models (e.g., Berry, Levinsohn, and Pakes 1995) since logit choice probabilities, conditional on simulated random coefficients, can be computed in closed form.

¹⁰Simulations in Roth and Peranson (1999) conducted with data reported to the NRMP suggest that multiplicity may not be empirically important. They find that almost all of the residents are matched to the same program across all of the stable matches.

This specification guarantees the existence and uniqueness of a stable match and a computationally tractable simulation algorithm that is described in Section VC. Existence is implied by the assumption that ordinal preferences over residents do not depend on the other residents hired by a program (see Roth and Sotomayor 1992).¹¹ Uniqueness results from the assumption that all programs have identical ordinal preferences (Clark 2006; Niederle and Yariv 2009; Pycia 2012).¹² Finally, Section IIIC notes that identifying a model with heterogeneity relies on exclusion restrictions, in this case an observable program characteristic that is excluded from the preferences of the residents for programs.

From the perspective of the application studied in the paper, the assumption that programs have vertical preferences is essential for calculating the implicit tuition. As we will see, a conservative estimate of the markdown from a perfect competition benchmark can be computed under this assumption using an estimated money-metric utility function for residents.¹³ One may then worry that misspecification of program preferences in the model will cause bias in estimates of resident preferences. Monte Carlo exercises conducted to address this issue are discussed briefly in Section VIB and presented in online Appendix B1.

III. Identification

This section describes how I leverage information in the sorting patterns and many-to-one matching to identify preference parameters. The discussion also guides the choice of moments used in estimation. Standard revealed preference arguments do not apply because “choice-sets” of individuals are unobserved and determined in equilibrium. The market index t is omitted in this section because the arguments are based on observing one market with many (interdependent) matches. For simplicity, I also assume that the number of residents is equal to the number of residency positions and treat all characteristics as exogenous. Identification of the case with endogenous salaries is discussed in Section IV, and does not require a reconsideration of arguments presented here.

A. Using Sorting Patterns: The Double-Vertical Model

Consider the simplified “double-vertical” model in which all residents agree upon the relative ranking of programs. In a linear parametric form for indirect utilities, preferences are represented with

¹¹This restriction may not be particularly strong because programs cannot submit ranks that depend on the rest of the team, indicating that incorporating such preferences was not an important design consideration in the NRMP. Further, this restriction on rankings that can be submitted limits the hope for identifying richer preference models given that the observed matches were generated using program ranks that don’t depend on the rest of the team.

¹²The ordinal preferences in the estimated model are consistent with any program-specific latent output function $F_j(h_{i_1}, \dots, h_{i_{c_j}})$ from a team of residents (i_1, \dots, i_{c_j}) at program j that is strictly increasing in each of its components. Although it allows for (limited) complementarities in this latent output, the assumption implies responsive ordinal preferences and hence, the existence of a many-to-one stable match (Roth and Sotomayor 1992, lemma 5.6). The counterfactuals in this paper will be based on the additional restriction that $F_j(h_{i_1}, \dots, h_{i_{c_j}}) = \sum_{k=1}^{c_j} f(h_{i_k})$, one that may not be necessary in other applications.

¹³Without the preference restriction on program preferences, it may be necessary to estimate a money-metric utility function for programs. This may not be feasible without further assumptions and/or data because a program pays each of its residents an identical salary and resident output is not observed.

$$u_j = \mathbf{z}_j \beta + \xi_j$$

$$h_i = \mathbf{x}_i \alpha + \varepsilon_i,$$

where x_i and z_j are observed and ξ_j and ε_i are standard normal random variables, distributed independently of the observed traits. Assume the location normalizations $E[u_j | \mathbf{z}_j = 0] = 0$ and $E[h_i | \mathbf{x}_i = 0] = 0$.

We first need a sign restriction on one parameter of the model to interpret sorting patterns in terms of preference parameters. Consider a model in which x measures the prestige of a resident's medical school and z measures the size of the hospital with which a program is associated. In this example, residents from prestigious medical schools sort into larger hospitals if the human capital distribution of residents from more prestigious medical schools is higher and hospital size is preferable. However, this sorting may also have been observed if these traits were undesirable to both sides. The observation necessitates restricting one characteristic of either residents or programs to be desirable. Throughout the paper, I assume that residents graduating from more prestigious medical schools, as measured by the NIH funding of the medical school, are more likely to have a higher human capital index. Under this sign restriction, the sorting patterns observed in column 1 of Table 2 can only be rationalized if a program's desirability is positively related to the NIH funding of its affiliates.

Now I argue that we can determine whether two observable types are equally desirable. Note that the set of programs with a higher value of $\mathbf{z}\beta$ have a higher distribution of utility to residents, and are therefore matched with residents with higher human capital. Therefore, if $\mathbf{z}\beta > \mathbf{z}'\beta$, the distribution of observable characteristics of residents matched with programs of type \mathbf{z} must be different than that of \mathbf{z}' . The sorting observed in the data thus informs us whether two observable types of programs (analogously residents) are equally desirable or not. For example, assume that there are two types of programs, one at larger but less prestigious hospitals than another program at a smaller hospital. The residents matched with these two hospital types have the same distribution of observable characteristics only if residents trade off hospital size for prestige.

Agarwal and Diamond (2014) formalizes this argument in a model with nonparametric functions of \mathbf{x} and \mathbf{z} , and nonparametric distributions for the additively separable errors ε and ξ . They prove that sorting patterns can be used to determine if \mathbf{x} and \mathbf{x}' (likewise, \mathbf{z} and \mathbf{z}') are equally desirable, but not the distribution of preferences.

B. Importance of Data from Many-to-One Matches

The preceding arguments using only sorting patterns do not contain information on the relative importance of observables on the two sides of the market. Specifically, we cannot determine the extent to which unobservable characteristics such as resident academic record or quality of program faculty are important in determining the observed matchings. To build intuition, consider an example in which x is a binary indicator for a resident graduating from a prestigious medical school and z is a binary indicator for a program at a large hospital. Assume that medical school

prestige and hospital size are desirable characteristics ($\alpha > 0$ and $\beta > 0$). Sorting patterns from such a model can be summarized in a contingency table in which residents from prestigious medical schools are systematically more likely to match with programs at large hospitals. For instance, consider the following table:

	$z = 1$	$z = 0$
$x = 1$	30 percent	20 percent
$x = 0$	20 percent	30 percent

These matches could result from parameters under which programs have a strong preference for residents from prestigious medical schools (large α) and residents have a moderate preference for large hospitals (small β). In this case, residents from more prestigious medical schools get their pick of programs, but often choose ones at small hospitals. On the other hand, the contingency table could have been a result of a strong preference for large hospitals (large β) but only a moderate preference for residents from prestigious medical schools (small α). There are a variety of intermediate cases that are indistinguishable from each other and either extreme. Intuitively, this ambiguity results from the ability of unobservable characteristics on either side to fit the observed sorting.

In addition to sorting patterns, data on many-to-one matches have information on the extent to which residents matched to the same program have similar characteristics. In a pairwise stable match, two residents at the same program must have similar human capital irrespective of the program's quality. Otherwise, either the program could replace the lower quality resident with a better resident, or the higher quality resident could find a more desirable program. Residents training at the same program have similar observables if x is highly predictive of human capital.¹⁴ Conversely, programs are not likely to match with multiple residents with similar observables if they placed a low weight on x . The variation in resident observable characteristics within programs is therefore a signal of the information observables contain about the underlying human capital of residents.

This information is not available in a one-to-one matching market because sorting patterns are the only feature known from the data. Agarwal and Diamond (2014) formally shows that having data from many-to-one matches is critical for identifying the parameters of the model, and provides simulation evidence to illustrate the limitations of sorting patterns and the usefulness of many-to-one matching data.¹⁵

Descriptive Statistics from Many-to-One Matching.—Table 4 shows the fraction of variation in resident characteristics that is within a program. Notice that almost all of the variation in the gender of the resident is within programs. In contrast, residents are more systematically sorted into programs where other residents have more

¹⁴The arguments suggest that there should be a low variance in $\mathbf{x}\alpha$, which has implications on the within-program covariance between various components of \mathbf{x} . The moments used for estimation will use this information. I treat \mathbf{x} as a scalar in this discussion for simplicity and because the sorting patterns can be used to determine α up to scale.

¹⁵In my experience, an optimization routine that did not use information from many-to-one matching had difficulty in converging.

TABLE 4—WITHIN-PROGRAM VARIATION IN RESIDENT CHARACTERISTICS

	Fraction of variation within program-year (percent)
log NIH fund (MD)	77.83
Median MCAT (MD)	72.09
US-born foreign graduate	79.01
Osteopathic/DO degree	85.16
Foreign degree	57.16
Allopathic/MD degree	64.81
Female	96.40

Notes: Adjusted R^2 from separate regressions of resident characteristic absorbing the program year fixed effects. Samples for “log NIH funding (MD)” and “Median MCAT (MD)” are restricted to residents with nonmissing values. Regression of “US-born foreign graduates” restricted to graduates of foreign medical schools.

similar academic qualifications. For instance, about 30 percent of the variation in the median MCAT score of the residents’ graduating medical school decomposes into across program variation. These summaries are consistent with academic qualifications being associated with human capital, but not resident gender. If gender were a strong determinant of a resident’s desirability to a program, in a double-vertical model, one would expect that programs would be systematically male or female dominated.

Table 5 presents another summary from many-to-one matching based on regressing the leave-one-out mean characteristic of a resident’s peer group in a program on the characteristics of the resident. Let $\bar{x}_{-i,k}^\mu$ be the average observable x_k of resident i ’s peers for a match μ , i.e., $\bar{x}_{-i,k}^\mu = \frac{1}{|\mu^{-1}(\mu(i))| - 1} \sum_{i' \in \mu^{-1}(\mu(i)) \setminus \{i\}} x_{i',k}$. I estimate the equation

$$\bar{x}_{-i,k}^\mu = \mathbf{x}_i \lambda + e_i,$$

where \mathbf{x}_i is a vector of resident i ’s observables. Each regression suggests that a resident’s characteristic is positively associated with the mean of the same characteristic of her peers. Viewing NIH funding, MCAT scores, and MD degree as quality indicators, there is a positive association between a resident’s quality and the average quality of her peer group. Further, the moderately high R^2 statistics for these regressions indicate that resident characteristics are more predictive of her peer groups than the results in Table 4 suggest. This is in part because, on average, residents that have low values on a given dimension should have compensating characteristics on another dimension. Taking the characteristics together improves the ability to predict the observable quality of the peer group.

C. Heterogeneity in Preferences

I now intuitively discuss the role of exclusion restrictions in learning about heterogeneity in preferences. These assumptions either restrict that a given agent characteristic is not valued by the other side of the market, or that agents with differing values of that characteristic have identical preferences.

TABLE 5—PEER SORTING

	Peer log NIH fund (1)	Peer log MCAT (2)	Peer foreign degree (3)	Peer DO degree (4)	Peer MD degree (5)
log NIH fund (MD)	0.2919*** (0.0132)	0.0103*** (0.0026)	-0.0249*** (0.0030)	-0.0043** (0.0019)	0.0293*** (0.0033)
log median MCAT (MD)	0.6449*** (0.1832)	0.0874 (0.0750)	-0.2000*** (0.0458)	0.0165 (0.0247)	0.1850*** (0.0499)
US born (for)	0.0403 (0.0421)	0.0141 (0.0103)	-0.1063*** (0.0091)	0.0394*** (0.0050)	0.0669*** (0.0079)
Observations	19,830	19,845	24,066	24,066	24,066
R ²	0.1280	0.6437	0.3632	0.0914	0.3197

Notes: Linear regression of average characteristics of peers on the characteristics of a resident. Column 1 restricts to the set of residents with at least one peer from a medical school with nonzero NIH funding. Column 2 restricts residents with at least one peer from a medical school with nonmissing MCAT score. All specifications include medical school type dummies and a dummy for residents from medical schools without NIH funding. Standard errors clustered at the program-year level in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

There are two sets of potential exclusion restrictions, one for each side of the market. I first discuss the exclusion restrictions on resident characteristics. The estimated model excludes the birth and medical school location of a resident from program preferences. Excluding the birth location, for example, implies that conditional on the quality of medical school, the propensity of residents for matching to programs close to their birthplace can only be a result of resident preferences, not the preferences of programs. Further, residents matching close to home do so at disproportionately lower quality programs only because they trade off program quality with preferences for location. The principle is similar to the use of variation excluded from one part of a system to identify a simultaneous equation model. The exclusion restriction in the example above isolates a factor influencing the demand for residency positions without affecting the distribution of choice sets faced by residents. Since programs have homogeneous preferences in the model, a violation of this assumption occurs if, conditional on medical school characteristics, residents born in a particular state are systematically more skilled than residents born elsewhere.

The empirical specifications also exclude the determinants of human capital (\mathbf{x}_i, ϵ_i) from preferences of residents for programs. This results in choice set variation across residents with different skills but similar preferences. Ideally, one would be able to estimate a preference distribution for programs that is heterogeneous across residents from different medical schools or with different skill levels. Richer specifications that allow for this type of preference heterogeneity were difficult to estimate, potentially because the available quality indicators of residents only include information on a resident's medical school, and do not vary at the individual level.

The second set of exclusion restrictions is on program characteristics. These restrictions are embedded in the homogeneous preference assumption discussed earlier. I do not exclude any program characteristics from resident preferences.

Existing formal analyses of identification of non-transferable utility matching models have either assumed that preferences are homogeneous on both sides of

the market (Agarwal and Diamond 2014) or have avoided exclusion restrictions but obtained results on identifying only the sum of utilities in a match (Menzel forthcoming).

IV. Salary Endogeneity

The salary offered by a residency program may be correlated with unobserved program covariates, resulting in biased estimates of residents' willingness to pay for programs. However, the bias cannot be signed ex ante in this setting. Programs with desirable unobserved traits may be able to pay lower salaries in equilibrium due to compensating differentials. Alternatively, desirable programs may be more productive or better funded, resulting in salaries that are positively associated with unobserved quality.

While it would be ideal to have a full model of salary setting, I avoid this for several reasons. First, the allegation of collusive salary setting in the lawsuit is unresolved. Second, hospitals tend to set identical salaries for residents in all specialties, suggesting that a full model should consider the joint salary setting decision across all residency programs at a hospital. Finally, a full model would need to account for accreditation requirements that require salaries to be adequate for a resident's living and educational expenses.¹⁶

A. A Control Function Approach

I propose a control function correction for potential bias due to correlation between salaries w_{jt} and program unobservables ξ_{jt} (see Heckman and Robb 1985; Blundell and Powell 2003; Imbens and Newey 2009). The method uses an instrument r_{jt} that is excluded from the utility function $U(\cdot)$. I describe the instrument in the next section.

I use a linear control function in which the salary, w_{jt} , is specified as

$$(4) \quad w_{jt} = \mathbf{z}_{jt}\gamma + r_{jt}\tau + \nu_{jt},$$

where \mathbf{z}_{jt} are program observable characteristics, r_{jt} is the instrument, and ν_{jt} is an unobservable. Endogeneity of w_{jt} is captured through correlation between the unobservables ν_{jt} and ξ_{jt} .

The approach requires (ξ_{jt}, ν_{jt}) to be independent of $(\mathbf{z}_{jt}, r_{jt})$. This assumption replaces the weaker conditional moment restriction needed in a linear instrumental variables approach.¹⁷ Under this independence, although w_{jt} is not (unconditionally) independent of ξ_{jt} , it is conditionally independent of ξ_{jt} given ν_{jt} and \mathbf{z}_{jt} . For

¹⁶The ACGME sponsoring institution requirements state that "Sponsoring and participating sites must provide all residents with appropriate financial support and benefits to ensure that they are able to fulfill the responsibilities of their educational programs."

¹⁷Imbens (2007) discusses these independence assumptions, noting that they are commonly made in the control function literature and are often necessary when dealing with a non-additive second stage. In this context, even though ξ_{jt} is additively separable from w_{jt} , the approach used in demand models pioneered by Berry (1994) and Berry, Levinsohn, and Pakes (1995), where an inversion can be used to estimate a variable with a separable form, is not available.

tractability given the limited salary variation, I further model the distribution of ξ_{jt} conditional on ν_{jt} as

$$(5) \quad \xi_{jt} = \kappa\nu_{jt} + \sigma\zeta_{jt},$$

where $\zeta_{jt} \sim N(0, 1)$ is drawn independently of ν_{jt} and (κ, σ) are unknown parameters. The specification allows the unobservable characteristic of the program ξ_{jt} to be correlated across time through ν_{jt} .

Given this specification, one can rewrite equation (1) as

$$(6) \quad u_{ijt} = \mathbf{z}_{jt}\beta_i^z + w_{jt}\beta_i^w + \kappa\nu_{jt} + \sigma\zeta_{jt},$$

and normalize the scale parameter σ to 1. Following common practice, I estimate equation (4) using OLS to obtain a consistent estimate of ν_{jt} as a conditioning variable in place of its true value. Hence, we can treat it as any other observed characteristic. Since variation in w_{jt} given ν_{jt} and \mathbf{z}_{jt} is due to r_{jt} , the assumptions above imply that ζ_{jt} is independent of w_{jt} , solving the endogeneity problem.

While it may be possible to relax the linear specification in principle, an important restriction in this approach is that unobservables of competitor programs cannot affect salaries (except through ν_{jt}). Nonetheless, a linear specification has been shown to substantially reduce bias in estimates even in models of oligopolistic competition in which the price has a nonlinear relationship with unobservables and the characteristics of competing products (Yang, Chen, and Allenby 2003; Petrin and Train 2010).

B. Instrument

Table 6 presents regression estimates of equation (4), except using a log-log specification so that coefficients can be interpreted as elasticities. The first four columns do not include the instrument r_{jt} , which is defined below. Column 1 shows limited correlation between salaries and observed program characteristics except rents and the Medicare wage index. The elasticity with respect to these two variables is small, at less than 0.15 in magnitude. This suggests that models that do not instrument for salaries may provide reasonable approximations.

To address potential correlation, however, I will also present estimates that use an instrument. Conversations with program directors suggest that salaries paid by competitors in a program's geographic area are viewed as benchmarks for setting salaries in the residency market.¹⁸ Column 2 shows that lagged salaries paid by other hospitals in the program's area are a significant predictor of salaries paid by a program, with an R^2 statistic of about 0.33. The use of this benchmark suggests instruments based on geographic aggregates that affect salaries but are unrelated to resident preferences. I will construct an instrument using Medicare reimbursement rates for residency training at competitor hospitals.

¹⁸The Council of Teaching Hospitals annually publishes statistics on the compensation paid to medical residents, disaggregated by geographical regions.

TABLE 6—WAGE REGRESSIONS

	log first year salary					
	(1)	(2)	(3)	(4)	(5)	(6)
log rent	-0.0373** (0.0177)	0.0179 (0.0140)	-0.0379** (0.0175)	-0.0378** (0.0160)	0.0172 (0.0143)	-0.0306 (0.0230)
Rural program	0.0065 (0.0081)	0.0103 (0.0071)	0.0110 (0.0080)	0.0104 (0.0076)	0.0103 (0.0069)	0.0055 (0.0079)
log wage index	0.1366*** (0.0307)	-0.0152 (0.0262)	0.1182*** (0.0302)	0.0806*** (0.0287)	-0.0167 (0.0263)	0.0809*** (0.0290)
log NIH fund (major)	0.0024 (0.0027)	0.0062*** (0.0021)	0.0023 (0.0026)	0.0034 (0.0025)	0.0062*** (0.0021)	0.0024 (0.0024)
log NIH fund (minor)	-0.0060* (0.0032)	-0.0005 (0.0029)	-0.0047 (0.0032)	-0.0040 (0.0031)	-0.0005 (0.0029)	-0.0041 (0.0031)
log number beds	0.0087* (0.0046)	0.0012 (0.0036)	0.0086* (0.0045)	0.0064 (0.0041)	0.0010 (0.0036)	0.0108** (0.0043)
log case-mix index	-0.0108 (0.0195)	0.0051 (0.0151)	-0.0046 (0.0195)	-0.0038 (0.0190)	0.0056 (0.0152)	-0.0065 (0.0191)
log competitor salary (lagged)		0.8779*** (0.0542)			0.8651*** (0.0683)	
log reimbursement			0.0227*** (0.0077)	0.0064 (0.0076)	-0.0002 (0.0063)	0.0050 (0.0070)
log competitor reimbursement				0.0968*** (0.0170)	0.0090 (0.0170)	0.0847*** (0.0178)
Location characteristics						Y
Observations	3,418	2,997	3,418	3,418	2,997	3,418
R ²	0.0452	0.3284	0.0640	0.1226	0.3294	0.1811

Notes: Regression of a program's first year salary on program characteristics. All columns include dummy variables for programs with no NIH funding at major affiliates, for no NIH funding at minor affiliates, and a dummy for missing Medicare ID for the primary institution. "Competitor salary (lagged)" is the average of lagged salaries of other family practice residency programs in the geographic area of the program hospital. "Competitor reimbursement" is a weighted average of the Medicare primary care per resident amounts of institutions in the geographic area of a program other than the primary institutional affiliate of the program. Geographic area defined as in Medicare DGME payments: MSA/NECMA unless less than three competitors with available data in the area, in which case the census region is used. For columns 3–6, a program's reimbursement rate is truncated at \$5,000 (46 observations) and a dummy for truncated observations is estimated. Sample restricted to programs for which salary was not imputed as described in the online Appendix. Standard errors clustered at the program level in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Medicare's reimbursement to residency programs for direct costs of training is based on cost reports submitted in 1984. Before the prospective payment system was established, the total payment made to a hospital did not depend on the precise classification of costs as training or patient care costs. The reimbursement system for residency training was severed from payments for patient care in 1985 because the two types of costs were considered distinct by the government. While patient care was reimbursed based on fees for diagnosis-related groups, reimbursements for residency training were calculated using cost reports in a base period, usually 1984. Line items related to salaries and benefits, and administrative expenses of residency programs were designated as direct costs of residency training. A per-resident amount was calculated by dividing the total reported costs on these line items by

the number of residents in the base period. Today, hospitals are reimbursed based on this per-resident amount, adjusted for inflation using CPI for all urban consumers.

This reimbursement system therefore uses reported costs from two decades prior to the sample period for this study. More importantly, the per-resident amount may not reflect costs even in the base period because hospitals had little incentive to account for costs under the correct line item. The mean reimbursement rate for family medicine residency programs is \$71,131, with a standard deviation of \$28,023 (in 2010 dollars). Newhouse and Wilensky (2001) argue that this variation in per-resident amounts is primarily driven by differences in hospital accounting practices rather than real costs because the distinction between patient care costs from those incurred due to residency training is arbitrary. In other words, whether a cost, say salaries paid to attending physicians, was accounted for in a line item later designated for direct costs can significantly influence reimbursement rates today. These reimbursements are earmarked for costs of residency training and are positively associated with salaries paid by a program today (column 3 of Table 6).

I instrument using a weighted average of reimbursement rates of other teaching hospitals in the geographic area of a program. The instrument is defined as

$$(7) \quad r_j = \frac{\sum_{k \in G_j} fte_k \times rr_k}{\sum_{k \in G_j} fte_k},$$

where rr_k and fte_k are the reimbursement rate and number of full-time equivalent residents at program k 's primary hospital in the base period, and G_j are the hospitals in program j 's geographic area other than j 's primary hospital. I base the geographic definitions on Medicare's physician fee schedule, i.e., the metropolitan statistical area (MSA) of the hospital or the rest of state if the hospital is not in an MSA. If less than three other competitors are in this area, define G_j to be the census division.¹⁹ The variation in competitor reimbursement rates is significant, with a standard deviation of \$13,429 (in 2010 dollars).

Consistent with the theory for the instrument's effect on salaries, column 4 shows that competitor reimbursements are positively related to salaries. When estimated in levels rather than logs, this specification represents equation (4), which is analogous to the first stage in a two-stage least squares method.²⁰ In column 5, I test the theory that competitor reimbursements affect salaries only through competitor salaries. Relative to column 4, controlling for the lagged average competitor salaries reduces the estimated effect of competitor reimbursements by an order of magnitude and results in a statistically insignificant effect.

The key assumption for validity of the instrument is that the program unobservable ξ_{jt} is conditionally independent of competitor reimbursement rates, given program characteristics and a program's own reimbursement rate, which is included in

¹⁹ Medicare's reimbursement scheme and the construction of the instrument are detailed in online Appendix C.

²⁰ Figure C1 in the online Appendix depicts an increasing relationship between salary and competitor reimbursements. Clustered at the program level, the first stage F -statistic for the coefficient on the instrument is 37.6. While greater than the commonly used threshold value of 10, Stock, Wright, and Yogo (2002) point out that the cutoff is only valid for linear models. Since the control function approach is based on assuming independence, I also tested for heteroskedasticity in the residuals from the first stage. The hypothesis that the residuals are homoskedastic cannot be rejected at the 90 percent confidence level using either the tests proposed by Breusch and Pagan (1979) or by White (1980).

z_{jt} for specifications using the instrument. This assumption is satisfied if variation in reimbursement rates is driven by an arbitrary classification of costs by hospitals in 1984 or if past costs of competitors are not related to residents' preferences during the sample period. The primary threat is that reimbursement rates are correlated with persistent geographic factors excluded from resident preferences. To some extent, this concern is mitigated by controlling for a program's own reimbursement rate. Additionally, column 6 shows that including location characteristics such as median age, household income, crime rates, college population, and total population changes the impact of competitor reimbursement rates on a program's salary by less than the standard error in the estimates.²¹ Another concern is the possibility that programs respond to the reimbursement rates of competitors by engaging in endogenous investment. A comparison of estimates from columns 1 and 4 shows little evidence of sensitivity of the coefficients on observed program characteristics (NIH, beds, case mix index) to the inclusion of reimbursement rate variables.

V. Estimation

This section defines the estimator, the moments used, and the simulation technique.

A. Simulated Minimum Distance

The estimation proceeds in two stages when the control function is employed. I first estimate the control variable ν_{jt} from equation (4) using ordinary least squares (OLS) to construct the residual

$$(8) \quad \hat{\nu}_{jt} = w_{jt} - \mathbf{z}_{jt}\hat{\gamma} - r_{jt}\hat{\alpha}.$$

Replacing this estimate in equation (6), we get

$$(9) \quad u_{ijt} \approx \mathbf{z}_{ijt}\beta_i^z + w_{jt}\beta_i^w + \kappa\hat{\nu}_{jt} + \zeta_{jt},$$

where the approximation is up to estimation error in ν_{jt} . The estimation of parameters determining the human capital index of residents and resident preferences over programs proceeds by treating $\hat{\nu}_{jt}$ like any other exogenous observable program characteristic. The error due to using $\hat{\nu}_{jt}$ instead of ν_{jt} , however, affects the calculation of standard errors. This first stage is not necessary in the model treating salaries as exogenous.

The distribution of preferences of residents and human capital can be determined as a function of observable characteristics of both sides and the parameter

²¹I also directly assess whether excluded geographical aggregates of population, income, share of population that is college educated, age, and crime rates are strong predictors of competitor reimbursements in online Appendix C. Estimates in Table C1 suggest that bias due to excluding these characteristics, if any, is likely small. First, most of the excluded characteristics are not statistically significant. While median age and property crime rates have significant coefficients, their economic magnitudes are small. Second, the excluded characteristics together explain less than 6 percent of the variation in addition to controls already included in the model. As noted earlier, these characteristics don't drive the effect of the instrument on resident salaries.

of the model, θ , collected from equations (2), (3), and (6). The second stage of the estimation uses a simulated minimum distance estimator (McFadden 1989; Pakes and Pollard 1989; Gouriéroux and Monfort 1997). The estimate $\hat{\theta}_{SMD}$ minimizes a simulated criterion function

$$(10) \quad \|\hat{\mathbf{m}} - \hat{\mathbf{m}}^S(\theta)\|_{\mathbf{W}}^2 = (\hat{\mathbf{m}} - \hat{\mathbf{m}}^S(\theta))' \mathbf{W} (\hat{\mathbf{m}} - \hat{\mathbf{m}}^S(\theta)),$$

where $\hat{\mathbf{m}}$ is a set of moments constructed using the matches observed in the sample, $\hat{\mathbf{m}}^S(\theta)$ is the average of moments constructed from S simulations of matches in the economy. The choice of weight matrix \mathbf{W} , the inference procedure, the optimization algorithm, and additional details on the estimator are in online Appendix B. For inference, the main challenge is in dealing with interdependence in the matches. I use a parametric bootstrap for the moments, and a subsequent delta method to construct standard errors. Finally, the choice of optimization algorithm must consider the fact that the simulated objective function is not smooth.

B. Moments

The vector $\hat{\mathbf{m}}$ consists of sample analogs of three sets of moments, calculated separately for each market and then averaged across markets. The simulated counterparts $\hat{\mathbf{m}}^S(\theta)$ are computed identically, but averaged across the simulations and markets.

For the match μ_t observed in market t , the set of moments are given by

- (i) Moments of the joint distribution of observable characteristics of residents and programs as given by the matches. For every pair of scalars $x_{k,i}$ and $z_{m_k,jt}$,

$$(11) \quad \hat{m}_{t,ov,k} = \frac{1}{N_t} \sum_{i \in \mathcal{N}_t} 1\{\mu_t(i) = j\} x_{k,i} z_{m_k,jt}.$$

- (ii) The within-program variance of resident observables. For each scalar $x_{k,i}$,

$$(12) \quad \hat{m}_{t,w,k} = \frac{1}{N_t} \sum_{i \in \mathcal{N}_t} \left(x_{k,i} - \frac{1}{|\mu_t^{-1}(\mu_t(i))|} \sum_{i' \in \mu_t^{-1}(\mu_t(i))} x_{k,i'} \right)^2.$$

- (iii) The covariance between resident characteristics and the average characteristics of a resident's peers. For every pair of scalars $x_{k,i}$ and $x_{m_k,i}$,

$$(13) \quad \hat{m}_{t,p,k} = \frac{1}{N_t} \sum_{i \in \mathcal{N}_t} x_{k,i} \frac{1}{|\mu_t^{-1}(\mu_t(i))| - 1} \sum_{i' \in \mu_t^{-1}(\mu_t(i)) \setminus \{i\}} x_{m_k,i'}.$$

The first set of moments include the covariances between program and resident characteristics. These moments are the basis of the regression coefficients presented in Tables 2 and 3. They quantify the degree of assortativity between resident and program characteristics observed in the data.²² The second and third set of moments

²²I include covariances for every pair of observed resident and program characteristics. Specifications employing random coefficients also use the square of the corresponding program characteristic. I also include the

take advantage of the many-to-one matching nature of the market.²³ Section IIIB presents summaries of these moments from the data. These moments cannot be constructed in one-to-one matching markets, such as the marriage market, but as formally discussed in Agarwal and Diamond (2014) are crucial to identify even the simpler double-vertical model. Since these moments extract information from within a peer group without reference to the program in which they are training, they effectively control for both observable and unobservable program characteristics.²⁴

C. Simulating a Match

Under the parametric assumptions made on ζ_{jt} , ε_i , and η_i in Section II, for a given parameter vector θ , a unique pairwise stable match exists and can be simulated. Because residents only participate in one market, matches in different markets can be simulated independently. For simplicity, I describe the procedure for only one market and omit the market subscript t . For a draw of the unobservables $\{\varepsilon_{is}, \eta_{is}\}_{i=1}^N$ and $\{\zeta_{js}\}_{j=1}^J$ indexed by s , calculate

$$h_{is} = \mathbf{x}_i \alpha + \varepsilon_{is},$$

and the indirect utilities $\{u_{ijs}\}_{i,j}$. The indirect utilities determine the program resident i picks from any choice set.

Begin by sorting the residents in order of their simulated human capital index, $\{h_{is}\}_{i=1}^N$, and let $i^{(k)}$ be the identity of the resident with the k th highest value.

- **Step 1:** Resident $i^{(1)}$ picks her favorite program. Set her simulated match, $\mu_s(i^{(1)})$, to this program and compute $J^{(1)}$, the set of programs with unfilled positions after $i^{(1)}$ is assigned.
- **Step $k > 1$:** Let $J^{(k-1)}$ be the set of programs with unfilled positions after resident $i^{(k-1)}$ has been assigned. Set $\mu_s(i^{(k)})$ to the program in $J^{(k-1)}$ most desired by $i^{(k)}$.

The simulated match μ_s can be used to calculate moments using equations (11)–(13). The optimization routine keeps a fixed set of simulation draws of unobservable characteristics for computing moments at different values of θ .

A model with preference heterogeneity on both sides requires a computationally more complex simulation method, such as the deferred acceptance algorithm (Gale and Shapley 1962), to compute a particular pairwise stable match.²⁵

probability that a resident is matched to a program located in the same state as her state of birth, or the same state as her medical school state.

²³ Alternatively, one could combine moments of type 2 and 3 to include all entries in (the upper triangular portion of) the within-program covariance matrix of characteristics. In estimation, the second set includes every resident characteristic and the third set includes all interactions.

²⁴ Note that the number of moments suggested increases rapidly as more characteristics are included in the preference models. If the covariance between each observed characteristic of the resident and of the program are included in the first set of moments, the number of moments is at least the product of the number of characteristics of each side while the number of parameters is the sum. This growth can create difficulties when estimating models with a very rich set of characteristics.

²⁵ In the deferred acceptance algorithm, each applicant simultaneously applies to her most favored program that has not yet rejected her. A set of applications are held at each stage while others are rejected and assignments are

VI. Empirical Specifications and Results

I present estimates from four models. The first model does not allow for unobserved heterogeneity in resident preferences but allows for heterogeneity in taste for program location based on a resident's birth and medical school location. The second model has the richest form of preferences as it allows for unobserved heterogeneity in preferences for diagnostic mix, research focus, and hospital size via normally distributed random coefficients on case mix index, NIH funds of major medical school affiliates, and the number of beds. The comparison allows us to assess robustness and the importance of unobserved preference heterogeneity. These two models treat salaries as exogenous. The third and the fourth models modify the first and second models respectively to address potential endogeneity in salaries using the instrument described in Section IVB. The last two specifications include a program's own reimbursement rate in addition to characteristics included in the other models.

Parameter estimates for residents' preferences for programs presented in the next section are translated into dollar equivalents for a select set of program characteristics. These are the most economically relevant statistics obtained from preference estimates. Online Appendix D briefly discusses the underlying parameters, which are not economically intuitive.

A. Preference Estimates

Panels A1 and A2 of Table 7 present the estimated preference parameters for programs in salary equivalent terms. These estimates are best interpreted as describing the types of programs residents value, which may derive from effects on future earnings or due to intrinsic preferences for the work or training environment. Panel B presents parameter estimates for the distribution of human capital, which determines ordinal rankings between residents.

Estimates without Salary Instruments.—Comparing Specifications (1) and (2), the estimated value of a 1 standard deviation higher case mix index at an otherwise identical program is about \$2,500–\$5,000 in annual salary for a typical resident. Likewise, residents are willing to pay for programs at larger hospitals as measured by beds, and for programs with higher NIH funded affiliates. The estimates from Specification (2) suggest a substantial degree of preference heterogeneity for these characteristics as well. The additional heterogeneity in preferences relative to Specification (1) results in a shift in the mean willingness to pay for NIH funding of major affiliates, the case mix index, and beds, but not whether they are desirable or not. The estimates also suggest that, despite the availability of an extensive set of program characteristics, program unobservables are significant. The willingness to pay for a standard deviation increase in the program unobservable has a similar order of magnitude as the case mix index.

made final only when no further applications are rejected. This temporary nature of held applications and the need to compute a preferred program for all applicants at each stage significantly increases the computational burden for a market with many participants such as the one studied in this paper.

TABLE 7—PREFERENCE ESTIMATES

	Geographic heterogeneity (1)	Full heterogeneity (2)	Geographic heterogeneity with instrument (3)	Full heterogeneity with instrument (4)
<i>Panel A1. Preference for programs (in US\$ for 1 standard deviation change)</i>				
Case mix index—coefficient	2,320.26 (335.74)	4,792.02 (918.88)	6,088.35 (1,216.81)	35,353.43 (23,061.80)
Sigma R.C.		4,502.71 (1,036.50)		14,336.18 (10,695.10)
log NIH fund. (major)— coefficient	6,498.66 (672.84)	490.67 (711.09)	4,402.22 (996.46)	−1,229.79 (3,422.43)
Sigma R.C.		5,497.70 (1,233.70)		29,159.39 (19,073.86)
log beds—coefficient	3,528.07 (341.98)	6,899.76 (1,245.58)	8,837.08 (1,616.13)	20,438.57 (13,799.62)
Sigma R.C.		11,107.09 (2,073.48)		47,944.02 (30,956.09)
log NIH fund. (minor)	5,559.78 (556.53)	4,993.23 (911.30)	7,619.88 (1,405.87)	30,385.76 (19,747.07)
Program unobservable	2,179.23 (213.72)	4,329.29 (600.79)	5,119.97 (926.17)	20,068.14 (12,781.22)
<i>Panel A2. Preference for programs</i>				
Rural program	5,610.75 (924.59)	7,327.49 (1,712.59)	17,313.61 (3,547.54)	25,539.74 (19,001.56)
University-based program	11,080.26 (1,482.34)	15,786.30 (2,337.62)	25,129.66 (5,279.00)	73,468.62 (48,398.50)
Community/university program	−2,217.05 (354.60)	−5,000.99 (1,039.03)	−7,506.98 (1,528.10)	−34,181.16 (22,490.08)
Medical school state	2,301.94 (218.21)	9,819.73 (1,386.89)	4,528.94 (821.05)	45,337.71 (29,632.43)
Birth state	1,319.99 (128.16)	6,342.26 (892.44)	2,450.98 (447.97)	29,386.24 (19,241.60)
Rural birth × rural program	108.89 (26.53)	1,188.76 (231.92)	232.98 (63.41)	4,984.95 (3,330.26)
<i>Panel B. Human capital</i>				
log NIH fund (MD)	0.1269 (0.0139)	0.1153 (0.0164)	0.0941 (0.0131)	0.1191 (0.0156)
Median MCAT (MD)	0.0666 (0.0038)	0.0814 (0.0070)	0.0413 (0.0030)	0.0797 (0.0056)
US born (foreign grad)	−0.2470 (0.0801)	0.1503 (0.1021)	0.2927 (0.0705)	0.2083 (0.0989)
Sigma (DO)	0.7944 (0.0285)	0.8845 (0.0359)	0.7275 (0.0292)	0.9321 (0.0370)
Sigma (foreign)	3.0709 (0.1102)	3.6190 (0.1469)	2.8215 (0.1131)	3.5549 (0.1411)

Notes: Detailed estimates in online Appendix Table D.1. Panel A1 presents the dollar equivalent for a 1 standard deviation change in a program characteristic. All columns include median rent in county, Medicare wage index, indicator for zero NIH funding of major associates, and for minor associates. Columns 3 and 4 include own reimbursement rates and the control variable. Point estimates using 1,000 simulation draws. Standard errors in parentheses. Optimization and estimation details described in online Appendix B.

Panel A2 presents estimates of preferences for program types and heterogeneity in preferences for program location. Both Specifications (1) and (2) estimate that, *ceteris paribus*, rural programs are preferable to urban programs. Although not

reported in the table, the typical rural program is less preferred to urban programs because rural programs tend to be associated with smaller hospitals and medical school affiliates with lower NIH funding.

Estimates from both specifications also suggest that residents prefer programs in their state of birth or in the same state as their medical school. For instance, estimates from Specification (2) imply that a typical resident is willing to forgo about \$10,000 in salary to match at a program in the same state as their medical school. These estimates are perhaps too large an estimate for moving costs alone. It may therefore be preferable to interpret these estimates as descriptive of a resident's desire to remain close to a significant other, friends and family, or due to initial sorting based on geographic preferences during admission into medical school. Although rural-born residents prefer rural programs more than other residents, they prefer rural programs at an estimated salary equivalent of under \$1,200. The estimated willingness to pay for these factors is smaller in Specification (1) although the relative importance for the different dimensions is similar.

Turning to panel B of Table 7, all specifications yield similar coefficients on the various resident characteristics and estimate that the unobservable determinants of human capital have larger variances for residents with foreign degrees. The estimated difference between a US-born foreign medical graduate and foreign graduates from other countries is an order of magnitude smaller than the standard deviation of unobservable determinants of human capital.

Estimates Using Salary Instruments.—As compared to estimates from Specifications (1) and (2), which treat salaries as exogenous, the estimated willingness to pay for program characteristics are generally larger in Specifications (3) and (4). The estimate for NIH funding of major medical school affiliates is the only exception. This increase in the estimated willingness to pay on including the instrument is driven by a fall in the coefficient on salaries but similar coefficient estimates for the other program characteristics. The estimates using salary instruments are also quite imprecise, as indicated by the typically larger standard errors. Point estimates in Specification (4) are particularly large and imprecise. The reason is that the specification results in a small, positive coefficient on salaries that is not statistically significant. The magnitudes of the instrumented specifications should be interpreted with caution given the lack of robustness, which is likely a consequence of the limited salary variation in the data.²⁶

Across specifications, however, the qualitative effect of including the salary instrument on parameter estimates is an increase in the point estimate in willingness to pay for program characteristics. This suggests that salaries are likely positively correlated with program unobservables and Specifications (1) and (2) likely underestimate the willingness to pay for more desirable programs. As discussed earlier, the direction of the bias in the estimates could not be determined based on theoretical arguments. The positive correlation between program unobservables and

²⁶The objective function for specifications using salary instruments is fairly flat along different combinations of coefficients on the salary and control variable. This indicates that the instrument may be weak. As Stock, Wright, and Yogo (2002) point out, the first stage diagnostics and computationally feasible weak-instrument robust inference methods are not well developed for nonlinear estimation techniques.

salaries could result from higher quality programs being more productive or better funded.

The lower estimated willingness to pay in Specifications (1) and (2), as we will see, will result in conservative estimates of the implicit tuition. My preferred approach is therefore to focus on Specification (2), which has the richest form of heterogeneity, for the counterfactual results and discuss the effect of the positive bias in the salary coefficient using the instrumented specifications. This approach may provide a reasonable approximation to the implicit tuition because, aside for controlled geographic covariates such as rent and wage index, estimates in column 1 of Table 6 do not show strong evidence of substantial correlation of salaries with program characteristics. Additionally, panels A1 and A2 indicate that in most cases a 95 percent confidence interval from the instrumented specification includes the point estimate from the corresponding specification that treats salaries as exogenous.

B. Model Fit and Sensitivity

In this section, I describe the in-sample and out-of-sample fit of estimates from Specification (2). The fit of the other specifications is qualitatively similar. The out-of-sample fit uses data from the 2011–2012 wave of the GME census, which was only accessed after parameter estimates were computed.

Estimates of the model only determine the probability that a resident with a given observable characteristic matches with a program with certain observables. The uncertainty in matches arises from unobservables of both the residents and the programs. Therefore, an assessment of fit must use statistics that average matches across groups of residents or programs.

For simplicity, I assess model fit using a single dimensional average quality of matched program for a group of residents with a similar observable component of human capital. I use the parameter estimates from the model to construct a quality index for each resident i and program j by computing $\mathbf{x}_i \hat{\alpha}$ and $\mathbf{z}_{ji} \hat{\beta}$, respectively. For each year t , I then divide the residents into ten bins based on $\mathbf{x}_i \hat{\alpha}$ and compute the mean quality of program with which residents from each bin are matched. Figure 1 presents a binned scatter plot of this mean quality of program as observed in the data and predicted by model simulations. Both the in-sample points and the out-of-sample points are close to the 45-degree line. The 90 percent confidence sets of the simulated means for several resident bins include the theoretical prediction.²⁷ This fit of the model provides confidence that parametric restrictions on the model are not leading to poor predictions of the sorting patterns in the market.

Additionally, one may be concerned about model misspecification due to idiosyncratic match-specific factors that are not included or due to violations from pairwise stability resulting from incomplete interviewing. Online Appendix B1 examines whether such misspecification are likely to result in large biases in the preference estimates using Monte Carlo simulations that vary the degree of misspecification. For the counterfactual considered in this paper, bias in the estimated preferences of residents may be particularly problematic. It appears that the estimates of resident

²⁷ A more model-free assessment of fit using sorting regressions only on observed covariates is presented in online Appendix Table D.2.

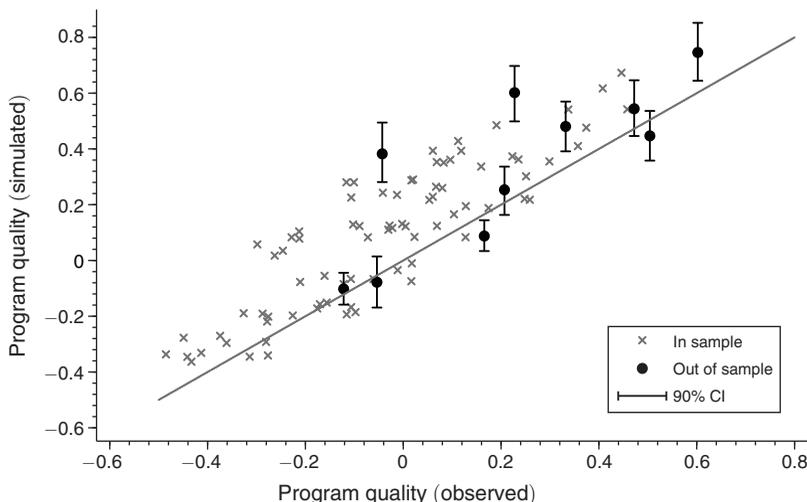


FIGURE 1. MODEL FIT: SIMULATED VERSUS OBSERVED MATCH QUALITY BY RESIDENT BINS

Notes: This scatter plot uses model estimates from specification (2) to first obtain predicted quality on observable dimensions of the residents and of the programs. Quality for the program is the “vertical component” $z_j\beta$ for the programs. The residents were binned into 10 categories, starting with foreign graduates, US-born foreign graduates, and Osteopathic graduates, and 7 quantile bins for MD graduates. The 7 MD bins are approximately equally sized, except for point masses at the cutoffs. An observation is defined at the bin-year level. Simulated means using the observed distribution of agent characteristics and 100 simulations of the unobserved characteristics. The 90 percent confidence set is constructed from these simulations.

preferences are not very sensitive to moderate levels of deviations from the assumptions made in the framework.

VII. Salary Competition

In 2002, a group of former residents filed a class-action lawsuit under the Sherman Act against the NRMP, alleging that the medical match is a competitive restraint used to depress salaries.²⁸ By replacing a traditional market in which residents could use multiple offers to negotiate with programs, they argued that the NRMP “enabled employers to obtain resident physicians without such a bidding war, thereby artificially fixing, depressing, standardizing and stabilizing compensation and other terms of employment below competitive levels.” A brief prepared by Orley Ashenfelter on behalf of the plaintiffs argued that competitive outcomes in this market would yield salaries close to the marginal product of labor, which the brief approximated using salaries of starting physicians, nurse practitioners, and physician assistants.²⁹

²⁸ Jung et al. v. AAMC et al. Class Action Complaint, No. 02-CV-00873, D.D.C. (May 5, 2002) states that “The NRMP matching program has the purpose and effect of depressing, standardizing and stabilizing compensation and other terms of employment.” The lawsuit was dismissed, overturning a previous opinion, following a provision in the Pension Funding and Equity Act of 2004 that disallowed evidence of participation in the medical match in antitrust cases.

²⁹ The expert report argues that these markets “are similar to the market for residents but operate without the anti-competitive conduct alleged against the Defendants.” A redacted copy of the report is available on request.

Physician assistants earned a median salary of \$86,000 in 2010³⁰ as compared to about \$47,000 earned by medical residents despite longer work hours.

Recent research has debated whether low salaries observed in this market are a result of the match. Using a stylized model, Bulow and Levin (2006) argue that salaries may be depressed in the match because residency programs face the risk that a higher salary may not necessarily result in matching with a better resident. Kojima (2007) uses an example to show that this result is not robust in a many-to-one matching setting because of cross-subsidization across residents in a program. Empirical evidence in Niederle and Roth (2003, 2009) suggests that medical fellowship salaries are not affected by the presence of a match. However, these studies do not explain why salaries remain low in decentralized markets, or lower than those paid to other health professionals.

The plaintiffs argued their case based on the classical economic model of homogeneous firms competing for labor in a market with free entry. However, this perfect competition benchmark may be misleading for an entry-level professional labor market. The data provide strong evidence that residents have preferences for program characteristics other than the salaries and may, thus, reject a higher salary offer from a less desirable program. Further, barriers to entry by residency programs and capacity constraints are imposed by accreditation requirements.³¹ A program must consider the option value of hiring a substitute resident when confronted with a competing salary offer. High quality programs may be particularly able to find other residents willing to work for low salaries. Conversely, highly skilled residents are scarce and they may be able to bargain for higher salaries. It is essential to consider these incentives in order to predict outcomes under competitive salary bargaining.

I model a “traditional” market using a competitive equilibrium, which is described by worker-firm specific salaries and an assignment such that each worker and firm demands precisely the prescribed assignment. Shapley and Shubik (1971) show that competitive equilibria correspond to core allocations and satisfy two conditions. First, allocations are individually rational for both workers and firms. Second, it must be that at the going salaries no worker-firm pair would prefer to break the allocation to form a (different) match at renegotiated salaries. This latter requirement ensures that further negotiations cannot be mutually beneficial. Kelso and Crawford (1982) show that competitive equilibria can result from a salary adjustment process in which the salaries of residents with multiple offers are sequentially increased until the market clears. The process embodies the “bidding war” plaintiffs suggest would arise in a “traditional” market. In fact, Crawford (2008) proposed redesigning the residency match based on this process in response to the lawsuit.

I first develop a stylized model to illustrate the dependence of competitive equilibrium salaries on both the willingness to pay for programs and the production technology of residency programs. For counterfactual simulations, I adopt an approach that does not rely on knowing the production technology of resident-program pairs

³⁰Source: Bureau of Labor Studies.

³¹The Accreditation Council for Graduate Medical Education places minimum standards on physical capital dedicated to a residency program, faculty and program personnel, hours dedicated to patient care, and diversity of patient population. Further, residency programs may not hire more residents than approved by the accreditation body.

because data on residency program output is not available. To quantify departures from the perfect competition benchmark, I use estimates of the resident preference distribution to compute the implicit tuition, which is a conservative estimate of the markdown from output net of training costs.

A. An Illustrative Assignment Model

I generalize the model of the residency market in Bulow and Levin (2006) which assumes that residents take the highest salary offer by allowing resident preferences to depend on program quality in addition to salaries and using a more flexible production function.

Consider an economy with N residents and programs in which each program may hire only one resident. Resident i has a human capital index, $h_i \in [0, \infty)$, and program j has a quality of training index, $q_j \in [0, \infty)$. To focus on salary bargaining, the training quality at a program is held fixed. Without loss of generality, index the residents and programs so that $h_i \geq h_{i-1}$, $q_j \geq q_{j-1}$, and normalize q_1 and h_1 to zero.

Residents have homogeneous, quasi-linear preferences for the quality of program, $u(q, w) = aq + w$ with $a \geq 0$. The value, net of variable training costs, to a program of quality q of employing a resident with human capital index h is $f(h, q)$ where $f_h, f_q, f_{hq} > 0$ and $f(0, 0)$ is normalized to 0. A program's profit from hiring resident h at salary level w is $f(h, q) - w$. I assume that an allocation is individually rational for a resident if $u(q, w) \geq 0$, and for a program if $f(h, q) - w \geq 0$.

A competitive equilibrium assignment maximizes total surplus. In this model, the unique equilibrium is characterized by positive assortative matching and full employment. Hence, in equilibrium, resident k is matched with program k and is paid a (possibly negative) salary w_k . The vector of equilibrium salaries is determined by the individual rationality constraints and incentive constraints that support this assignment. There is a range of salaries that are a part of a competitive equilibrium. Shapley and Shubik (1971) shows that there exists an equilibrium salary vector that is weakly preferred by all residents to all other equilibria, and another that is preferred by all programs. Since the plaintiffs alleged that salaries are currently much lower than in a "traditional" market, I focus on the worker-optimal equilibrium which has higher salaries for every worker than any other equilibrium. This outcome is unanimously preferred by all residents to other competitive equilibria, and can be solved for in this model:

PROPOSITION 1: *The worker optimal competitive equilibrium salaries are given by*

$$(14) \quad w_k = -aq_k + \sum_{i=2}^k [f(h_i, q_i) - f(h_{i-1}, q_i)].$$

PROOF:

Corollary to the equilibrium characterization in Proposition E.1 (online Appendix E1). ■

Resident 1 receives her product of labor $f(h_1, q_1)$ (normalized to 0), the maximum her employer is willing to pay. For resident 2, the first term $-aq_2$ represents an implicit price for the difference in the value of training received by her compared to that of program 1 (with $q_1 = 0$). If a resident were to use a salary offer of w by program 1 in a negotiation with program 2, the resident would accept a counter-offer of $w - aq_2$. The second term in resident 2's salary, $f(h_2, q_2) - f(h_1, q_2)$, is program 2's maximum willingness to pay for the difference in productivity of residents 1 and 2, which accrues entirely to the resident in the worker-optimal equilibrium. The sum of these two terms measures the impact of the outside option of each party on the salary negotiation determining w_2 . For $k > 2$, these (local) differences in the productivity of residents add up across lower matches to form the equilibrium salary.

B. Implicit Tuition

In models of general training that use a perfect competition framework, such as Rosen (1972) and Becker (1975), the implicit price for training is the marginal cost of training alone because free entry prevents firms from earning rents due to their quality. When entry barriers are large due to fixed costs or restrictions from accreditation requirements (see footnote 31), firms can earn additional profits due to their quality. Ruling out entry also allows us to focus on the role of salary inflexibility. Equation (14) shows that under these assumptions, program k can levy the implicit tuition aq_k on residents.³² This implicit tuition results from a force similar to compensating differentials (Rosen 1987), while allowing for heterogeneity in resident skill. Equilibrium salaries are the sum of the implicit tuition and a split of the value f produced by a resident-program pair.

As mentioned earlier, the data does not allow us to determine f but the implicit tuition can be calculated using residents' preference distribution alone. The next result shows that the implicit tuition bounds the markdown in salaries from below, implying that a gap between f and equilibrium salaries exists as a result of market fundamentals. Under free entry by firms, salaries would be equal to f because any profits earned by firms would be competed away.

PROPOSITION 2: *For all production functions f with $f_h, f_q, f_{hq} \geq 0$, the profit of firm k is bounded below by the implicit tuition aq_k in any competitive equilibrium.*

PROOF:

Corollary to Theorem E.3 stated and proved in online Appendix E2. ■

Hence, the implicit tuition aq_k is a markdown in salaries that is independent of the output. If residents have a strong preference for program quality, this implicit tuition will be large and salaries in any competitive equilibrium will be well below the product $f(h_k, q_k)$.

³²The term $f(h, q)$ can be viewed as output net of training costs. Note that the implicit price aq_k does not depend on the number of residents and programs N , which could be very large, or on the distribution of program quality. The important difference from perfect competition is that the number of firms is not disproportionately larger than the number of workers.

To build intuition, consider two particular limiting cases for the production function. If $f(h, q)$ depends only on h so that the value of a resident, denoted $\bar{f}(h)$, does not vary across programs, the worker-optimal salaries are given by

$$(15) \quad w_k = \bar{f}(h_k) - aq_k.$$

Under this production function, the resident claims the value of her labor and salaries equal her product net of the implicit tuition. Residents can engage programs in a bidding war until their salary equals the output less the implicit tuition because all programs value resident k equally.

On the other hand, if $f(h, q)$ depends only on q so that all residents produce $\underline{f}(q)$ irrespective of their human capital, the worker-optimal salaries are

$$(16) \quad w_k = \underline{f}(q_k) - aq_k.$$

In this case, programs do not share the product $\underline{f}(q_k)$ with residents since any two residents are equally productive. The resident still pays an implicit tuition for training.³³

While the production function directly influences competitive salaries, Proposition 2 shows that in all cases resident k pays the implicit tuition aq_k . Equations (15) and (16) highlight that the side of the market that owns the factor determining differences in f is compensated for their productivity. Residents are compensated for their skill if human capital is an important determinant of f . For this reason, using a production function of the form $\bar{f}(h)$ results in a markdown in salaries from f that is only due to the implicit tuition.

The key difference from the results derived using models with free entry is because firms enter and bid for labor services until a zero profit condition is met. Hence, compensation close to marginal productivity of labor results from free entry rather than negotiations between a fixed set of agents.

C. Generalizing the Implicit Tuition

The expression for the implicit tuition derived above relied on the assumption that residents have homogeneous preferences for program quality and do not directly speak to competitive outcomes in a model with heterogeneous preferences. This section generalizes the definition of implicit tuition to make it applicable to the estimated model.

Notice that the profit earned by program k in a worker-optimal equilibrium under a production function of the form $f(h)$ is precisely the implicit tuition aq_k because this production function does not provide programs with infra-marginal productive rents. Under this production function, markdowns from output are determined only by residents' preferences for programs. Consequently, calculating firm profits using a production function of this form provides a conservative approach to estimating payoffs to programs. The next result shows that even under heterogeneous

³³To ensure assortative matching in these limiting cases, I assume that a program (resident) with equally attractive offers breaks the tie in favor of the resident (program) with the higher human capital (quality).

preferences for programs, the difference between salaries and output is the same for all production functions of the form $f(h)$, circumventing the need for estimating f .

I state the result for a one-to-one assignment model, and the general result for many-to-one setting is stated and proved in online Appendix E3.³⁴ With a slight abuse of notation, let the total surplus from the pair (i, j) be $a_{ij}^f = u_{ij} + f(h_i) \geq 0$.³⁵ The equilibrium for a modified assignment game in which the surplus produced by the pair is $a_{ij}^{\tilde{f}} = u_{ij} + \tilde{f}(h_i) \geq 0$ can be characterized terms of the equilibria of the game with surplus a_{ij}^f as follows:

PROPOSITION 3: *The equilibrium assignments of the games defined by a_{ij}^f and $a_{ij}^{\tilde{f}}$ coincide. Further, if u_i^f and v_j^f are equilibrium payoffs for the surplus a_{ij}^f , then $u_i^{\tilde{f}} = u_i^f + \tilde{f}(h_i) - f(h_i)$ and $v_j^{\tilde{f}} = v_j^f$ are equilibrium payoffs under the surplus $a_{ij}^{\tilde{f}}$. Hence, a firm's profit in a worker-optimal equilibrium depends on $\{u_{ij}\}_{i,j}$ but is identical across production functions of the form $f(h)$.*

PROOF:

See online Appendix E3 for the proof of the general case with many-to-one matches. ■

As in the illustrative model, under a production technology that depends only on human capital, the residents are the residual claimants of output. A change in the productivity of human capital is reflected in the salaries, one for one. The firms' profits depends only on the preferences of the residents. Thus, I refer to the difference between output and salaries in the worker-optimal competitive equilibrium for a model in which f depends only on h as the implicit tuition. This definition uses the assumption that preferences of the programs can be represented using a single human capital index in the empirical model but also makes the additional restriction that the productivity of human capital, in dollar terms, does not depend on the identity of the program.

Since a closed form expression for competitive equilibrium salaries is not available with heterogeneous resident preference, I calculate the implicit tuition implied by estimated preferences using a two-step procedure.³⁶ Each step solves a linear program based on the approach developed in Shapley and Shubik (1971):

- **Step 1:** Solve the optimal assignment problem, modified from the formulation by Shapley and Shubik (1971) to allow for many-to-one matching.

³⁴In the general formulation, I assume that the total output from a team of residents (h_1, \dots, h_{q_j}) is $F(h_1, \dots, h_{q_j}) = \sum_{k=1}^{q_j} f(h_k)$, where $f(h_k) = 0$ if position k is not filled.

³⁵This formulation implicitly assumes that, at every program, it is individually rational for a worker to accept a salary equal to her product. It further assumes that the output of every resident is nonnegative.

³⁶Since the total number of positions in the market exceeds the number of residents and the value of options outside the residency market are difficult to determine, I assume that the equilibrium is characterized by full employment. This property follows if, for instance, it is individually rational for all residents to be matched with their least desirable program at a salary equal to the value of the resident to this program.

- **Step 2:** Calculate the worker-optimal element in the core given the assignments above.

Online Appendix E4 describes the procedure in more detail. All calculations are done with the 2010–2011 sample of the data.

D. *Estimates of Implicit Tuition*

Estimates presented in Section VI suggest that residents are willing to take large salary cuts in order to train at preferred programs, which can translate into a large implicit tuition. Table 8 presents summary statistics of the distribution of implicit tuition using estimates from various specifications. I estimate the average implicit tuition to be about \$23,000 for Specifications (1) and (2). As mentioned in Section VI, the results from Specifications (3) and (4) indicate that salaries are positively correlated with program unobservables. We therefore expect larger estimates of the implicit tuition with these specifications, but ones that are imprecise since the instrument appears weak. Specification (3) yields an estimated average implicit tuition of \$43,500 but with a larger standard error of \$13,700. Nonetheless, this specification rules out an average implicit tuition smaller than \$17,000. These estimates are economically large in comparison to the mean salary of about \$47,000 paid to residents. Specification (4) results in a still larger point estimate of \$120,000. The standard error cannot be computed for this specification because the coefficient on salary was not significantly different from zero.³⁷ Given that the salaries in the data range from about \$31,000 to \$65,000, these calculations, particularly those with Specification (4), involve significant parametric extrapolations on resident choices.

The results also show significant dispersion in the implicit tuition across residents and programs. The standard deviation in the implicit tuition from Specifications (1)–(3) is between \$12,000 and \$25,000. The seventy-fifth percentile of implicit tuition can be about three times higher than the twenty-fifth percentile, with even higher values at the ninety-fifth percentile. This dispersion arises from the differences in program quality, which allows higher quality programs to lower salaries more than relatively low quality programs.

The average estimated implicit tuition is upward of 50 percent of the \$40,000 salary difference between medical residents and physician assistants. This finding refutes the plaintiffs' argument that the salary gap would not exist if residents' salaries were set competitively and physician assistant salaries approximated the productivity of residents. However, the lower end of the estimates cannot explain the salary gap between starting physicians and medical residents, which is approximately \$90,000.³⁸ As discussed earlier, the implicit tuition is a conservative estimate of

³⁷I used a simulation approach to compute the standard errors because a delta method using finite difference derivatives is computationally prohibitive as an estimate for each simulation involves solving a large linear program. To assess whether a one-sided confidence interval using Specification (4) rules out small estimates of the average implicit tuition, I simulated the distribution of the estimated average implicit tuition by truncating the asymptotic distribution of the salary coefficient below at 0.01 (the point estimate is 0.498 and the standard error is 0.317). This truncation should affect the right tail of the estimates, but not the lower end of the estimates. The estimated one-sided 95 percent confidence interval places a lower bound of \$39,631 on the average implicit tuition.

³⁸I use Mincer equation estimated using interval regressions on confidential data from the Health Physician Tracking Survey of 2008 to calculate the average salaries for starting family physicians. Details available on request.

TABLE 8—IMPLICIT TUITION

	Geographic heterogeneity (1)	Full heterogeneity (2)	Geographic heterogeneity with instrument (3)	Full heterogeneity with instrument (4)
Mean	\$22,625.38 (3,498.41)	\$23,753.54 (5,461.19)	\$43,470.39 (13,678.08)	\$119,622.30 —
Median	\$21,166.40 (3,266.20)	\$21,194.27 (4,985.89)	\$40,606.85 (12,847.51)	\$107,689.70 —
Standard deviation	\$12,279.77 (1,779.57)	\$16,651.24 (3,932.44)	\$24,792.30 (7,485.20)	\$80,400.55 —
Fifth percentile	\$5,160.78 (1,449.57)	\$2,737.88 (982.41)	\$7,912.03 (3,246.19)	\$14,480.97 —
Twenty-fifth percentile	\$14,069.19 (2,365.88)	\$11,618.22 (2,780.80)	\$24,853.10 (8,299.05)	\$58,470.95 —
Seventy-fifth percentile	\$28,901.58 (4,352.64)	\$31,439.99 (6,991.71)	\$58,354.66 (18,134.03)	\$162,742.00 —
Ninety-fifth percentile	\$45,784.76 (6,921.96)	\$55,165.06 (1,2591.74)	\$92,343.91 (2,8071.67)	\$273,583.30 —

Notes: Based on 100 simulation draws. For columns 1–3, each simulation draws a parameter from a normal with mean $\hat{\theta}_{SMD}$ and variance $\hat{\Sigma}$, where $\hat{\Sigma}$ is estimated as described in Section IIB. Column 4 is computed at $\hat{\theta}_{SMD}$. All columns simulate the unobservables ξ_{ji} , ε_i , and β_i and use the 2010 sample of residents and programs. Standard errors in parentheses.

the salary markdown and part of this salary gap may be due to differences in the productivity of medical residents and starting physicians.

One can also interpret the implicit tuition as a bound on the markdown from the marginal productivity if the medical residency market were to adopt the flexible-salary match proposed by Crawford (2008) since the suggestion is intended to implement competitive equilibrium outcomes. The results suggest that a large gap between salaries and productivity would still remain even if salaries were flexible and set competitively.

VIII. Conclusion

While preferences of agents determine outcomes in matching markets, a common constraint is that only data on employer-employee matches or student enrollment records, rather than stated preferences, are available. This paper develops an empirical framework for estimating the distribution of preferences of agents in large two-sided markets with non-transferable utility using only data on final matches. I use pairwise stability together with a vertical preference restriction on one side of the market to estimate preference parameters using simulated minimum distance. The empirical strategy is based on using moments from the sorting patterns observed in the data and information available only in a market with many-to-one matching.

I use these methods to quantitatively analyze whether centralization in the medical residency market is responsible for low salaries. I find that heterogeneity in program types and capacity constraints result in quantitatively large departures from the perfect competition model suggested by the plaintiffs in the lawsuit. Theoretical

results presented in Section VII show that equilibrium salaries can be well below the product of labor (net of training costs) when residents value the quality of a program. Counterfactual estimates show that the willingness to pay for programs results in salary markdowns (implicit tuition) upward of \$23,000 in any competitive equilibrium. My estimates also show that higher quality programs would earn a larger implicit tuition than less desirable programs. To the extent that higher quality programs are matched with higher skilled residents, the implicit tuition is a countervailing force to high dispersion salaries driven by productivity differences. The implicit tuition may therefore explain the empirical observations of Niederle and Roth (2003, 2009) in fellowship markets.

The result suggests that the limited supply of heterogeneous residency positions, due to barriers to entry such as accreditation requirements, can cause significant salary depression, and weighs against the view the match is primarily responsible for low resident salaries. In this market, salaries may also be influenced by the previously mentioned guideline requiring minimum financial compensations for residents. These forces are not directly linked to the presence of a match.

REFERENCES

- Agarwal, Nikhil.** 2015. "An Empirical Model of the Medical Match: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer20131006>.
- Agarwal, Nikhil, and William Diamond.** 2014. "Identification and Estimation in Two-sided Matching Markets." Yale University Cowles Foundation Discussion Paper 1905.
- Akkus, Oktay, J. Anthony Cookson, and Ali Hortacsu.** 2015. "The Determinants of Bank Mergers: A Revealed Preference Analysis." <http://ssrn.com/abstract=2459727> (accessed March 20, 2015).
- Becker, Gary S.** 1975. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. 2nd ed. New York: National Bureau of Economic Research.
- Berry, Steven T.** 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics* 25 (2): 242–62.
- Berry, Steven, Amit Gandhi, and Philip Haile.** 2013. "Connected Substitutes and Invertibility of Demand." *Econometrica* 81 (5): 2087–2111.
- Berry, Steven T., and Philip A. Haile.** 2014. "Identification in Differentiated Products Markets Using Market Level Data." *Econometrica* 82 (5): 1749–97.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–90.
- Berry, Steven, and Ariel Pakes.** 2007. "The Pure Characteristics Demand Model." *International Economic Review* 48 (4): 1193–1225.
- Blundell, Richard, and James L. Powell.** 2003. "Endogeneity in Nonparametric and Semiparametric Regression Models." In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, vol. 2, edited by Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, 312–57. New York: Cambridge University Press.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff.** 2013. "Analyzing the Determinants of the Matching of Public School Teachers to Jobs: Disentangling the Preferences of Teachers and Employers." *Journal of Labor Economics* 31 (1): 83–117.
- Breusch, T. S., and A. R. Pagan.** 1979. "A Simple Test for Heteroscedasticity and Random Coefficient Variation." *Econometrica* 47 (5): 1287–94.
- Bulow, Jeremy, and Jonathan Levin.** 2006. "Matching and Price Competition." *American Economic Review* 96 (3): 652–68.
- Chiappori, Pierre-André, Bernard Salanié, and Yoram Weiss.** 2015. "Partner Choice and the Marital College Premium: Analyzing Marital Patterns over Several Decades." CEPR Discussion Paper DP10403. <http://ssrn.com/abstract=2562928> (accessed February 1, 2015).
- Choo, Eugene, and Aloysius Siow.** 2006. "Who Marries Whom and Why." *Journal of Political Economy* 114 (1): 175–201.
- Clark, Simon.** 2006. "The Uniqueness of Stable Matchings." *B. E. Journal of Theoretical Economics: Contributions to Theoretical Economics* 6 (1): 1–30.

- Crawford, Vincent P.** 2008. "The Flexible-Salary Match: A Proposal to Increase the Salary Flexibility of the National Resident Matching Program." *Journal of Economic Behavior and Organization* 66 (2): 149–60.
- Fox, Jeremy T.** 2009. "Matching Models: Empirics." In *The New Palgrave Dictionary of Economics, Online Edition*, edited by Steven N. Durlauf and Lawrence E. Blume. New York: Palgrave Macmillan. http://www.dictionaryofeconomics.com/article?id=pde2009_M000426 (accessed May 3, 2015).
- Fox, Jeremy T.** 2010. "Estimating Matching Games with Transfers." www.personal.umich.edu/~jtfox/working-papers/matching_fox.pdf (accessed May 3, 2015).
- Fox, Jeremy T., and Patrick Bajari.** 2013. "Measuring the Efficiency of an FCC Spectrum Auction." *American Economic Journal: Microeconomics* 5 (1): 100–46.
- Gale, David, and Lloyd S. Shapley.** 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly* 69 (1): 9–15.
- Galichon, Alfred, and Bernard Salanié.** 2015. "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models." <http://ssrn.com/abstract=1804623> (accessed February 18, 2015).
- Gordon, Nora, and Brian Knight.** 2009. "A Spatial Merger Estimator with an Application to School District Consolidation." *Journal of Public Economics* 93 (5-6): 752–65.
- Gouriéroux, Christian, and Alain Monfort.** 1997. *Simulation-Based Econometric Methods*. New York: Oxford University Press.
- Heckman, James J., and Richard Robb, Jr.** 1985. "Alternative Methods for Evaluating the Impact of Interventions: An Overview." *Journal of Econometrics* 30 (1-2): 239–67.
- Imbens, Guido W.** 2007. "Nonadditive Models with Endogenous Regressors." In *Advances in Economics and Econometrics: Theory and Applications*, vol. 3, edited by Richard Blundell, Whitney Newey, and Torsten Persson, 17–46. New York: Cambridge University Press.
- Imbens, Guido W., and Whitney K. Newey.** 2009. "Identification and Estimation of Triangular Simultaneous Equations Models without Additivity." *Econometrica* 77 (5): 1481–1512.
- Kelso, Alexander S., Jr., and Vincent P. Crawford.** 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica* 50 (6): 1483–1504.
- Kojima, Fuhito.** 2007. "Matching and Price Competition: Comment." *American Economic Review* 97 (3): 1027–31.
- Kojima, Fuhito, Parag A. Pathak, and Alvin E. Roth.** 2013. "Matching with Couples: Stability and Incentives in Large Markets." *Quarterly Journal of Economics* 128 (4): 1585–1632.
- Logan, John Allen, Peter D. Hoff, and Michael A. Newton.** 2008. "Two-Sided Estimation of Mate Preferences for Similarities in Age, Education, and Religion." *Journal of the American Statistical Association* 103 (482): 559–69.
- McFadden, Daniel.** 1989. "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration." *Econometrica* 57 (5): 995–1026.
- Menzel, Konrad.** Forthcoming. "Large Matching Markets as Two-Sided Demand Systems." *Econometrica*.
- National Resident Matching Program (NRMP).** 2006, 2007, 2009, 2011. "Charting Outcomes in the Match." <http://www.nrmp.org/match-data/nrmp-historical-reports> (accessed June 12, 2012).
- Newhouse, Joseph P., and Gail R. Wilensky.** 2001. "Paying For Graduate Medical Education: The Debate Goes On." *Health Affairs* 20 (2): 136–47.
- Niederle, Muriel, and Alvin E. Roth.** 2003. "Relationship between Wages and Presence of a Match in Medical Fellowships." *Journal of the American Medical Association* 290 (9): 1153–54.
- Niederle, Muriel, and Alvin E. Roth.** 2009. "The Effects of a Centralized Clearinghouse on Job Placement, Wages, and Hiring Practices." In *Studies of Labor Market Intermediation*, edited by David Autor, 235–71. Chicago: University of Chicago Press.
- Niederle, Muriel, and Leeat Yariv.** 2009. "Decentralized Matching with Aligned Preferences." National Bureau of Economic Research Working Paper 14840.
- Pakes, Ariel, and David Pollard.** 1989. "Simulation and the Asymptotics of Optimization Estimators." *Econometrica* 57 (5): 1027–57.
- Petrin, Amil, and Kenneth Train.** 2010. "A Control Function Approach to Endogeneity in Consumer Choice Models." *Journal of Marketing Research* 47 (1): 3–13.
- Pycia, Marek.** 2012. "Stability and Preference Alignment in Matching and Coalition Formation." *Econometrica* 80 (1): 323–62.
- Rosen, Sherwin.** 1972. "Learning and Experience in the Labor Market." *Journal of Human Resources* 7 (3): 326–42.
- Rosen, Sherwin.** 1987. "The Theory of Equalizing Differences." In *Handbook of Labor Economics*, vol. 1, edited by Orley C. Ashenfelter and Richard Layard, 641–92. New York: Elsevier.
- Roth, Alvin E.** 1984. "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy* 92 (6): 991–1016.

- 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 (4): 748–80.
- Roth, Alvin E., and Marilda Sotomayor.** 1992. *Two-Sided Matching. Handbooks in Economics*, vol. 11. New York: Elsevier Science.
- Roth, Alvin E., and Xiaolin Xing.** 1994. "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions." *American Economic Review* 84 (4): 992–1044.
- Shapley, L. S., and M. Shubik.** 1971. "The Assignment Game I: The Core." *International Journal of Game Theory* 1 (1): 111–30.
- Signer, Mona M.** 2012. "2012 Main Residency Match and SOAP." National Resident Matching Program (accessed June 12, 2012).
- Sorensen, Morten.** 2007. "How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital." *Journal of Finance* 62 (6): 2725–62.
- Stock, James H., Jonathan H. Wright, and Motohiro Yogo.** 2002. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments." *Journal of Business and Economic Statistics* 20 (4): 518–29.
- Uetake, Kosuke, and Yastora Watanabe.** 2012. "Entry by Merger: Estimates from a Two-Sided Matching Model with Externalities." <http://ssrn.com/abstract=2188581> (accessed August 12, 2012).
- Weese, Eric.** Forthcoming. "Political Mergers as Coalition Formation: An Analysis of the Heisei Municipal Amalgamations." *Quantitative Economics*.
- White, Halbert.** 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica* 48 (4): 817–38.
- Yang, Sha, Yuxin Chen, and Greg M. Allenby.** 2003. "Bayesian Analysis of Simultaneous Demand and Supply." *Quantitative Marketing and Economics* 1 (3): 251–75.