

DISCUSSION PAPER SERIES

IZA DP No. 11230

**Gender Homophily in Referral Networks:
Consequences for the Medicare Physician
Earnings Gap**

Dan Zeltzer

DECEMBER 2017

DISCUSSION PAPER SERIES

IZA DP No. 11230

Gender Homophily in Referral Networks: Consequences for the Medicare Physician Earnings Gap

Dan Zeltzer

Tel Aviv University and IZA

DECEMBER 2017

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Gender Homophily in Referral Networks: Consequences for the Medicare Physician Earnings Gap*

In this paper, I assess the extent to which the gender gap in physician earnings may be driven by physicians' preference for working with specialists of the same gender. By analyzing administrative data on 100 million Medicare patient referrals, I provide robust evidence that doctors refer more to specialists of their same gender, a tendency known as *homophily*. I propose a new measure of homophily that is invariant to differences between the genders in the propensity to refer or receive referrals. I show that biased referrals are predominantly driven by physicians' decisions rather than by endogenous sorting of physicians or patients or by gender differences in the labor supply. As 75% of doctors are men, estimates suggest biased referrals generate a 5% lower demand for female relative to male specialists, pointing to a positive externality for increased female participation in medicine.

JEL Classification: I11, J16, L14

Keywords: networks, referrals, gender, physician markets

Corresponding author:

Dan Zeltzer

Tel Aviv University

School of Economics

Tel Aviv

Israel 6997801

E-mail: dzeltzer@tauex.tau.ac.il

* I especially thank Janet Currie, Ilyana Kuziemko, Bo Honore, and Sylvain Chassang for invaluable advice throughout the research process. For helpful comments and discussions, I thank Liran Einav, Matt Jackson, Alexandre Mas, Jessica Pan, Aureo de Paula, Tom Vogl, Juan Pablo Xandri, and participants at the Econometric Society Meeting in Milan, the University of Pennsylvania, The University of Texas at Austin, The University of Maryland, Einaudi Institute for Economics and Finance, the IDC, Ben-Gurion University, the University of Haifa, The Hebrew University, Tel Aviv University, the AEA Meeting in Chicago, University of Vienna, and the Barcelona GSE Summer Forum. I am also grateful to Jean Roth and Mohan Ramanujan for their assistance in obtaining and managing the data. Drs. Bon Ku, Maria Maguire, and Steven Vogl provided helpful practitioners' perspectives. Financial support for this research was provided by the Program for U.S. Healthcare Policy Research of the Center For Health and Wellbeing at Princeton University.

1 Introduction

It remains unclear why women are still underrepresented in top career positions and earn less than their male counterparts. Persistent gender earnings gaps have attracted particular attention in research and policy debates, and many channels for explaining them have been proposed (for surveys see Blau and Kahn, 2017; Niederle and Vesterlund, 2007; Bertrand, 2011; Azmat and Petrongolo, 2014). In this paper, I study in detail one such channel: biased professional networks.

It is well documented that individuals tend to develop relationships with others similar to themselves, a tendency known as *homophily*. This paper studies the extent to which this tendency contributes to the gender earnings gap in medicine, through its effect on patient referral volumes, and therefore on the demand for female physicians. Using Medicare data on 100 million patient referrals among half a million U.S. physicians, I find that doctors refer more to specialists of their same gender.¹ Furthermore, I show that this gender bias in referrals is predominantly driven by doctors' decisions rather than by endogenous sorting of physicians or patients. My results imply that because most U.S. physicians are men, gender-biased referrals make demand for female specialists lower than demand for male specialists, contributing substantially to the earnings gap among physicians. I directly test this implication and show that it cannot be explained by differences between the genders in labor supply.

Medicare data are particularly suited for studying the contribution of biased referral networks to the gender earnings gap, for several reasons. First, the Medicare reimbursement structure is fixed, so differences in earnings are due solely to the quantity of work, not wage differentials. Second, referrals and payments are jointly observed, directly revealing the impact of doctors' referral decisions on specialists' workload and revenue. Third, observing the gender of both doctors and specialists allows me to isolate the demand-side gender bias in referrals from a range of other observed and unobserved differences between the genders, such as differences in labor supply, because I can see not only whether men work more or

¹Throughout this paper the terms *doctor* and *specialist* are used to denote the role of a physician as a referral origin or target, regardless of their medical specialties, which are accounted for in the empirical analysis. In practice, most referrals are from primary care to other specialties.

receive more referrals, but also whether referrals originating from men are disproportionately targeted towards men. Fourth, rich patient and physician data allow me to account for multiple dimensions of potential sorting that could result in a correlation between the gender of the referring and receiving physicians, including distance and shared institutional affiliations. Finally, longitudinal data allow me to test whether biased referrals actually translate into disparity in earnings.

To account for gender differences that may make male specialists send or receive more referrals than their female counterparts, such as differences in labor supply related to men working longer hours, I define a new homophily measure. This measure, *directed homophily*, compares the fraction of referrals made to male specialists between male and female *doctors*. I show that if referral rates were fully explained by supply-side differences between male and female specialists, then referral decisions should be independent of the gender of the doctor making them, all else being equal. Namely, there should be no directed homophily.² Labor supply differences are only one of many potential differences between the genders accounted for by this approach to measuring homophily. For example, male physicians are also on average eight years more experienced than female physicians, as female entry to medicine is rather recent.

I find that across the U.S., female doctors refer to female specialists a third more than male doctors do (19% women-to-women compared with 15% men-to-women—a 4 percentage-point difference; the difference has an age gradient, and is slightly bigger among older doctors). Specifically, Medicare referrals exhibit directed homophily. Homophily estimates decline only modestly even when fixed-effects are used to narrow the comparison to that between doctors of the same specialty, within the same hospital, and with similar patients.

However, evaluating the contribution of doctors' choices to these gender differences in referral rates further requires accounting for potential sorting. Because different doctors choose specialists from different pools, if physicians endogenously sorted by gender into market segments (e.g. hospitals or medical specialties), doctors would be more exposed to—and thus more likely to refer

²Oft-used homophily measures disregard this concern, as highlighted in Graham (2014). Cf., Coleman (1958); Currarini, Jackson and Pin (2009). A similar approach is used for studying discrimination by Anwar and Fang (2006) and Antonovics and Knight (2009).

to—specialists of their own gender. Even absent sorting, quantifying the underlying bias in referrals based on observed doctors’ choices requires accounting for how restricted these choices were (e.g., even a doctor preferring same-gender specialists may end up referring mostly to specialists of the opposite gender because of availability constraints).

I therefore develop and estimate a discrete choice model, where doctors choose specialists from local pools. This model allows for—and exploits—variation between doctors’ choice sets. The determinants of referral relationships are identified by comparing the characteristics of chosen specialists with those of unchosen ones. This model serves two purposes. First, it is used as an empirical framework for estimating the gender bias in referrals, all other things being equal. Second, it is used as a quantifying framework for assessing the impact of this bias on gender disparity in specialist demand. To overcome the computational hurdle of estimating the model with very large potential choice sets, I use choice-based sampling (McFadden, 1984; Manski and Lerman, 1977).

Estimates of the discrete choice model suggest that most of the bias in referrals is due to doctors’ choices: faced with otherwise similar male and female specialists, doctors are 10% more likely to refer to the one of their own gender, even controlling for the fact that doctors refer more to specialists with whom they have common institutional affiliations, who are nearby, who are of similar experience, and who have attended their same medical school. Additionally, using longitudinal data I find that referral relationships between physicians of the same gender are more persistent than referral relationships between physicians of the opposite gender.

Turning to the the impact of biased referrals on earnings, I show that because most referrals are currently made by men, biased referrals imply that more patients are referred to male specialists and fewer to female specialists. To quantify the impact of the estimated bias in referrals on specialists’ earnings, I use the model to calculate the average demand, by gender, given the bias and the gender composition of the current U.S. physician population. As 75% of referring physicians are men, the estimated 10% bias in referral probabilities makes demand for female specialists 5% lower than demand for male specialists, all else being equal. This gender disparity in demand would vanish if doctors’ choices were unbiased,

or, alternatively, if their population were gender-balanced.

I show that the evidence can rule out several potential confounding factors: gender differences in labor supply, physician sorting, and patient sorting. Gender differences in labor supply, while an important contributor to gender differences in earnings in many other settings, fail to explain the observed correlation between the gender of doctors and specialists. Specifically, supply-side differences in workload cannot explain why male doctors refer a greater *fraction* of their patients to male specialists, and vice versa. Physician sorting is another concern. However, sorting on multiple observed factors, including location, institutional affiliation, and medical school attended is directly accounted for by the discrete choice model. To further check the robustness of my results to sorting on unobservable dimensions and to unobserved heterogeneity in labor supply, I directly test whether specialists' earnings depend on how their gender matches that of nearby doctors. I find that they do.³ This evidence also suggests that female workload varies with the fluctuations in demand, implying that female specialists are not capacity constrained. A separate concern is that referrals could appear biased if patients, rather than physicians, exhibit homophily, or if more complex cases are more likely to be seen by male physicians. However, accounting for the gender of patients and their age—a strong predictor of case complexity—fail to account for nearly any of the observed bias in referrals. I therefore argue that the totality of the evidence suggests that biased referrals lead to a lower stream of incoming referrals to female doctors, and thus, their livelihood is hurt.⁴

Previous research has shown that gender gaps in earnings exist among highly skilled occupations, such as those in the financial and corporate sectors (Bertrand, Goldin and Katz, 2010) and the legal profession (Azmat and Ferrer, 2016). Such gaps often remain unexplained even when gender differences in many individual

³Specifically, I construct a monthly panel of physician payments and test the correlation of specialist monthly payments with the gender composition of nearby primary-care doctors. I find that the fraction of primary-care claims handled by male physicians is positively correlated with male specialists' earnings and negatively correlated with female specialist earnings in the same market, even when specialist-fixed effects and controls for patient gender are used to account for heterogeneity in supply and for patient sorting, respectively.

⁴Note that this fact does not imply that referrals are discriminatory, as there could be alternative explanations for biased referrals, not all implying discriminatory preferences. For example, men may be more familiar with or feel more comfortable working with other men. While not discriminatory per se, these mechanisms are still detrimental for the minority gender.

characteristics are accounted for, particularly in medicine (Baker, 1996; Weeks, Wallace and Wallace, 2009; Lo Sasso et al., 2011; Esteves-Sorenson, Snyder et al., 2012; Seabury, Chandra and Jena, 2013). In my Medicare sample, the raw earnings gap between male and female physicians is 66 log points. Individual characteristics, including specialty, experience, medical school attended, and no-work spells, explain half of this initial gap. Biased referrals explain approximately 15% of the remaining 33 log points gap—the within-specialty gender gap in workload. My analysis therefore shows that a substantial part of the previously unexplainable gender earnings gap among physicians is explained by disparity in demand due to gender-biased referrals putting women at a disadvantage relative to otherwise similar men. Recent work by Sarsons (2017), which builds on and extends the current framework, further suggests that in the case of referrals to surgeons, referring doctors interpret patient outcomes differently, depending on the surgeon’s gender.

The main implication for medicine is that women in specialties that make the most referrals, such as primary care, induce positive externalities for women in other specialties by increasing the demand for their services through referrals. This effect is particularly important in specialties in which much of the work depends on referrals, such as most surgical specialties—an area in which there are currently very few women. Furthermore, although eventually the part of the earnings gap due to homophily is expected to vanish or even reverse as recent female entrants gradually transform the gender composition of the physician labor force, homophily is still a hindrance to pay convergence. The contribution of homophily to the earnings gap could be even larger if demand disparities also discourage female entry to higher-paying specialties.

Results shed light on gender differences more generally, highlighting a mechanism that may limit the earnings and career opportunities for women in other contexts, particularly where networking is important. These results thus relate to three existing lines of work: the influence of networks, homophily, and earnings inequality. Networks are shown to influence hiring (Hellerstein, McInerney and Neumark, 2011; Burks et al., 2014), compensation (Renneboog and Zhao, 2011), access to freelance jobs (Ghani, Kerr and Stanton, 2014), venture capital funding (Hochberg, Ljungqvist and Lu, 2007), and managerial positions (Zimmerman,

2014), among other things. Separately, homophily has been widely documented in various networks and in different dimensions such as age, race, and political attitudes (McPherson, Smith-Lovin and Cook, 2001), and has been shown to contribute to network segregation and limit the flow of information, both in theory (Bramoullé et al., 2012*a*; Golub and Jackson, 2012), and in practice (Currarini, Jackson and Pin, 2009; Himmelboim, McCreery and Smith, 2013; Halberstam and Knight, 2016). This paper demonstrates that homophily may bias the professional interactions between individuals and can provide a key to understanding propagation and perpetuation of gender inequalities in medicine and beyond.

This paper proceeds as follows: Section 2 describes the data and decomposes the gender earnings gap among physicians in Medicare. Section 3 discusses the homophily measure and documents homophily patterns in Medicare referrals. Section 4 develops the model and uses it to estimate the underlying gender bias in physicians' choices. Section 5 studies the impact of this bias on earnings disparity. Section 6 discusses alternative explanations and implications. Section 7 concludes.

2 Data and Background

2.1 Data Sources

The main data source for this study is the Carrier database, a panel of all physician-billed services for a random sample of 20% of Medicare beneficiaries for 2008–2012.⁵ Data encode the gender of doctors, specialists, and patients, as well as payments, and are linked to rich data on physician characteristics. The sample contains patients with traditional fee-for-service Medicare, which accounts for two-thirds of all Medicare beneficiaries, with a total of 35 million covered persons, and more than half a million doctors, all across the U.S.

I use a confidential version of the data, which contains both payment and referral information for each claim, and thus allows for studying homophily and

⁵Medicare is the federal health insurance program for people who are age 65 or older, certain younger people with disabilities, and people with end-stage renal disease. It is run by a government agency, the Centers for Medicare and Medicaid (CMS).

its impact on pay disparities.⁶ For each encounter of a patient with a physician, the data contain the following: the date of service and its location, the type of service, patient gender, the physician specialty, and payments made to the physician by all payors; data also record the referring provider, if there was one. Non-physician providers (such as nurse practitioners) are excluded, based on CMS specialty code. A small number of services are excluded, such as lab tests, which are often ordered by physicians directly, in which case the ordering physician is reported instead of the referring physician.⁷ To protect the privacy of patients, no reported statistics are based on fewer than 11 individual patients. Thanks to the large sample size, such cells are never encountered.

Data on physician gender and other characteristics were obtained from Physician Compare, a public CMS database that tracks information on physicians who provide Medicare services.⁸ The included characteristics are: sex, specialty, hospital and group practice affiliations, medical school attended, and year of graduation (used to calculate experience). Panel A of Table 1 summarizes physician characteristics (for additional descriptive information see the supplementary materials). These data are further combined with beneficiary gender and summary cost and utilization from the Master Beneficiary Summary File.

The sample is fairly representative of U.S. physicians, as more than 90% of U.S. physicians provide Medicare services.⁹ By volume, Medicare-billed physician

⁶For a detailed description of these data see “Carrier RIF Research Data Assistance Center (ResDAC),” <http://www.resdac.org/cms-data/files/carrier-rif>. Accessed December 2017.

⁷Claims are reported using CMS Health Insurance Claim Form 1500, which contains a fields (17, 17a) for the name and identifier of the referring or ordering provider. For details see CMS Claims Processing Manual (Rev. 3103, 11-03-14) Chapter 26, 10.4, Item 17. Services are excluded with BETOS codes for Tests, Durable Medical Equipment, Imaging, Other, and Unclassified Services. For a detailed description of these codes see <https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/MedicareFeeForSvcPartsAB/downloads/BETOSDescCodes.pdf>. Accessed December 2017. About a third of the remaining claims record a referring physician provider.

⁸Physician Compare Database, <https://data.medicare.gov/data/physician-compare> Accessed December 2017.

⁹In a national representative survey of non-pediatrician primary care providers conducted by the Kaiser Family Foundation and The Commonwealth Fund in 2015, 94% of male and 93% of female physicians reported seeing Medicare patients. These fraction are similar to the 91% percent of all office-based physicians who report accepting new Medicare patients in the National Ambulatory Medical Care Survey (NAMCS) in 2012. The percentage of physicians reported as accepting new Medicare patients in the NAMCS was also similar to the percentage accepting

Table 1: Descriptive Statistics: Physicians and Referrals

	All	Men	Women
A. All Physicians			
Male Physician	0.723		
Experience (years)	22.4	24.2	17.9
Patients*	311	346	219
Claims*	755	850	515
Pay*	\$106,112	\$121,997	\$64,620
Obs. (All Physicians)	530,357	383,525	146,832
B. Doctors (any outgoing referrals)			
Male Physician	0.734		
Avg. Outgoing Referral Volume*	\$43,925	\$48,315	\$31,810
Fraction Male Patients	0.430	0.463	0.339
Links (out-degree)	16.2	17.1	13.7
Outgoing Referrals to Men:	0.834	0.848	0.795
Obs. (Doctors)	383,173	281,238	101,935
C. Specialists (any incoming referrals)			
Male Physician	0.755		
Avg. Incoming Referral Volume*	\$48,730	\$55,405	\$28,155
Fraction Male Patients	0.412	0.433	0.348
Links (in-degree)	18.0	19.9	12.3
Incoming Referrals from Men:	0.777	0.795	0.719
Obs. (Specialist)	345,390	260,795	84,595

Notes: 20% sample of patients; *volume variable, multiplied by 5 to adjust for sampling; Physician demographics and average work volume are for all sampled physicians (Part A). Referred work volume (Parts B and C) are for Doctors and Specialists, namely physicians with at least one outgoing referral (Part B) or incoming referral (Part C) and complete demographic characteristics. The terms *doctor* and *specialist* reflect roles in referrals, not physician specialty. Experience is years since medical school graduation. Pay is average annual Medicare payments by all payors in current dollars. Claims and Patients are average counts. Links is the number of distinct physicians with whom the physician had referral relationships. Incoming and outgoing referrals fractions are of fraction of referral dollar volume.

services are a quarter of the market for physician services in the U.S., which has an annual volume of half a trillion dollars, about 3% of the U.S. GDP.¹⁰ Even though claims for 20% of all patients are observed, selection of physicians into the sample based on their workload is negligible: even for those with a minimal workload, the probability of being sampled is close to 1. The average physician sees hundreds of Medicare patients every year. Seeing 30 patients is enough to be sampled with a probability of 0.999. For the same reason, the probability of missing links between physicians drops sharply as long as they see more than just a few patients.

Referrals in Medicare For the study of homophily, I construct the network of physician referrals from referral information recorded on claims. If one physician referred patients to another during the year, a link is recorded, with the link weight depending on the volume of the relationship, measured as one of the following: the number of patients, the number of claims, or total dollar value of services referred during the year. Table 1 (Panels B and C) shows that there are differences in the number of colleagues men and women physicians work with. Conditional on making any referrals, doctors refer to 16 specialists on average; conditional on receiving any referrals, specialists receive referrals from 18 doctors on average. But compared with women, men send referrals to 5 more specialists and receive referrals from 6 more doctors. These differences are explained in part by significant gender differences in specialization and the fact that the average number of referral relationships varies by medical specialty (see Appendix D for details). Therefore, it is important to control for specialty when studying homophily in referrals.

For the purpose of this study, it is useful that referrals in traditional Medicare are not limited or driven by institutional constraints, as beneficiaries can see any provider that accepts them. Unlike some managed care private insurance plans, there is no formal requirement to obtain a referral in order to see a specialist. Thus referrals are not mechanically constrained in that way. However, evidence suggests that referrals are still an important determinant of care tra-

new privately insured patients. Most physicians accept new patients of either insurance type.

¹⁰2012 National Health Expenditure Accounts (NHEA).

jectory (Johnson, 2011; Barnett et al., 2012; Choudhry, Liao and Detsky, 2014; Agha, Frandsen and Rebitzer, 2017).

Referrals are mostly made to nearby specialists. To study the implications of homophily on the pay gap, I therefore relate physician participation and pay at the local market level. I define local markets based on Hospital Referral Regions (HRR) from the 2012 Dartmouth Atlas of Healthcare.¹¹ There are in total 306 HRR, corresponding roughly to a metropolitan area. Each zip-code maps to one and only one HRR. HRR are the smallest geographical areas that represent nearly isolated networks: less than 15% of referrals cross their boundaries.

2.2 Decomposing The Gender Earnings Gap in Medicare

In 2012, the average female physician in the sample received a total of \$64,620 from Medicare, compared with the average male physician, who received \$121,995—48% less (66 log points). Figure 1 shows a gap in pay exists at every experience level and reaches its peak for mid-career physicians. Medicare has standardized payments per service and pays men and women equally. Therefore, this gap only reflects disparities in work quantity and type.

To quantify the contribution to the gap of previously studied mechanisms, I decompose the gross earnings gap by estimating a standard (log) annual pay equation:

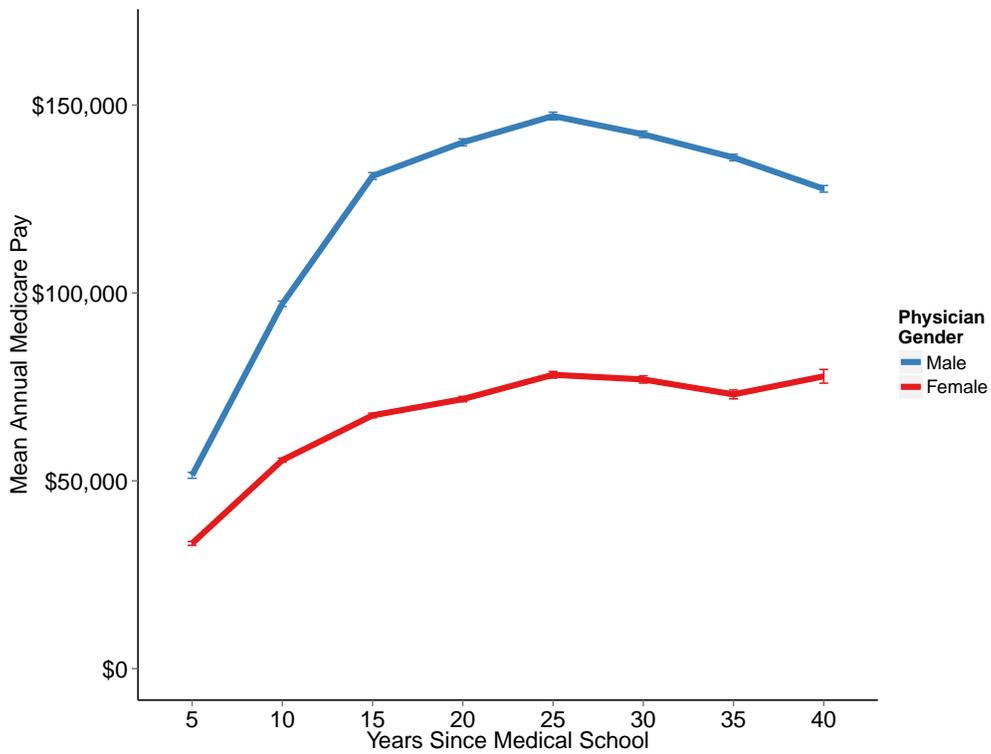
$$\log(\text{Pay}_k) = \beta \mathbb{1}_{g_k=M} + \delta X_k + \varepsilon_k \quad (1)$$

where k index physicians, $\mathbb{1}_{g_k=M}$ is a dummy for male physician, and X contains a constant and the following characteristics: physician specialty dummies; physician experience in years, including a quadratic term; previous no-work spells, defined as the fraction of quarters with zero claims in the observed history (i.e., during all sampled years, excluding the year of graduation and the current year), and dummies for local market (HRR) and for medical school attended. Results of this analysis are shown in Table 2.

About half of the Medicare physician pay gap is accounted for by known factors. The largest part (20 log points, or about a third of the overall gap) is due to women practicing lower-paying specialties. For example, 51% of active

¹¹“Dartmouth Atlas of Healthcare”, <http://www.dartmouthatlas.org/tools/downloads.aspx?tab=39>. Accessed December 2017

Figure 1: The Unadjusted Gender Pay Gap, by Experience Level



Notes: Source: 20% sample of Medicare physician claims for 2012. Mean Annual Medicare Pay is total annual payments (by all payers) to physicians for Medicare services, multiplied by 5 to adjust for sampling. Years are since medical school graduation (bin maximum, e.g. 10 stands for 6-10).

Table 2: The Gender Pay Gap for Medicare Physicians

	<i>Dependent variable:</i>					
	Log(Annual Pay)					
	(1)	(2)	(3)	(4)	(5)	(6)
Male Physician	0.668 (0.005)	0.654 (0.005)	0.468 (0.005)	0.361 (0.004)	0.337 (0.004)	0.340 (0.005)
Experience Quadratic	No	Yes	Yes	Yes	Yes	Yes
Specialty	No	No	Yes	Yes	Yes	Yes
No-Work Spells	No	No	No	Yes	Yes	Yes
HRR	No	No	No	No	Yes	Yes
Med. School	No	No	No	No	No	Yes
Constant	Yes	Yes	Yes	Yes	Yes	Yes
Observations	498,580	447,863	447,863	424,361	420,319	296,199
Adjusted R ²	0.033	0.052	0.290	0.407	0.441	0.471

Notes: Estimates from an OLS regression of annual pay on physician attributes. Experience is years since graduation. Specialty is a dummy for 54 CMS specialty code. No-work spells are previous quarters with no claims. HRR is a dummy for one of 306 Dartmouth Hospital Referral Regions. Med. School is Physician Compare medical school ID. The number of observations varies due to incomplete data on some characteristics.

obstetrician-gynecologists are women, but less than 6% of active orthopedic surgeons are women (Figure A3). Consistent with previous works, female physicians also have more career interruptions, which explain an additional 10 log points (Table 2, Columns 1–3). Differences in experience, location, and medical school attended explain a little more. But the remaining half of the gross gap (34 log points), reflecting the within-specialty gender gap in workload, remains largely unexplained. Understanding the causes for this large difference in workload is important, as beyond its direct effect on pay, lower workload by women could feed back to their specialization and career choices (Chen and Chevalier, 2012).

The earnings gap documented here for Medicare physicians conforms with previous studies of gender earnings gaps for physicians and other highly skilled professionals. Seabury, Chandra and Jena (2013) use Current Population Surveys (CPS) from 1987–2010 and estimate a median gap of 16%–25% (18–30 log points) among U.S. physicians. Weeks, Wallace and Wallace (2009) use survey data from 1998–2005 and find women earn about a third less than men. The gender pay gap in Medicare is also on par with pay gaps in other highly skilled occupations.

Bertrand, Goldin and Katz (2010) find a gross gap of almost 60 log points 10 to 16 years after graduation for MBA graduates working in the financial and corporate sectors, and Azmat and Ferrer (2016) find large gender gaps in hours billed and new client revenue among lawyers.

Like previously documented earnings gaps in other occupations, the earnings gap in medicine partly reflects known differences between the genders, including differences in labor supply. However, much of it remains unexplained. The rest of this paper exploits data on referral relationships among physicians to study in detail and quantify one additional contributing channel: how biased referrals make demand for female physicians lower than demand for male physicians.

3 Homophily in Physician Referrals

In this section, I define a measure of homophily that accounts for potential differences between the genders and show that physician referrals exhibit gender homophily beyond such differences. That is, doctors disproportionately refer more patients to specialists of their same gender.

3.1 Measuring Homophily with Unobserved Gender Differences

In order to examine evidence for gender homophily in physician referrals, I first define a new homophily measure, *directed homophily*. Directed homophily compares the fraction of referrals to male specialists between male and female doctors. Unlike previous homophily measures, directed homophily is insensitive to unobserved differences between the genders in the propensity to refer or receive referrals, and thus better reflects underlying bias in referrals towards same-gender others.

Consider the network of physician referrals in a given market, where a link exists between *doctor* j and *specialist* k if j referred any patients to k ; in such case we say j *refers* to k .¹² (Throughout, I use lowercase and uppercase to index

¹²Let G_g be the average fraction of referrals doctors of gender $g \in \{m, f\}$ send to specialists of gender $G \in \{M, F\}$.¹³ For example, the fraction of referrals male doctors send to female specialists is $F_m = \frac{n_{mF}}{n_{mF} + n_{mM}}$, where n_{gG} is the average number of referrals doctors of gender

doctors and specialists, respectively.) I define directed homophily as follows:

Definition 1 (Directed Homophily). *Directed homophily* is the difference between the fraction of outgoing referrals of male and female doctors to male specialists (or equivalently, to female specialists):

$$DH := M_m - M_f = F_f - F_m$$

That is, referrals exhibit directed homophily ($DH > 0$) if male doctors refer to male specialists more than their female counterparts. Table 3 illustrates this definition using Medicare data. In Medicare, male doctors refer 85% of their patients to male specialists, compared to female doctors, who refer 80% of their patients to male specialists, so $DH = 85 - 80 = 5p.p.$ (figures are rounded to the nearest integer). Instead of comparing outgoing referrals, one could define homophily based on the difference in incoming referral rates. It is easy to verify that both measures always have the same sign. Preserving link direction is important (e.g., see Figure A4). Directed homophily can also be redefined to admit weighted links (Appendix A). Weights reveal whether same-gender referrals are not only more likely but also more voluminous.

Table 3: Overall Directed Homophily (DH) in Medicare

		To (Specialist)	
		Female (F)	Male (M)
From (Doctor)	Female (f)	20%	80%
	Male (m)	15%	85%

$$DH = 85\% - 80\% = 20\% - 15\% = 5p.p.$$

g send to specialists of gender G .

Gender Imbalances in the Labor Force, Gender Differences in Labor Supply, and Other Gender Differences

Directed homophily is driven by neither baseline imbalance in the gender distribution of doctors or specialists nor other differences between male and female physicians. Since most specialists are men, both male and female doctors are expected to refer more to male specialists. Furthermore, because male physicians entered the profession eight years earlier on average, they are more experienced and thus may attract more referrals. In addition, differences in labor supply between male and female specialists could make women less available to receive referrals. Therefore, simply controlling for the population fraction would not be enough. However, directed homophily does not use population gender shares as a benchmark against which referral rates are compared. Instead, it measures whether a greater *fraction* of referrals made by men go to other men compared to the fraction of referrals made by women. Therefore, it does not reflect differences in referrals related to gender differences in experience or in labor supply. Instead, directed homophily captures a correlation between the gender of doctors and specialists, which would not be explained by such differences. More generally, directed homophily is insensitive to any differences between the genders in the propensity to send or receive referrals, observed or not.

Directed homophily contrasts with an oft-used homophily measure, *inbreeding homophily*, that uses population shares as a baseline.¹⁴ For example, if both genders referred more patients to male specialists only because males were a more appropriate target or preferred to work longer hours, unlike inbreeding homophily, directed homophily would still be zero. Directed homophily is positive only if there is a correlation between the gender of the referring doctor and the receiving specialist. Put differently, directed homophily measures relative, not absolute,

¹⁴Male doctors exhibit *inbreeding homophily* if $M_m > M$, where M is the fraction of male specialists. Likewise, female doctors exhibit inbreeding homophily if $F_f > F$, where $F = 1 - M$ is the fraction of female specialists (or equivalently, if $M_f < M$). Inbreeding homophily has long been used in sociology (see Coleman, 1958; McPherson, Smith-Lovin and Cook, 2001), and more recently in economics (e.g., Currarini, Jackson and Pin, 2009; Bramoullé et al., 2012b; Currarini and Vega-Redondo, 2013) (normalized or approximated variants are often used). Note that inbreeding homophily by both genders immediately implies directed homophily, while the reverse is not true, e.g. if $M_m > M_f > M$.

gender differences in referrals. Note that with unobserved heterogeneity, such absolute differences are generally not identified.

3.2 Patterns of Homophily in Physician Referrals

Overall, Medicare referrals exhibit significant directed homophily (Table A2). Before accounting for any differences in characteristics, 84.77% of the referrals made by male doctors are to male specialists, compared with only 80.27% of the referrals made by female doctors. Consequently, a greater fraction of (incoming) referrals to male specialists come from male doctors relative to the incoming referrals of female specialists.

Homophily is also significant within geographic segments. Figure 2 plots the average fraction of referrals to male specialists over the fraction of male specialists in 306 local U.S. markets (HRRs). Each market is represented by two vertically aligned points, which capture referral rates to male specialists by male and female doctors in the market. Unsurprisingly, the overall relationship between the fraction of male specialists and the fraction of referrals they receive is positive. However, two additional facts are apparent. First, the fraction of referrals going to male specialists is greater than the fraction of male specialists in the population (most points are above the 45-degree line). Second, even within local markets, referrals exhibit directed homophily (the vertical difference between the fitted curves).

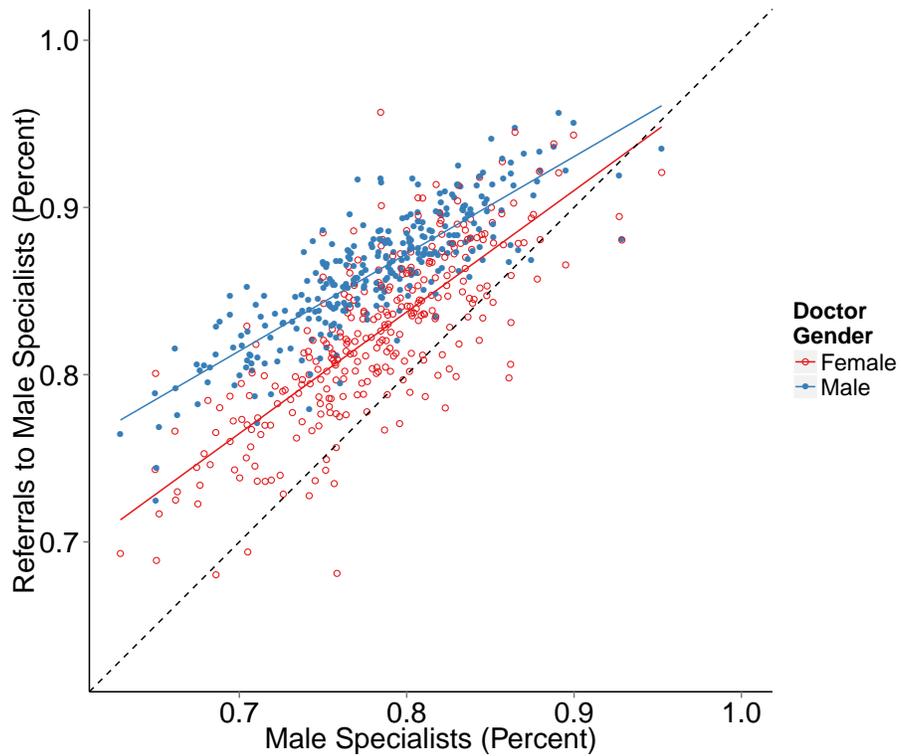
I further measure directed homophily for different bins by regressing the fraction of patients each doctor j referred to male specialists, M_j , on the doctor's gender g_j and other characteristics X_j :

$$M_j = \alpha_1 + \beta_1 g_j + \delta_1 X_j + \varepsilon_j, \quad (2)$$

for doctors with any referrals. The ordinary least squares estimate of β_1 measures average directed homophily: how much more often men refer to men on average across markets.

An alternative explanation for observed homophily, other than gender bias, is sorting of physicians into market segments by gender. Sorting would make the pool of specialists biased towards the doctors' own gender. Some control for

Figure 2: Referrals to Male Specialists Over Their Population Fraction, by Doctor Gender



Notes: For each local physician market (Dartmouth Hospital Referral Region), average fractions of referrals from male and from female doctors to male specialists are plotted over the fraction of male specialists in the market. Each of these 306 local U.S. markets is thus represented by two vertically aligned data points. On average, men refer more to men than women do, even after accounting for the variation between markets in the availability of male specialists. The proposed measure, *directed homophily*, represents the vertical difference between the fitted curves.

sorting can be obtained from a variant of (2):

$$M_j = \beta_2 g_j + \delta_2 X_j + \gamma_{h(j)} + \varepsilon_{jh} \quad (3)$$

where $\gamma_{h(j)}$ is a fixed-effect for the hospital with which doctor j is affiliated (for robustness, hospital interacted with medical specialty is also considered). The estimate β_2 measures the average directed homophily within market segments defined by hospital and specialty bins. It captures the differences in referral rates to men between male and female doctors of the same medical specialty that are affiliated with the same hospitals. To the extent such doctors face similar pools of specialists, β_2 indicates a bias in referrals.

An alternative explanation for homophily is that patients may prefer seeing physicians of their same gender.¹⁵ If such preferences affect both patients' choice of doctors and their choice of specialists, they may yield directed homophily among physicians, even absent any gender bias among physicians. One way to account for patient preferences is to control for each doctor's gender mix of patients in (2) or (3); another is to estimate similar specifications with disaggregated data, where each observation represents one referral of a patient by a doctor to a specialist and directly includes the gender of all parties involved.

Accounting for each doctor's gender and age mix of patients hardly changes estimated directed homophily, suggesting homophily in physician referrals is not driven by homophily on behalf of patients. Table 4 shows directed homophily measures obtained from individual physician data. There is a 4.3 percentage points difference between referral rate of male and female doctors of same specialty and experience (Column 2); This difference is 4.0 percentage points when the gender mix of patients of each doctor is controlled for (Column 3). Physicians with more male patients do refer more patients to male specialists on average. But doctors with similar fractions of male patients still refer more to specialists of their same gender. Results are very similar when more disaggregated data are used instead (Table A4) and are virtually unchanged when additional controls for patient age are included.

¹⁵For example, Reyes (2006) shows that female patients are more likely to visit female obstetrician-gynecologists.

Table 4: Estimates of Directed Homophily

	Percent of Referrals to Male Specialists					
	OLS (No FE)				OLS (With FE)	
	(1)	(2)	(3)	(4)	(5)	(6)
Male Doctor	0.053 (62.7)	0.043 (49.1)	0.040 (44.8)	0.040 (44.0)	0.029 (30.5)	0.030 (32.6)
Percent Male Patients			0.029 (16.5)	0.028 (14.7)	0.031 (16.1)	0.043 (23.4)
Cons.	0.79 (1027.6)	0.81 (263.8)	0.80 (254.3)	0.81 (256.9)	0.82 (249.4)	0.78 (589.1)
Specialty (Doctor)	No	Yes	Yes	Yes	Yes	No
Experience (Doctor)	No	Yes	Yes	Yes	Yes	Yes
Obs. (Doctors)	385,104	384,985	384,985	347,534	347,534	347,534
Groups (Hospital/Specialty)					4,819	66,563
Rank	2	56	57	57	57	4
Mean Dep. Var.	0.82	0.82	0.82	0.83	0.83	0.83
R^2	0.012	0.038	0.039	0.041		
R^2 Within					0.034	0.0079

Notes: t statistics in parentheses. Estimates of equations (2) and (3) for the sample of doctors with any referrals. The dependent variable, percent of referral to male specialists, is the fraction of referrals that are made to male specialists. Percent Male Patients is the fraction of referred patients who are male. Column 4 shows estimates of the same specification as Column 3 using the subsample of doctors with at least one hospital affiliation, used also in Columns 5 and 6. For sample and variable definitions, see Section 2.

Doctors refer more to specialists of their same gender even when comparison is restricted to market segments defined by doctors' hospital affiliation and their medical specialty (Columns 5 and 6). That is, male doctors affiliated with the same hospital, and of the same medical specialty, and with the same gender mix of patients still refer more to male specialists (a 3 percentage points difference). These estimates rule out the possibility that homophily only reflects homophily on behalf of patients or sorting of physicians by gender into hospitals.

Directed homophily is greater among older doctors (with above-median experience) than among younger ones (Table A5). Homophily estimates are virtually unchanged when links are weighted by different measures of referral volume: number of patients, number of claims, or overall dollar value of services (Table A6).

To conclude, directed homophily estimates reveal a correlation between the gender of doctors and specialists they refer patients to, even within different market segments defined by location, hospital, and specialty, and even when patient gender and age mix is accounted for. Such correlation is hard to explain in terms of unobserved differences between specialists of opposite genders, as any such difference should affect all referring doctors similarly, independently of their gender. Directed homophily thus suggests that referrals may be biased.

However, directed homophily varies with the fraction of male specialists available (note that the fitted curves in Figure 2 are not parallel). Therefore, it does not directly reflect the magnitude of any underlying bias in referrals. For example, in markets where nearly all specialists are male doctors, choices are constrained and reveal little about their preferences. In such cases, directed homophily would be close to zero even if referrals were biased. Therefore, identifying and quantifying the underlying bias in referral decisions requires accounting for differences between doctors in the fraction of available specialists, to which I turn next.

4 The Link Between Homophily in Referrals and Gender Biases in Physician Choices

This section examines the relationship between the observed homophily in referrals and the underlying gender bias in doctors' referral choices, and estimates the latter directly. Analyzing a discrete choice model where doctors choose spe-

cialists, I show that homophily decomposes to gender-biased preferences within market segments and sorting across such segments. I then use this model to separately estimate the gender bias in doctors’ choice of specialists as referral targets, which is identified by comparing the average characteristics between chosen and unchosen specialists, accounting for the variation in choice sets between doctors and for sorting on multiple potential dimensions. I find that faced with a choice between otherwise similar male and female specialists, doctors are 10% more likely to refer to a specialist of their same gender.

4.1 Homophily, Gender-Biased Referrals, and Sorting

Consider a model where doctors $j \in J$ choose specialists to refer patients to, from an opportunity pool $k \in K_j$. Denote the gender of doctors and specialists by $g_j \in \{f, m\}$, and $g_k \in \{F, M\}$. Doctors maximize a gender-sensitive utility function and choose a specialist:

$$\operatorname{argmax}_{k \in K_j} U_j(k) = \beta \mathbb{1}_{g_j=g_k} + \delta X_{jk} + \varepsilon_{jk} \quad (4)$$

where $\mathbb{1}_{g_j=g_k}$ indicates both physicians are of the same gender, i.e., $(g_j, g_k) \in \{(f, F), (m, M)\}$. The choice of specialists depends on individual and specialist attributes (X_{jk} ; e.g., specialist experience or distance between clinics), but may also depend on gender, if $\beta > 0$. This case represents *gender-biased preferences*. If ε_{jk} is an independently and identically distributed Gumbel extreme value, equation (4) yields the conditional logit probability for a referral from j to k , given gender and other characteristics:

$$p_{jk} := P(Y_{jk} = 1 | g_j, g_k, X) = \frac{e^{\eta_{jk}}}{\sum_{k' \in K_j} e^{\eta_{jk'}}} \quad (5)$$

where $Y_{jk} = 1$ if j refers to k and $Y_{jk} = 0$ otherwise, and $\eta_{jk} := \beta \mathbb{1}_{g_j=g_k} + \delta X_{jk}$. That is, referral relationships are determined by pairwise characteristics. This excludes more strategic setups where links are formed in response to other links or in anticipation of such links.

Homophily due to Gender-Biased Preferences Biased preferences lead to homophily. To see how, first consider the case where there is one market with one common pool of specialists $K_j = K$, for all doctors $j \in J$, and where $X_{jk} = X_k$, namely, it includes only specialist characteristics but not pairwise ones, so all doctors face essentially the same choice set. Let $M = \frac{1}{|K|} \sum_k \mathbb{1}_{g_k=M}$ be the fraction of male specialists in this set (with slight abuse of notation: M is also used throughout to label male specialists). The first proposition shows that gender-biased preferences lead to homophily.

Proposition 1 (Preference-Based Homophily). *Within a market, referrals exhibit directed homophily if and only if preferences are gender-biased. Namely, for $M \in (0, 1)$, $DH > 0$ if and only if $\beta > 0$.*

To see why Proposition 1 is true, first consider the homogeneous case: $\delta = 0$, and note that the conditional probabilities of referrals to a male specialist are:

$$P(M_m) = \frac{M}{M + \omega(1 - M)} \geq M \geq \frac{\omega M}{\omega M + (1 - M)} = P(M_f) \quad (6)$$

where $P(G_g) := P(g_k = G | g_j = g)$ denotes the probability that the chosen specialist's gender is G conditional on doctors' gender being g , and $\omega = e^{-\beta} \in (0, 1]$.¹⁶ Equation (6) shows that for all $M \in (0, 1)$, biased preferences result in directed homophily, $P(M_m) > P(M_f)$: doctors of each gender slightly discount the other (by a factor ω).¹⁷ Conversely, with unbiased preferences ($\beta = 0$), directed homophily is zero, as referral rates to men are common to doctors of both genders:

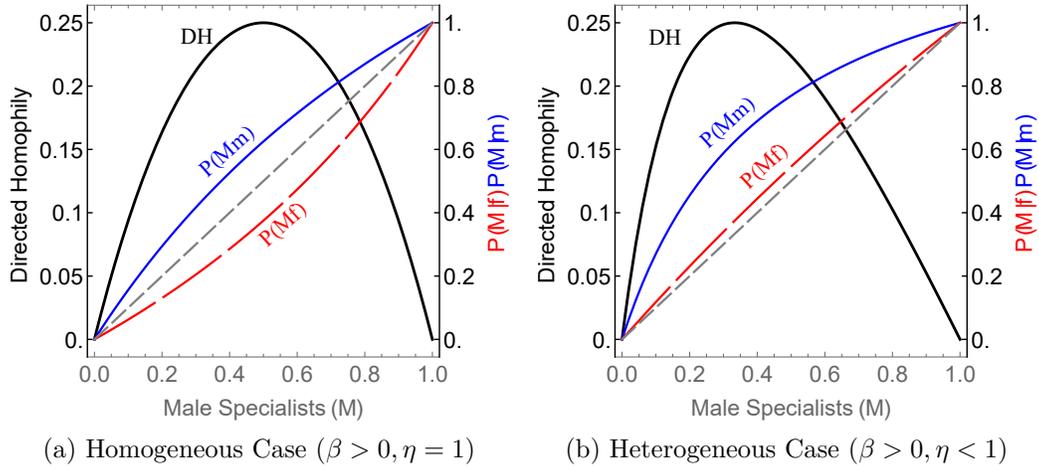
$$P(M_m) = M = P(M_f). \quad (7)$$

In case systematic differences exist between male and female specialists ($\delta \neq 0$), they affect both genders (Figure 3b). Even in this case, $P(M_m) = P(M_f)$ if and only if preferences are unbiased (see Appendix A for the proof). For simplicity, for the rest of this section, again assume homogeneity.

¹⁶See Appendix A for derivation of this and other results

¹⁷Clearly, if specialists are mostly men then men refer more to men than to women: $P(M_m) > P(F_m)$, which is not to be confused with $P(M_m) > P(M_f)$.

Figure 3: Preference-Based Homophily, With and Without Heterogeneity



Notes: Model-calculated probabilities of referrals to male specialists by male and female doctors and their difference—directed homophily—as functions of the fraction of specialists who are male (the dashed line is the 45-degree line). Case (a) shows gender-biased preferences and homogeneous specialist population: Male specialists receive more referrals than their fraction in the population from males and less than this fraction from females. Case (b) shows gender-biased preferences and heterogeneous specialist population: Doctors of both genders refer to male specialists more than their population fraction, but male doctors refer even more than female doctors to male specialists.

An important implication of (6) is that for a given bias in preferences, observed homophily is greater when the pool of available specialists is more gender-balanced (as illustrated in Figure 3, and as seen before in Figure 2; in the extreme cases when $M \in \{0, 1\}$ there is no homophily even when preferences are biased). When doctors face a balanced pool, their choices are less constrained and therefore more strongly reflect their preferences. Therefore, the preference bias β is a more portable parameter than directed homophily, which in part reflects differences among doctors in the mix of available specialists.

Homophily due to Sorting by Gender into Market Segments Apart from preferences, physician sorting by gender into local markets also leads to homophily.¹⁸ Sorting generates homophily by making doctors more exposed to specialists of their same gender. Formally, suppose that there is a set of separate markets indexed by $c \in C$, each with its own set of doctors J^c and specialists K^c , and the corresponding fractions of male doctors, m^c , and male specialists, M^c , assumed throughout to be in $(0, 1)$. Referrals only occur within markets. That is, $K_j = K$ for all $j \in J^c$. Markets may also vary in size $\mu^c = \frac{J^c}{J}$ (so $\sum_c \mu^c = 1$). The conditional probabilities of referrals to men now vary by market and are denoted $P(M_m|c)$ and $P(M_f|c)$. One way to define *sorting* is as a positive correlation between the genders of doctors and specialists across markets $\text{Cov}(m^c, M^c) > 0$.

Proposition 2 (Sorting-Based Homophily). *With sorting, referrals exhibit homophily when pooled across all markets:*

$$P(M_m) > M > P(M_f)$$

for all $\beta \geq 0$.

Intuitively, if fractions of male doctors and specialists are correlated then referrals coming from male doctors are more likely to occur in markets with more male specialists. Homophily then appears at the aggregate level, even when referrals are unbiased ($\beta = 0$) so there is no homophily within each market.

¹⁸Sorting could also be into market segments, like institutional affiliations, that determine the scope of referrals.

Sorting and biased preferences are in fact exhaustive: combined, they fully account for the overall homophily observed. The following proposition decomposes homophily into these two causes: preferences (within market) and sorting (across markets). For clarity, the proposition is stated here for inbreeding homophily. For its equivalent for directed homophily, see Appendix A.

Proposition 3 (Homophily Decomposition). *Homophily observed across all markets decomposes to preferences and sorting as follows:*

$$\overbrace{P(M_m) - M}^{\text{Overall Homophily}} = \frac{1}{m} \left(\overbrace{\mathbb{E}[m^c(P(M_m|c) - M^c)]}^{\text{Biased Preferences}} + \overbrace{\text{Cov}[m^c, M^c]}^{\text{Sorting}} \right).$$

That is, the homophily observed when all markets are pooled is the sum of two terms: (a) the average market-specific (preference-based) homophily, weighted by market size μ^c and share of doctors, $\frac{m^c}{m}$, and (b) sorting into markets.¹⁹ Note that gender differences (e.g., in labor supply) would not show up in this decomposition, because they would only generate different *levels* of referrals to men and women, not *correlations* between the gender of doctors and specialists, which is what directed homophily measures.

Proposition 3 shows that observed homophily in each market is a combination of preferences and sorting into unobserved market segments. A corollary of this proposition is that when market boundaries are observed, homophily observed within each market identifies the presence of a bias in preferences. When market boundaries are imperfectly observed (e.g., if physicians sort by gender into hospitals, but hospital affiliation is not observed), observed homophily is a combination of preferences and sorting. Accounting for sorting is therefore required to identify preference bias. I now turn to estimate this bias.

4.2 Identification and Estimation of Preference Bias

My primary concern is to identify the gender bias in doctors' preferences separately from sorting and from individual differences between the genders in the propensity to refer or receive referrals. In Section 3, homophily was measured by

¹⁹The proof of Proposition 3 only uses Bayes' Rule to relate aggregate and market-specific referral probabilities and does not rely on a specific parameterization of these probabilities.

comparing referrals between genders. This section takes a more direct approach and estimates the underlying bias in referrals using the specialist choice model specified above. The main parameter of interest is the gender-bias in preferences β , which quantifies how much doctors are more likely to refer to specialists of their own gender, all else being equal.

Identification is based on comparing gender and other characteristics between the set of specialists that were chosen by each doctor and the set of specialists that were available but not chosen. The model allows different doctors to face different pools of specialists. In fact, such variation in choice sets helps identify the model parameters. Potential sorting is mitigated by including controls for multiple factors that are expected to impact the likelihood of referrals between pairs of physicians, including: location (distance), specialty, experience, patient gender, shared medical school, and shared affiliations. The residual threat is from sorting on unobserved attributes, namely, from factors correlated with the gender of both doctors and specialists and that are relevant for the choice of referrals. Note that for an omitted factor to confound the estimates, it must not only be related to referrals, but also correlated with the genders of both doctors and specialists. Furthermore, I argue that characteristics unrelated to referral appropriateness that might be shaping preferences are not confounders, but rather underlying mechanisms (e.g., if men refer to men because they golf together, golf club affiliation explains homophily, but does not explain it away). The identification assumption is therefore that no clinically related factors correlated with both the probability of a referral and with the gender of both physicians are omitted. In Section 5, I further mitigate concerns for residual sorting by directly testing for a correlation between the gender mix of doctors and specialist demand, a correlation that is expected to exist if and only if preferences are gender-biased.

I estimate gender bias in preferences using the conditional logit model specified in (5) for the probability of referrals from doctor j to specialist k , conditional on gender g and other specialist and pairwise characteristics X . The identifying variation comes from differences within each doctor's choice set; thus, any doctor-level attributes are differenced out, as is clear from comparing the log of

the ratio of probabilities:

$$\log \frac{p_{jk}}{p_{jk'}} = \beta(\mathbf{1}_{g_j=g_k} - \mathbf{1}_{g_j=g_{k'}}) + \delta(X_{jk} - X_{jk'}). \quad (8)$$

The data consist of an observation for each dyad (j, k) , with associated physician and dyad (pairwise) characteristics X_{jk} , and a binary outcome standing for whether the dyad is linked. To account for differences between opportunity pools, X_{jk} includes specialist gender. The main parameter of interest is β , the average gender bias in preferences.

Because the opportunity pool of specialists is very large, considering all possible dyads is computationally unfeasible. I therefore use choice-based sampling (also known as case-control sampling). That is, instead of considering all possible pairs of doctors and specialists in the U.S., each chosen specialist k is compared against two randomly sampled unchosen specialists k' and k'' from the same HRR and of the same specialty of k , the specialist to which j actually referred. This choice makes a conservative (weak) assumption about substitutability. Specifically, specialists in the same city and medical specialty are not assumed to be perfect substitutes. Rather, it only assumes that specialists from different markets or from different medical specialties are not substitutes. Variation in multiple other characteristics that are included as controls (e.g., distance) capture the actual substitutability within those cells. This sampling scheme yields consistent estimates (McFadden, 1984).

4.3 Results: Homophily and Gender-Biased Preferences

On average, doctors more frequently choose specialists that are of their same gender and that are similar to themselves on other observed characteristics. Table 5 describes the sample used for estimating preference bias. It compares the average characteristics of specialists that are chosen with two randomly sampled specialists from the same market (HRR) and medical specialty who were not chosen. Chosen specialists are much more likely to have common institutional affiliations with the referring doctor, to be located nearby, to be of similar age, and to have gone to the same medical school.

Table 5: Average Characteristics of Chosen versus Unchosen Specialists

Doctor and Specialist:	Doctor Referred to Specialist	
	Yes	No [†]
Same Gender	0.712	0.678
Same Zip Code	0.280	0.0824
Same Hospital	0.778	0.298
Same Group	0.191	0.052
Same Med. School ⁺	0.107	0.0817
Experience Difference (years)	11.25	12.16
Observations (Dyads)	5,632,166	9,635,750
	2,852,950 ⁺	4,685,218 ⁺
Clusters (Doctors)		375,440
		242,579 ⁺

Notes: † Two specialists not chosen for referrals were randomly sampled from the set of specialists with the same HRR and medical specialty as those of each chosen specialist by each doctor. + is for the sample with non-missing school data. All differences are significant ($p < 0.001$).

All else being equal, doctors are 10% more likely to refer to specialists of the same gender, as seen in Table 6, which shows estimates of the specialist-choice model (5).²⁰ Distance (proximity) and hospital affiliation are the strongest determinants of referrals, with referrals far more likely between providers sharing an affiliation and within the same zip code. Modest sorting on location and hospital affiliation is confirmed by the slight decrease in same-gender estimates when these characteristics are included as controls. Doctors are also more likely to refer to specialists of similar experience. Estimates suggests that there is a comparable effect on the likelihood of referrals for being of different gender and for having an age difference of about ten years.²¹ Doctors are also more likely to refer to specialists that attended the same medical school.

Results from including interaction terms in the estimated model suggest that the preference bias is somewhat stronger within hospitals and somewhat weaker

²⁰Estimates represent odds ratios, but due to sparsity, estimates close to zero approximately equal the percentage increase in probability, i.e., around zero $\beta \approx \frac{p_{jk}|g_j = g_k}{p_{jk'}|g_j \neq g_{k'}} - 1$.

²¹Note however, that unlike gender, which shows imbalance with more male than female doctors, as long as no age group dominates the market by making the most referrals, homophily on age does not create a real disadvantage for any age group. It rather implies that work for specialists is more likely to be coming from doctors of similar age.

Table 6: Conditional-Logit Estimates: Referral Probability

Doctor and Specialist:	Doctor Referred to Specialist					
	(1)	(2)	(3)	(4)	(5)	(6)
Same Gender	0.104 (55.27)	0.0983 (41.96)	0.0841 (35.81)	0.105 (29.79)	0.104 (29.65)	0.0844 (27.71)
Male Specialist	0.197 (103.46)	0.169 (71.43)	0.175 (73.70)	0.165 (46.09)	0.165 (46.03)	0.194 (63.15)
Same Hospital		3.116 (721.31)	3.114 (720.15)	2.945 (541.87)	2.941 (540.97)	2.803 (568.47)
Same Practice Group		1.346 (178.27)	1.346 (178.27)	1.320 (135.27)	1.320 (135.26)	1.652 (142.36)
Similar Experience			0.128 (131.66)	0.132 (93.47)	0.131 (92.95)	0.136 (110.80)
Same Medical School					0.209 (49.96)	
Specialist Experience	Yes	Yes	Yes	Yes	Yes	Yes
Same Zip Code	No	Yes	Yes	Yes	Yes	No
Zip Code Distance	No	No	No	No	No	Yes
Obs. (Dyads)	14,793,483	14,559,311	14,555,821	6,712,241	6,712,241	8,915,969
Clusters (Doctors)	375,440	367,479	367,370	242,579	242,579	285,448
Pseudo R Sqr.	0.002	0.360	0.361	0.346	0.347	0.331

Notes: t statistics in parentheses. Results of conditional logit estimates of (5) for 2012. Data consist of all linked dyads and a matched sample of unlinked dyads, by location and specialty (see text for details). The dependent binary variable is 1 if there were any referrals between the doctor and the specialist during the year. Same Gender is a dummy for the specialist and doctor being of the same gender. Male Specialist is a dummy for the specialist being male. Same hospital and practice groups mean the doctor and the specialist have at least one common affiliation of either type. Similar Experience is the negative of the absolute difference in physicians' year of graduation. Zip Code Distance is a quadratic polynomial in the distance between zip code centroids for zip codes with a distance of 50 or fewer miles between them. Because medical school data are partial, Column (4) estimates the same specification as (3) with non-missing school data.

within groups, albeit only slightly. It does not depend on whether doctors and specialists graduated from the same medical school (Table A7). Estimated separately by specialty, gender bias is positive for all but a few small specialties where doctors are likely too restricted in their choices to be able to express any gender preference (Figure A7).

The estimated gender bias is comparable with the previously estimated homophily (Table 4). Substituting $\hat{\beta} = 0.1$ in equation (6) shows that facing an opportunity pool with 80% male specialists (roughly the U.S. average), the estimated gender bias of 10% implies directed homophily of 3.2 percentage points, net of sorting. Were the fractions of male and female specialists equal, this same bias would yield directed homophily of 5 percentage points.

Further considering the persistence of referral relationships, I find that same-gender physicians are also more likely to maintain referral relationships over time (Appendix B). Same-gender links are between 1.5–4.5% relatively more likely to persist (i.e., stay active the year after). This suggests same-gender referrals may be more common as a consequence of a dynamic process in which same-gender relationships are more likely to survive over time.

In sum, estimates point to a significant gender bias in referral choices even when multiple and detailed characteristics are included that account for possible sorting on multiple dimensions. Combined with earlier findings of robust directed homophily, evidence suggests that homophily is predominantly driven by preferences rather than sorting. Results imply that increasing the diversity of the opportunity pools would increase homophily rather than decrease it: gender-biased individual preferences are more manifest in diverse pools, which permit choice. Another implication—one that is central to this paper—is that homophily diverts demand away from females (the gender minority) and generates a gap in earnings. I next turn to formalize and test this implication directly.

5 The Contribution of Gender-Biased Preferences to the Gender Earnings Gap

In this section, I study the impact of homophily on the gender earnings gap. Using the model introduced in Section 4 as a quantification framework, I show that be-

cause most referring physicians are currently men, biased referrals make demand for female specialists 5% lower than demand for male specialists. This difference amounts to 15% of the current gender gap in physician specialist earnings. The resulting gap depends on the gender distribution of doctors and specialists—the stronger the majority of male doctors, the greater the impact on the gap—so one way to reduce the contribution of biased referrals to gender disparity in earnings is to balance the gender of referring doctors. I further test and confirm this predicted correlation between doctors’ gender and specialists’ demand using monthly data on Medicare payments for each specialist over 2008–2012.

Intuitively, when referrals exhibit homophily, specialists receive fewer referrals when fewer doctors share their gender. The following proposition shows that this relationship holds if and only if preferences are gender-biased. Furthermore, whether same-gender specialists substitute or complement each other depends on the gender mix of doctors.

Proposition 4 (Demand for Specialists by Gender). *All else being equal, average specialist demand depends on gender as follows:*

- i With gender-neutral preference ($\beta = 0$), specialist demand is invariant to gender.*
- ii With gender-biased preference ($\beta > 0$):*
 - (a) Each specialist’s demand is higher the more referring doctors are of his (or her) gender.*
 - (b) Same-gender specialists substitute for each other when most referring doctors are of their same gender, and complement each other when most doctors are of the opposite gender.*

The proof is by noting that the demand for male specialists—the average number of referrals received (denoted by D^M)—is a weighted average of doctors’ respective probability of referring to a male:

$$D^M = mP(M_m) + (1 - m)P(M_f) \tag{9}$$

(Demand for female specialists, D^F , is derived similarly.) Substituting (6) into (9) and differentiating by m and M yields:

$$\frac{\partial D^M}{\partial m} = \overbrace{P(M_m) - P(M_f)}^{\text{directed homophily}} \quad (10)$$

$$\frac{\partial D^M}{\partial M} = (1 - m) \overbrace{\frac{w(1 - w)}{(1 - M(1 - w))^2}}^{\text{Complements (+)}} + m \overbrace{\frac{-(1 - w)}{(M + w(1 - M))^2}}^{\text{Substitutes (-)}}. \quad (11)$$

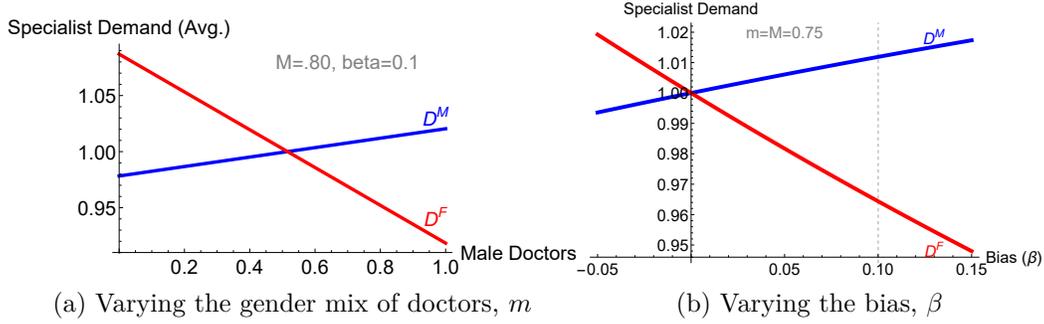
Equation (10) shows that demand for male specialists increases in the fraction of male doctors if and only if doctors are gender-biased. Equation (11) shows that male specialists are substitutes if most doctors are male and complements if most doctors are female. Figure A2 illustrates this general relationship for different values of M and m . As females are still the minority of both doctors and specialists in most markets (the darker area of the surface in the figure), demand for female specialists is lower due to both these effects: fewer doctors favor them, and male substitutes are easily found.

The same framework used to estimate the gender bias in referrals Section 4 (and further developed in (9)–(11)) can be used to quantify its impact on the difference in demand between male and female specialists. I consider two alternative counterfactual scenarios: (1) eliminating the gender bias in referrals, holding fixed the current fractions of male doctors and male specialists in the United States, and (2) balancing the gender mix of doctors, holding fixed the current gender bias in referrals.

Figure 4a shows that average demand for male and female specialists for different levels of bias, β , for average fractions of male doctors and specialists in the current U.S. physician markets. Overall, at the estimated bias in referrals, $\hat{\beta} = 0.1$, the fact three quarters of referrals are made by men result in demand for female specialists that is 5% lower than demand for male ones. This bias amounts to 15% of the unexplained, 33 log points gap in earnings. Eliminating the bias would result in a 1% decrease in male specialist earnings and a 4% increase in female specialist earnings (the asymmetry stems from females being the minority of specialists). These differences currently amount to thousands of dollars a year due to missing referrals to women. The contribution of biased referrals on the

earnings gap through limiting female demand is comparable in magnitude to the contribution to the gap of gender differences due to no-work spells (cf., Table 2).

Figure 4: Specialist Demand, Referrals Bias, and the Gender of Doctors



Notes: Predicted average demand for female and male specialists, given current fractions of male doctors and specialists (m , and M) and the bias in referrals β . Average demand is normalized to one. Figure (4a) shows the relationship between average demand for male and female specialists and the gender mix of doctors, given the estimated bias in referrals and the average gender mix of specialists. Conversely, Figure (4b) shows the relationship between demand and the bias, holding fixed the gender mix of doctors and specialists. Eliminating the bias or balancing the gender of doctors is predicted to reduce the gap in demand between male and female specialists by 5% (in popular terms, restoring 5 cents per dollar to women).

The impact of a given bias in preferences on the earnings gap is mediated by the gender composition of doctor and specialist populations (Proposition 4). Therefore, an alternative way to assess the impact of the estimated bias in referrals on gender earnings disparity is to predict how average demand would change if the fractions of male doctors and specialists changed, holding the bias constant. This relationship between demand for specialists of both genders and the fraction of male doctors, holding the fraction of specialists at its current average is illustrated in Figure 4b. The stronger the male majority among doctors, the greater the impact of any given bias on specialists' gender earnings gap.

An alternative way to reduce gender demand disparity among specialists (without eliminating the bias) is therefore to have more female doctors. Table 7 shows the counterfactual earnings gaps associated with different fractions of both male doctors (m) and male specialists (M). At a gender composition similar to the current one, $M = m = 0.8$, the estimated bias of $\hat{\beta} = 0.1$ is asso-

Table 7: Counterfactual Earnings Gap (Female-to-Male Difference) with Current Bias and Different Gender Mixes

Male	Males Specialists (M)						
Doctors (m)	0.4	0.5	0.6	0.7	0.8	0.9	1
0.4	0.0190	0.0200	0.0210	0.0220	0.0230	0.0240	0.0250
0.5	-0.0010	0	0.0010	0.0020	0.0030	0.0040	0.0050
0.6	-0.0210	-0.0200	-0.0190	-0.0180	-0.0170	-0.0160	-0.0150
0.7	-0.0410	-0.0400	-0.0390	-0.0380	-0.0370	-0.0360	-0.0351
0.8	-0.0610	-0.0600	-0.0590	-0.0580	-0.0570	-0.0560	-0.0551
0.9	-0.0809	-0.0799	-0.0789	-0.0780	-0.0770	-0.0761	-0.0751
1	-0.1009	-0.0999	-0.0989	-0.0980	-0.0970	-0.0961	-0.0952

Notes: Using estimated bias $\hat{\beta} = 0.1$, the table shows calculated earnings gaps: $D^F - D^M$, a function of m , M , and β , due to homophily-related workload differences, for different gender distributions of doctors and specialists. At $M = m = 0.75$ the gender bias in referrals is contributing 4.75 percentage points (or "cents per dollar") to the physician gender earnings gap.

ciated with a 5.7% lower demand for female, relative to male, specialists. With equal gender fractions ($M = m = 0.5$) there is no gap, even when preferences are biased. In this case, homophily only affects the composition of demand, not its level. While both the fraction of male doctors and the fraction of male specialists matter, the former is far more important at the estimated level of bias. (With extreme (lexicographic) bias, both matter equally; see Appendix Section C.) This relationship, between demand for specialists of either gender and the gender mix of doctors, is directly tested and confirmed in the following section.

5.1 Testing the Impact of Homophily on Earnings Using Longitudinal Data

Proposition 4 provides a testable prediction: If referrals are gender-biased, the fraction of male doctors should be positively correlated with the earnings of male specialists and negatively correlated with the earnings of female specialists. Such correlation is expected only if preferences are gender-biased and therefore doubles as a test for both a presence of gender bias in referrals and for the impact of such bias on the gender gap in specialist earnings.

I test this prediction using a monthly panel of physician payments. I split all physicians into two mutually exclusive groups, by their medical practice: primary-care physicians, who handle most outgoing referrals, and all other medical spe-

cialists. I estimate:

$$\log(\text{Pay}_{k,t}) = (\beta_M \mathbf{1}_{g_k=M} + \beta_F \mathbf{1}_{g_k=F}) m_{c(k,t),t} + \gamma_t + \alpha_k + \varepsilon_{k,t} \quad (12)$$

for all non-primary-care specialists, denoted by k , and months, denoted by t . The dependent variable $\text{Pay}_{k,t}$ is the specialist total monthly Medicare payments; the variable $m_{c(k,t),t}$ is the fraction of claims at specialist k 's market at month t handled by male primary-care doctors; specialist and time fixed effects, α_k and γ_t , are included. Of interest is the difference $\beta_M - \beta_F$: the differential impact a higher fraction of male doctors has on male and female specialists' pay, tested against the null of unbiased referrals, where this difference is zero.

Including specialist fixed effects allows for unobserved differences that likely exist between male and female specialists. This includes possible differences in labor supply (e.g., due to maternity-related leaves). This specification also allows for workload to be correlated across specialties. Indeed, it is likely that when primary care doctors see more patients, so do specialists due to, for example, seasonality. The identifying assumption is that no omitted factors simultaneously boost the monthly workload of male primary-care physicians and decrease the workload of female non-primary-care specialists. Controls are also included for the monthly fraction of services incurred by male patients to rule out patient homophily as an explanation.²²

I estimate (12) using a monthly panel of individual-physician pay for the period of 2008–2012. That is, I calculate for each market and month the fraction of primary-care claims handled by male doctors and, separately, the fraction of services incurred by male patients. Claim payments are aggregated to obtain total monthly payments for each physician specialist and each month.

Results show that the more referrals are handled by male primary-care physicians, the greater the demand for male specialists, and the smaller the demand for female specialists. Specifically, each 1.0% monthly increase in the fraction of referrals handled by male doctors is associated with a 0.47% increase in male workload and a 0.27% decrease in female workload. Results barely change when

²²To control for patient homophily, a term $(\delta_M \mathbf{1}_{g_k=M} + \delta_F \mathbf{1}_{g_k=F}) \mu_{c(k,t),t}$ is included, where μ is the percent of services incurred by male patients at k 's market at t . Here too the effect is allowed to differ by specialist gender.

controls for patient gender are included, suggesting the effect is not due to endogenous sorting by patients. These results are all identified from within-specialist variation in workload, so they are not an artifact of systematic labor supply differences between male and female specialists. Results support the presence of a direct link between gender bias in referrals and a disparity in demand between the genders that contributes to the gender earnings gap. Such correlation between the monthly workloads of specialists of opposite genders and the gender of nearby doctors is hard to explain in terms of gender differences in labor supply.

The magnitude of the estimated effect of biased referrals on gender pay disparities is fairly large. It suggests that were females to handle half of outgoing primary-care referrals instead of their current share, the pay gap among non-primary-care specialists would decrease by an estimated $(.50 - .35) \times (0.47 + 0.27) = 11\%$. However, unlike the previous counterfactual calculations, such back-of-the-envelope calculation does not restrict the overall volume and therefore should be taken as suggestive only.

Table 8: Male Fraction of Primary Care and Specialist Workload

	(1)	(2)
	log(Monthly Pay)	log(Monthly Pay)
Female specialist x pct. male PCP (HRR)	-0.26 (0.054)	-0.27 (0.054)
Male specialist x pct. male PCP (HRR)	0.49 (0.029)	0.47 (0.029)
Month Dummies	Yes	Yes
M,F x Pct Male patients (HRR)	No	Yes
Obs. (Physician x Month)	18,087,629	18,087,629
Clusters (Physician)	418,939	418,939
R Sqr.	0.0323	0.0322

Notes: Fixed-effect estimates of (12) with and without controls for patient gender. For each non-primary-care physician specialist, monthly pay is the the total monthly pay for Medicare services billed. Specialist gender is interacted with the fraction of claims handled by male primary-care physicians in the same market during the month. In Column 2 it is also interacted, separately, with the percent of services incurred by male patients in the market (as controls). Standard errors are clustered by specialist.

As Proposition 4 shows, were homophily solely due to sorting, a correlation like the one documented, between specialist workload, their gender, and the gender of nearby doctors, is not to be expected. Thus, results also suggest that the

observed homophily estimated in Section 3 is at least in part due to biased preferences and further support the identification of the preference bias in Section 4.

6 Mechanisms, Alternative Explanations, and Implications

Whether homophily implies that referrals are biased and whether this bias contributes to the earnings gap depends on ruling out potential confounding factors. I reevaluate three such factors: gender differences in labor supply, physician sorting, and patient sorting. I suggest that these factors fail to explain the totality of evidence, namely: directed homophily in referrals that reflects a correlation between the gender of doctors and specialists, even within narrowly defined cells; doctors' estimated tendency to refer to specialists of their own gender, even holding other factors equal; and the correlation over time between specialist workload and the gender of nearby doctors, whose sign depends on the congruence of specialists' own gender with the gender of nearby doctors. I therefore argue that evidence suggests that referrals are indeed biased, and that this bias explains a substantial fraction of the current gap in workload between male and female specialists. I then go on to discuss potential implications of the bias in referrals.

Earlier work and the decomposition from Section 2.2 suggest that significant differences might still exist between male and female physician specialists in their preferences for flexibility or in their overall desired workload. However, such differences fail to explain the other evidence, which suggests female specialists also face lower demand due to fewer referrals. Specifically, that female specialists may choose to work less does not explain why they receive a different fraction of their referrals from male doctors than their male counterparts (Section 3). Furthermore, labor supply differences fail to explain why doctors of similar location, affiliation, tenure, and medical school background refer more to specialists of their same gender (Section 4). Finally, had female specialists limited their capacity voluntarily, one would not expect their workload to increase with the fraction of referrals handed out by nearby female doctors (Section 5.1).

As discussed in Section 4.1, sorting is an alternative explanation for homophily in referrals. However, several results suggest that sorting fails to explain the ob-

served patterns of homophily and that a gender bias in referrals is contributing to them substantially. First, homophily exists even when comparing doctors and specialists with joint affiliations, same specialty, and the same experience. Second, homophily estimates hardly change when patient gender and age are accounted for. Third, the discrete model estimates suggest that gender determines referral target choice even when multiple potential dimensions of sorting, including location, hospital and group practice affiliations, medical school attended, and experience are accounted for (Table 6) when interactions are considered (Table A7), and when considered separately across specialties (Figure A7). The idea of sorting being the sole explanation for homophily is hard to reconcile with evidence of the dependency of workload on nearby doctors' gender and the persistence of same-gender work relationships over time (A1). The totality of evidence therefore suggests that a significant portion of the current gap is due to differences in demand, not just supply.

This evidence for the impact of biased referrals on the earnings gap does not account for indirect benefits specialists might accrue from having more referred patients, such as having more returning patients (in specialties where patients are repeatedly seen) or having additional patients through word of mouth. This analysis also does not cover the potential long-term effects biased referrals may have on female entry into medicine and into specific medical specialties. Therefore, the overall impact of gender-biased referrals could be greater.

One potential concern is that female physicians make up for missing referrals of Medicare patients through a higher workload outside of Medicare. However, there is evidence to the contrary. First, as discussed in Section 2.2, substantial gender differences in physician earnings have been documented in other settings, including when non-Medicare payments are considered.²³ Second, the high frequency correlation between female specialist workload and nearby gender of doctors (Table 8) would require a rather specific compensatory patient load outside of Medicare. Finally, there are no obvious reasons why demand disparity due to referral patterns would not be similar in other parts of the healthcare system.

²³For example, in a survey of 24,216 U.S. physicians across 25 specialty areas conducted in 2012, Medscape found that male physicians earned on average about 40% more than female physicians. (Medscape Physician Compensation Report 2012 <http://www.medscape.com/features/slideshow/compensation/2012/public>.) Accessed May 2017.

The analysis has focused on the intensive margin of workload rather than on the extensive margin of entry. The question remains as to whether women, who were significantly underrepresented in medicine, refrained from entering the profession in earlier periods, or refrain from entering specific specialties to this day because of homophily. The absence of women from medicine has dramatically changed in recent decades: slightly more than half of current medical school graduates are female, and medicine is increasingly a feminine profession. But women are still grossly underrepresented in many lucrative specialties, many of which, such as surgical specialties, rely greatly on referrals. The short time span of the data, and the presence of other differences between specialties, such as training duration and schedule flexibility, make it hard to identify the effects homophily may have on the extensive margins of female participation and specialization. But results hint of the possibility that homophily is an impediment to female entry, both into medicine in general and into particular specialties. If true, it would imply an even greater contribution of biased referrals to the gender gap.

7 Conclusion

This paper examines the contribution of gender biases in professional networks to the gender earnings gap. The focus is on the medical profession, where data on patient referrals from Medicare reveal the gender of both the referring and the receiving physicians as well as associated payments. Such data allow me to study in detail and assess the contribution to the earnings gap of one particular channel: homophily—the tendency of people to connect to others similar to themselves.

I define a new homophily measure that compares the outgoing referrals rates to male specialists between male and female doctors (rather than to population fractions). Such comparison “differences out” any systematic differences between the genders that could result in more referrals to men, capturing only a disproportional tendency to refer within gender. I further estimate a discrete-choice model in which doctors choose specialists from local pools to quantify the contribution of doctors’ preferences to the observed homophily, net of potential sorting. I use the same model to quantify the contribution of biased referrals to the earnings gap.

Using data on referrals among half a million U.S. physicians in 2008–2012, I find robust evidence for the presence of gender homophily in physician referrals that contributes substantially to the gender earnings gap in Medicare payments. Homophily is driven predominantly by gender-biased doctors’ choices, not sorting. Because most referring doctors are currently men, biased referrals generate gender differences in demand. This channel is separate from known gender differences in labor supply that explain other parts of the gender earnings disparity.

The empirical evidence suggests that a positive externality is associated with increased female participation in medicine. While medicine is a small fraction of the labor market and the particulars of referral practices in this market may not entirely generalize to other fields, the evidence highlights the possibility that homophily—a widely documented tendency in numerous types of interactions—contributes to gender disparities in other professions, particularly where networking is important. Therefore, homophily can provide a key to understanding the persistence of gender inequality.

References

- Agha, Leila, Brigham Frandsen, and James B Rebitzer.** 2017. “Causes and Consequences of Fragmented Care Delivery: Theory, Evidence, and Public Policy.” *NBER Working Paper*.
- Antonovics, Kate, and Brian G Knight.** 2009. “A new look at racial profiling: Evidence from the Boston Police Department.” *The Review of Economics and Statistics*, 91(1): 163–177.
- Anwar, Shamena, and Hanming Fang.** 2006. “An alternative test of racial prejudice in motor vehicle searches: Theory and evidence.” *The American Economic Review*, 96(1): 127–151.
- Azmat, Ghazala, and Barbara Petrongolo.** 2014. “Gender and the labor market: What have we learned from field and lab experiments?” *Labour Economics*, 30: 32–40.
- Azmat, Ghazala, and Rosa Ferrer.** 2016. “Gender gaps in performance: Evidence from young lawyers.” *Journal of Political Economy*.
- Baker, Laurence C.** 1996. “Differences in earnings between male and female physicians.” *New England Journal of Medicine*, 334(15): 960–964.
- Barnett, Michael L, Nancy L Keating, Nicholas A Christakis, A James OMalley, and Bruce E Landon.** 2012. “Reasons for choice of referral physician

- among primary care and specialist physicians.” *Journal of general internal medicine*, 27(5): 506–512.
- Bertrand, Marianne.** 2011. “New perspectives on gender.” *Handbook of labor economics*, 4: 1543–1590.
- Bertrand, Marianne, and Sendhil Mullainathan.** 2004. “Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination.” *The American Economic Review*, 94(4): 991–1013.
- Bertrand, Marianne, Claudia Goldin, and Lawrence F Katz.** 2010. “Dynamics of the gender gap for young professionals in the financial and corporate sectors.” *American Economic Journal: Applied Economics*, 228–255.
- Blau, Francine D, and Lawrence M Kahn.** 2017. “The gender wage gap: Extent, trends, and explanations.” *Journal of Economic Literature*, 55(3): 789–865.
- Bramoullé, Yann, Sergio Currarini, Matthew O Jackson, Paolo Pin, and Brian W Rogers.** 2012a. “Homophily and long-run integration in social networks.” *Journal of Economic Theory*, 147(5): 1754–1786.
- Bramoullé, Yann, Sergio Currarini, Matthew O Jackson, Paolo Pin, and Brian W Rogers.** 2012b. “Homophily and long-run integration in social networks.” *Journal of Economic Theory*, 147(5): 1754–1786.
- Burks, Stephen V, Bo Cowgill, Mitchell Hoffman, and Michael Gene Housman.** 2014. “The value of hiring through employee referrals.” *Quarterly Journal of Economics*.
- Chen, M Keith, and Judith A Chevalier.** 2012. “Are women overinvesting in education? Evidence from the medical profession.” *Journal of Human Capital*, 6(2): 124–149.
- Choudhry, Niteesh K, Joshua M Liao, and Allan S Detsky.** 2014. “Selecting a Specialist: Adding Evidence to the Clinical Practice of Making Referrals.” *JAMA*, 312(18): 1861–1862.
- Coleman, James.** 1958. “Relational analysis: the study of social organizations with survey methods.” *Human organization*, 17(4): 28–36.
- Currarini, Sergio, and Fernando Vega-Redondo.** 2013. “A simple model of homophily in social networks.” *University Ca’Foscari of Venice, Dept. of Economics Research Paper Series*, , (24).
- Currarini, Sergio, Matthew O Jackson, and Paolo Pin.** 2009. “An economic model of friendship: Homophily, minorities, and segregation.” *Econometrica*, 77(4): 1003–1045.
- Esteves-Sorenson, Constança, Jason Snyder, et al.** 2012. “The gender earnings gap for physicians and its increase over time.” *Economics Letters*, 116(1): 37–41.

- Ghani, Ejaz, William R Kerr, and Christopher Stanton.** 2014. "Diasporas and outsourcing: evidence from oDesk and India." *Management Science*, 60(7): 1677–1697.
- Goldin, Claudia.** 2004. "The long road to the fast track: Career and family." *The Annals of the American Academy of Political and Social Science*, 596(1): 20–35.
- Goldin, Claudia, Lawrence F Katz, and Ilyana Kuziemko.** 2006. "The Homecoming of American College Women: The Reversal of the College Gender Gap." *The Journal of Economic Perspectives*, 20(4): 133–156.
- Golub, Benjamin, and Matthew O Jackson.** 2012. "How homophily affects the speed of learning and best response dynamics."
- Graham, Bryan S.** 2014. "An empirical model of network formation: detecting homophily when agents are heterogenous."
- Guvenen, Fatih, Greg Kaplan, and Jae Song.** 2014. "The Glass Ceiling and The Paper Floor: Gender Differences among Top Earners, 1981–2012." NBER Working Paper.
- Halberstam, Yosh, and Brian Knight.** 2016. "Homophily, group size, and the diffusion of political information in social networks: Evidence from twitter." *Journal of Public Economics*, 143: 73–88.
- Hellerstein, Judith K, Melissa McInerney, and David Neumark.** 2011. "Neighbors and Coworkers: The Importance of Residential Labor Market Networks." *Journal of Labor Economics*, 29(4): 659–695.
- Himmelboim, Itai, Stephen McCreery, and Marc Smith.** 2013. "Birds of a feather tweet together: Integrating network and content analyses to examine cross-ideology exposure on Twitter." *Journal of Computer-Mediated Communication*, 18(2): 40–60.
- Hochberg, Yael V, Alexander Ljungqvist, and Yang Lu.** 2007. "Whom you know matters: Venture capital networks and investment performance." *The Journal of Finance*, 62(1): 251–301.
- Johnson, Erin M.** 2011. "Ability, Learning and the Career Path of Cardiac Specialists." *Cambridge, Massachusetts Institute of Technology*.
- Lo Sasso, Anthony T, Michael R Richards, Chiu-Fang Chou, and Susan E Gerber.** 2011. "The \$16,819 pay gap for newly trained physicians: the unexplained trend of men earning more than women." *Health Affairs*, 30(2): 193–201.
- Manski, Charles F, and Steven R Lerman.** 1977. "The estimation of choice probabilities from choice based samples." *Econometrica*, 1977–1988.
- McFadden, Daniel.** 1984. "Econometric analysis of qualitative response models." *Handbook of econometrics*, 2: 1395–1457.
- McPherson, Miller, Lynn Smith-Lovin, and James M Cook.** 2001. "Birds of a feather: Homophily in social networks." *Annual Review of Sociology*, 415–444.

- Niederle, Muriel, and Lise Vesterlund.** 2007. “Do women shy away from competition? Do men compete too much?” *The Quarterly Journal of Economics*, 122(3): 1067–1101.
- Renneboog, Luc, and Yang Zhao.** 2011. “Us knows us in the UK: On director networks and CEO compensation.” *Journal of Corporate Finance*, 17(4): 1132–1157.
- Reyes, Jessica Wolpaw.** 2006. “Do female physicians capture their scarcity value? The case of OB/GYNs.” NBER Working Paper.
- Sarsons, Heather.** 2017. “Interpreting Signals in the Labor Market: Evidence from Medical Referrals.” Job Market Paper.
- Seabury, Seth A, Amitabh Chandra, and Anupam B Jena.** 2013. “Trends in the earnings of male and female health care professionals in the United States, 1987 to 2010.” *JAMA internal medicine*, 173(18): 1748–1750.
- Weeks, William B, Tanner A Wallace, and Amy E Wallace.** 2009. “How do race and sex affect the earnings of primary care physicians?” *Health Affairs*, 28(2): 557–566.
- Zimmerman, Seth D.** 2014. “The returns to college admission for academically marginal students.” *Journal of Labor Economics*, 32(4): 711–754.

Online Appendix

A Mathematical Appendix

Defining Directed Homophily with Weighted Links

As discussed in Section 3.1, directed homophily can be easily adapted to accommodate weighted links. This section details this definition. First, define n_{gG} using weighted degrees, as follows: Let n_{jk} be the weight of the link from j to k (e.g. number of patients referred). The weighted out-degree of j is $d(j) = \sum_k n_{jk}$. The weighted out-degree to females is $d^F(j) = \sum_k \mathbb{1}_{g_k=F} n_{jk}$. Now n_{mF} is the average of $\frac{d^F}{d}$ over all male j , and so on for n_{gG} . The rest of the definition is as previously indicated in Section 3.1.

Proofs

Proof. (Proposition 1) Pick any j such that $g_j = m$. Summing up probabilities of referrals to all available specialists gives:

$$\begin{aligned} P(M_m) &= \sum_{k:g_k=M} P(Y_{jk} = 1) = \frac{\sum_{k:g_k=M} e^{\beta \mathbb{1}_{g_j=g_k}}}{\sum_k e^{\beta \mathbb{1}_{g_j=g_k}}} \\ &= \frac{\sum_{k:g_k=M} e^{\beta}}{\sum_{k:g_k=M} e^{\beta} + \sum_{k:g_k \neq M} e^0} = \frac{M e^{\beta}}{M e^{\beta} + 1 - M}. \end{aligned}$$

The probability $P(M_f)$ is similarly derived.

Consider next the case: $\delta \neq 0$, where a correlation exists between gender and decision-relevant specialist characteristics (e.g., men may be more experienced, or women may be available for fewer hours). In this case, (6) becomes:

$$P(M_m) = \frac{M}{M + \omega\eta(1 - M)} \geq \frac{\omega M}{\omega M + \eta(1 - M)} = P(M_f). \quad (13)$$

which holds, because:

$$P(M_m) = \frac{\sum_{k:g_k=M} e^{\beta \mathbb{1}_{g_j=g_k} + \delta X_k}}{\sum_k e^{\beta \mathbb{1}_{g_j=g_k} + \delta X_k}} = \frac{\sum_{k:g_k=M} e^{\beta + \delta X_k}}{\sum_{k:g_k=M} e^{\beta + \delta X_k} + \sum_{k:g_k \neq M} e^{\delta X_k}}$$

$$\xrightarrow{P} \frac{M\eta_M e^\beta}{M\eta_M e^\beta + (1-M)\eta_F} = \frac{M e^\beta}{M e^\beta + \eta(1-M)}$$

where $\eta_G = E[e^{\delta X_k} | g_k = G]$ for $G \in \{M, F\}$, and $\eta = \frac{\eta_F}{\eta_M}$ (so $\eta \gtrless 1$ when $E[e^{\delta X_k} | g_k = F] \gtrless E[e^{\delta X_k} | g_k = M]$). The convergence is by the Law of Large Numbers, assuming characteristics are independent across specialists.

Regardless of gender-biased preferences, if $\eta < 1$ male specialists attract a disproportionately high fraction of referrals from both genders (Figure 3b). Conversely, if $\eta > 1$, female specialists attract more referrals, so whether $P(M_m)$ and $P(M_f)$ are each greater or smaller than M depends on η . In (13) too, $P(M_m) = P(M_f)$ if and only if preferences are unbiased, i.e., $\beta = 0$. So Proposition 1 also holds for the heterogeneous case. \square

Proof. (Proposition 2) The overall conditional probability is a weighted average of market-specific conditional probabilities (weights are proportional to both market size and the relative share of male doctors in each market). Using Bayes' rule:

$$\begin{aligned} P(M_m) &= \sum_{c \in C} P(c|m)P(M|m, c) = \sum_{c \in C} \mu^c \frac{m^c}{m} P(M|m, c) \\ &\geq \sum_{c \in C} \mu^c \frac{m^c}{m} M^c = \frac{1}{m} E[m^c M^c] \\ &> \frac{1}{m} E[m^c] E[M^c] = M. \end{aligned}$$

The first inequality is due to preferences: $P(M|m, c) \geq M^c$ (equality being the case $\omega = 1$), and the second is due to segregation. By symmetry, the same proof works for females. \square

Note that the definition of sorting extends to the more general case where K_j is specific to each doctor as: $\text{Cov}(m^j, M^{K_j}) > 0$, where $m^j = \mathbf{1}_{g_j=m}$ and M^{K_j} is the fraction of male in K_j . (this definition is indeed more general, as by covariance decomposition, $\text{Cov}[m_j, M^j] = \text{Cov}[m^c, M^c]$ under separate markets with common $K_j = K^c$ in each.) For this more general definition of sorting the proof follows immediately from Proposition 1: with unbiased preferences $P(M_m) = E[M^j | g_j = m] > M$, by $\text{Cov}[m_j, M^j] > 0$.

Proof. (Proposition 3)

$$\begin{aligned}
P(M_m) - M &= \sum_{c \in C} \mu^c \left(\frac{m^c}{m} P(M|m, c) - \frac{m^c}{m} M^c + \frac{m^c}{m} M^c - M^c \right) \\
&= \sum_{c \in C} \mu^c \left(\frac{m^c}{m} (P(M|m, c) - M^c) + M^c \left(\frac{m^c}{m} - 1 \right) \right) \\
&= E \left[\frac{m^c}{m} (P(M|m, c) - M^c) \right] + \text{Cov} \left[\frac{m^c}{m}, M^c \right].
\end{aligned}$$

(Note that this proof only uses Bayes' rule to relate aggregate and market-specific referral probabilities and does not rely on a specific parameterization of these probabilities: it only requires relevant moments to exist.) \square

See below for a statement and proof of this proposition for directed homophily.

Proof. (Proposition 4) Pick any male specialist k . The demand k faces in market c is obtained by aggregating over all doctors in that market (as all variables are market specific, I suppress the superscript c):

$$\begin{aligned}
D_M &= \sum_{j \in J} p_{jk} = \sum_{j \in J} \frac{e^{\beta s(j,k)}}{\sum_{k' \in K} e^{\beta s(j,k')}} \\
&= \sum_{j \in J, g_j=1} \frac{e^{\beta s(j,k)}}{\sum_{k' \in K} e^{\beta s(j,k')}} + \sum_{j \in J, g_j=0} \frac{e^{\beta s(j,k)}}{\sum_{k' \in K} e^{\beta s(j,k')}} \\
&= \frac{1}{N} \left(\sum_{j \in J, g_j=1} \frac{1}{M + \omega(1 - M)} + \sum_{j \in J, g_j=0} \frac{\omega}{\omega M + (1 - M)} \right) \\
&= \frac{n}{N} \left(\frac{m}{M + \omega(1 - M)} + \frac{\omega(1 - m)}{\omega M + (1 - M)} \right).
\end{aligned}$$

Where $n = |J|$ and $N = |K|$. When $\omega = 1$ then $D_M = \frac{n}{N}$, which is independent of both M and m . Suppose $\omega < 1$. To see that ii(a) is true, rewrite:

$$\begin{aligned}
D_M &= \frac{n}{NM} \left(mP(M_m) + (1 - m)P(M_f) \right) \\
&= \frac{n}{NM} \left(P(M_f) + m(P(M_m) - P(M_f)) \right)
\end{aligned}$$

and note that $\partial D_M / \partial m > 0$ since $P(M_m) - P(M_f) > 0$ for every $\beta > 0$. To see that ii(b) is true take the derivative of D_M with respect to M :

$$\frac{\partial D_M}{\partial M} = \frac{n(1-w)}{N} \left(\underbrace{\frac{(1-m)w}{(1-M(1-w))^2}}_{\text{Complements}} - \underbrace{\frac{m}{(M+w(1-M))^2}}_{\text{Substitutes}} \right).$$

The denominators of the terms labeled ‘‘Complements’’ and ‘‘Substitutes’’ are both positive. Therefore, for m near enough zero, Complements dominates and the derivative $\partial D_M / \partial M$ is positive, whereas for m near enough one Substitutes dominates and the derivative is negative. For intermediate values of m , the sign of the derivative may depend on M . \square

Proposition 5 (Directed Homophily Decomposition). *The overall directed homophily decomposes as follows:*

$$P(M_m) - P(M_f) = E\left[\frac{m^c}{m}P(M|m, c) - \frac{1-m^c}{1-m}P(M|f, c)\right] + \frac{1}{m(1-m)}\text{Cov}[m^c, M^c] \quad (14)$$

Proof. (Proposition 5) Applying the proof of Proposition 3 to female (by symmetry) and substituting $P(M_f) = 1 - P(F_f)$ yields :

$$M - P(M_f) = E\left[\frac{1-m^c}{1-m}(M^c - P(M|f, c))\right] + \text{Cov}\left[\frac{m^c}{1-m}, M^c\right]$$

Hence

$$\begin{aligned} P(M_m) - P(M_f) = & E\left[\frac{m^c}{m}(P(M|m, c) - M^c) + \frac{1-m^c}{1-m}(M^c - P(M|f, c))\right] \\ & + \frac{1}{m(1-m)}\text{Cov}[m^c, M^c] \end{aligned}$$

\square

B Homophily and Relationship Persistence

The above analysis relied on a cross-section data. In this section, longitudinal data on the evolution of the network of referrals over several years are used to

estimate the dynamics in referral relationships with respect to gender. I find same-gender links persist longer in time, suggesting a dynamic foundation for the static excess of same-gender links.

For the study of the persistence of referral relationships, I estimate the following specification:

$$P(Y_{jk,t+1} = 1 | Y_{jk,t} = 1, g, X) = \frac{e^{\eta_{jkt}}}{1 + e^{\eta_{jkt}}} \quad (15)$$

using data on all dyads (j, k) such that $Y_{jk,t} = 1$, where $Y_{jk,t} = 1$ if j referred to k at period t and $Y_{jk,t} = 0$ otherwise, and $\eta_{jkt} := \alpha_j + \beta \mathbf{1}_{g_j = g_k} + \delta X_{jkt}$. That is, (15) estimates the probability of links (referral relationships) existing at t would still exist at $t + 1$. Each dyad is included only once, for the first year it is observed. Because only existing links are considered, no sampling is necessary for estimating this specification: all observed dyads are used.

Existing relationships are relatively more likely to persist between same-gender providers. Table A1 shows different estimates of link persistence, obtained from the sample of all initially connected dyads (physicians with referral relationships at the base year, defined as the first year they were observed in the data). Both logit and linear estimates with two-way fixed effects (for doctors and for specialists) show that same-gender links are more likely than cross-gender links to carry on to the following year (Columns 1–2). Columns (3) and (4) estimate separately for male and female doctors the probability of links persisting, again using physician fixed-effects to account for individual heterogeneity in the persistence of relationships. Consistent with the findings above—that male are much more likely to receive referrals, both male and female doctors’ relationships with male specialists are more persistent, but persistence is significantly higher for male doctors than it is for female doctors ($p < 0.001$). That is, same-gender relationships persist relatively longer in time.

Table A1: Estimates: Link Persistence

	Link Persists Next Year			
	(1)	(2)	(3)	(4)
	Logit	FE	FE	FE
Same Gender	0.044	0.014		
	(16.2)	(24.0)		
Male Doctor	0.069			
	(16.3)			
Male Specialist	0.16		0.029	0.0062
	(57.4)		(50.4)	(5.89)
Similar Experience	0.0046	0.0011	0.0016	0.00085
	(38.1)	(39.5)	(55.3)	(15.8)
Same Hospital	0.12	0.027	0.030	0.027
	(28.5)	(29.5)	(31.6)	(14.3)
Same Zip Code	0.16	0.097	0.092	0.076
	(55.1)	(145.1)	(129.9)	(56.3)
Same School	0.088	0.013	0.015	0.014
	(26.9)	(17.1)	(20.0)	(9.09)
Constant	-0.81			
	(-193.7)			
Specialty (Specialist)	No	No	Yes	Yes
Obs. (j,k)	7,255,778	7,204,471	5,734,596	1496658
Rank	8	5	58	58
R^2		0.20	0.10	0.11
N. Cluster	280,750	255,507	191,647	64,579
FE1 (Doctors)		255,507	191,647	64,579
FE2 (Specialists)		237,363		

Notes: t statistics in parentheses. Results of link persistence estimates. Column (1) shows estimates (5) for 2008–2012. Data consist of an observation for each linked dyad (j, k) , for the first year it was observed in the data. The dependent binary variable is 1 if the link between the doctor j and the specialist k continued during the subsequent year. Same gender is a dummy for the specialist and doctors being of the same gender. Male specialists/doctor is a dummy for the specialist/doctor being male. Similar Experience is negative the absolute difference in physicians' year of graduation. Column (2) shows linear estimates with two-way fixed effects (for doctor and for specialist) using the same data. Columns (3) and (4) show linear estimates with fixed-effects only for doctor, estimated separately for female (3) and male (4) doctors. Sample size is restricted by the availability of medical school data. Results excluding school affiliation are very similar. All standard errors are clustered by doctor.

C The Earnings Gap with Extreme Bias

In this section, I study the relationship between the earnings gap and the gender mix of physicians for different levels of gender bias in referrals. For small to moderate levels of gender bias, what determines the sign and size of the gender gap in earnings is mostly the gender distribution of doctors: the more of them are male, the greater the gap in favor of male specialists. As seen in Table 7, the gender gap in earnings for the estimated bias of 10% depends mostly on m , the fraction of males upstream, and varies only a little with M , the fraction of males downstream. This fact is more generally true for small levels of bias, as can be seen by linearly approximating the gap, i.e., the difference in average demand between the genders, around $\beta = 0$:

$$Gap = D^F - D^M \approx (2m - 1)\beta + O(\beta^2) \quad (16)$$

That is, what matters for the size (and the sign) of the earnings gap is the fraction of males upstream: when they are the majority, men get more work downstream, and vice versa.

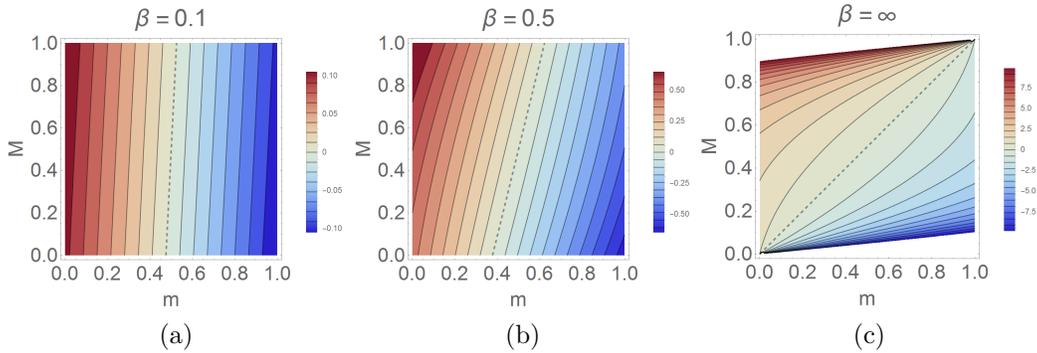
In fact, the gap mostly depends on the gender distribution upstream even for relatively high levels of bias (Figure A1). However, for extremely high levels of gender bias, both upstream and downstream majorities matter:

$$\lim_{\beta \rightarrow \infty} Gap = \frac{m - M}{M(1 - M)} \quad (17)$$

Specifically, when doctors refer *only* to specialists of their own gender, then the gender whose upstream fraction is greater than its downstream fraction gets more referrals.²⁴

²⁴I thank Alexander Frankel for bringing this case to my attention.

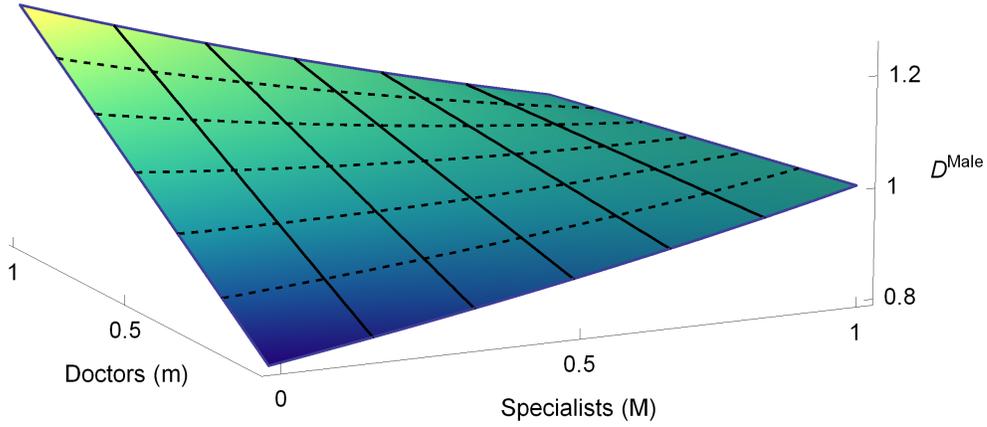
Figure A1: The Gender Earnings Gap With Different Levels of Bias



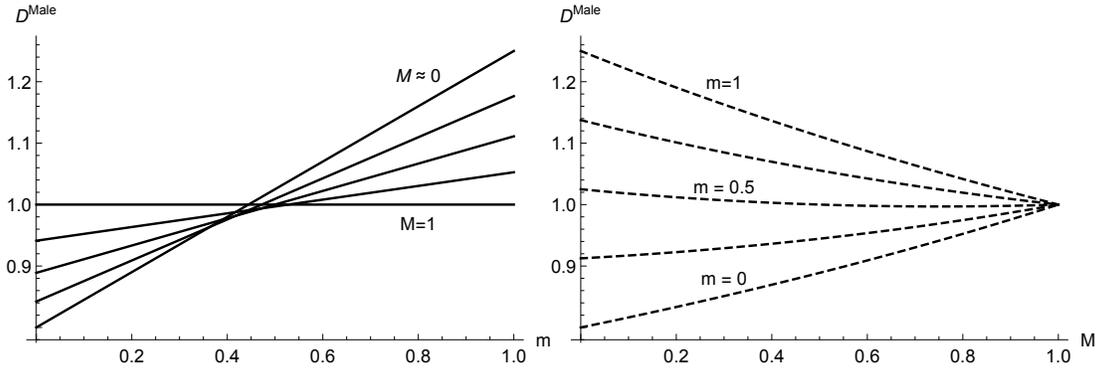
Colored contour plots of the gender earnings gap, $D^F - D^M$, for different fractions of males upstream m and downstream M , each for a different level of bias β . Blue (right) and red (left) darker shades reflect greater demand for male and female specialists, respectively. The zero-gap contours are dashed. For (a) the estimated level of bias for U.S. physicians ($\beta = \hat{\beta} = 0.10$), and even for (b) much higher levels of bias ($\beta = 0.50$), the sign and size of the gender earnings gap mostly depend on the fraction of males upstream. In contrast, for (c) extreme bias ($\beta = \infty$), a bias that reflects lexicographic preferences, the gap depends on the relative fractions of males (females) upstream and downstream.

D Additional Tables and Figures

Figure A2: Average Specialist Demand with Gender-Biased Preferences



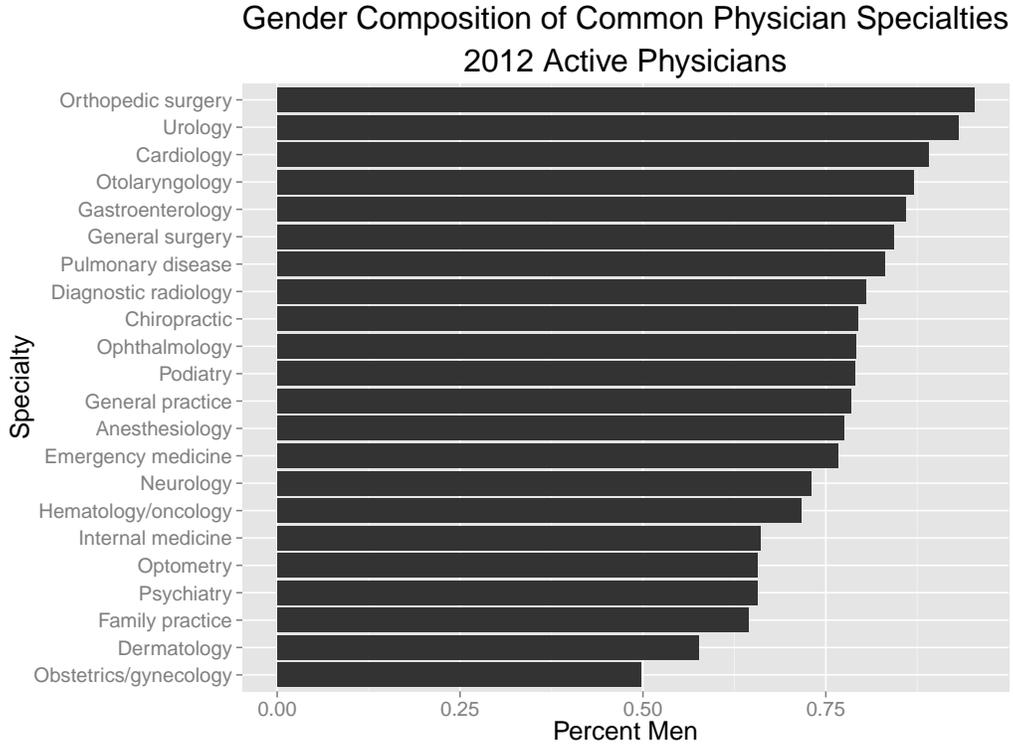
(a) Demand for Male Specialists over the Fractions of Male Doctors and Male Specialists



(b) Demand for Male Specialists and Male Doctors (c) Demand for Male Specialists and Male Specialists

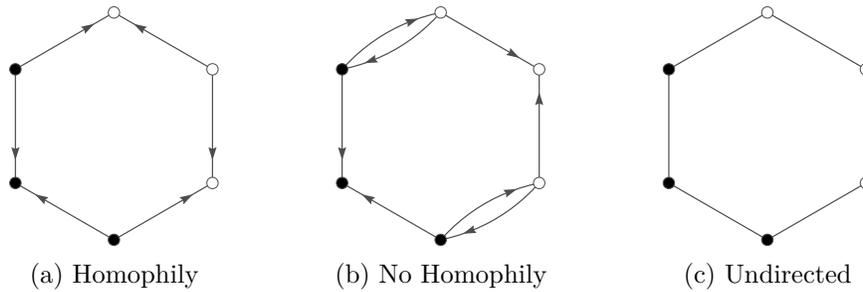
Notes: Average male specialist demand as a function of the fractions of male doctors and male specialists, with gender-biased preferences, i.e. $\beta > 0$ (calculated from the model with $\omega = 0.8, \eta = 1$). The surface in Panel (a) depicts the average demand D^{Male} , a function of the fractions of both male doctors, m , and male specialists, M . Panel (b) shows different cross sections of D^{Male} for different levels of M . Panel (c) shows different cross sections D^{Male} for different levels of m . Demand for male specialists is increasing the more doctors upstream are male. Whether specialists of the same gender substitute or complement each other depends on whether they are of the same gender as the upstream majority.

Figure A3: Male Fraction of Physicians in Common Medical Specialties



Notes: Percent of active physicians (with any claims) who are male, for the most common specialties by overall number of physicians. Columns are sorted so specialties with the greatest male shares are at the top.

Figure A4: Homophily and Link Direction



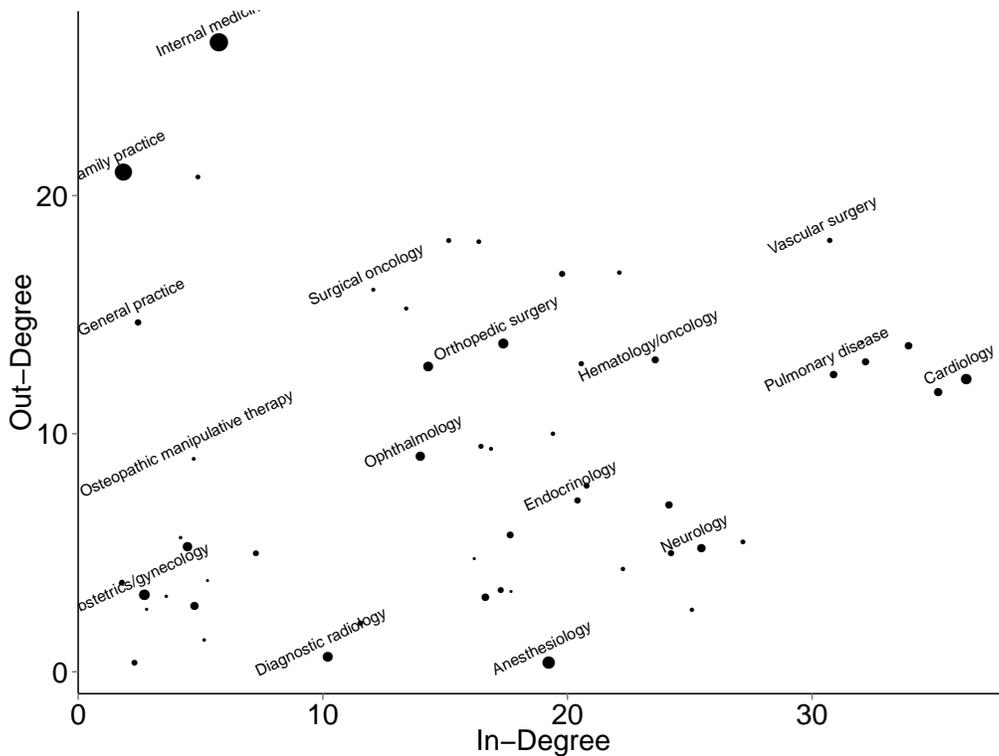
Because referrals are asymmetric, link direction is important in defining homophily: the network (a) exhibits homophily while (b) does not, a difference concealed in their undirected counterpart (c). More generally, this example speaks against treating asymmetric relationships as if they were symmetric when studying homophily.

Table A2: Medicare Referrals by Gender

From	A. Referrals			B. Percent of Outgoing			C. Percent of Incoming		
	To		To			To			
	F	M	F	M	Total	F	M	Total	
f	420,976	1,712,510	f	19.73	80.27	100	f	24.74	19.36
m	1,280,691	7,130,872	m	15.23	84.77	100	m	75.26	80.64
							Total	100	100

Notes: Referral counts and percentages, by gender of referring and receiving physician. Because services are sometimes billed on several separate claims, multiple referrals of the same patient from a doctor to a specialist are counted as one. Source: 20% sample of Medicare physician claims for 2012.

Figure A5: Average Number of Referral Relationships by Medical Specialty



Notes: Degree-heterogeneity is to be expected because doctors in different specialties play different roles in routing patients: some mostly diagnose and refer out, others mostly receive referrals and treat. The figure shows degree distribution by specialty for 2012 referrals: Out-degree is the average number of physicians to whom a physician referred patients during the year. In-degree is the average number of physicians from whom a physician received referrals. Physicians with neither incoming nor outgoing referrals during the year were excluded. Point diameter is proportional to the square root of the number of practitioners in a specialty. Common specialties are labeled. See Table A3 for the data used to generate this figure.

Table A3: 2012 Average Degree by Specialty

	Specialty	In-Degree	Out-Degree	Physicians
1	Internal medicine	5.8	26.4	86,220
2	Family practice	1.9	21.0	74,638
3	Anesthesiology	19.2	0.4	33,434
4	Obstetrics/gynecology	2.7	3.2	22,871
5	Cardiology	36.3	12.3	21,714
6	Orthopedic surgery	17.4	13.8	19,411
7	Diagnostic radiology	10.2	0.6	18,768
8	General surgery	14.3	12.8	18,011
9	Emergency medicine	4.5	5.2	16,065
10	Ophthalmology	14.0	9.1	15,702
11	Neurology	25.5	5.2	11,469
12	Gastroenterology	35.2	11.7	11,178
13	Psychiatry	4.8	2.7	10,861
14	Dermatology	16.6	3.1	8,624
15	Pulmonary disease	30.9	12.5	8,272
16	Urology	33.9	13.7	8,234
17	Otolaryngology	24.2	7.0	7,666
18	Nephrology	32.2	13.0	7,105
19	Hematology/oncology	23.6	13.1	7,019
20	Physical medicine and rehabilitation	17.7	5.7	6,224
21	General practice	2.5	14.7	4,853
22	Endocrinology	20.4	7.2	4,534
23	Infectious disease	24.2	5.0	4,492
24	Neurosurgery	19.8	16.7	4,010
25	Radiation oncology	17.3	3.4	3,933
26	Rheumatology	20.8	7.8	3,765
27	Plastic and reconstructive surgery	7.3	5.0	3,759
28	Pathology	2.3	0.4	3,627
29	Allergy/immunology	11.5	2.0	2,768
30	Pediatric medicine	1.8	3.8	2,695
31	Medical oncology	20.6	12.9	2,507
32	Vascular surgery	30.7	18.1	2,486
33	Critical care	16.5	9.5	2,046
34	Thoracic surgery	15.2	18.1	1,886
35	Interventional Pain Management	27.2	5.5	1,655
36	Geriatric medicine	4.9	20.8	1,597
37	Cardiac surgery	16.4	18.0	1,526
38	Colorectal surgery	22.1	16.8	1,161
39	Pain Management	22.3	4.3	1,055
40	Hand surgery	19.4	10.0	1,047

Notes: A link represents referral relationships with another physician from any specialty; specialties with less than 1,000 are included but not shown, due to space constraints.

Table A4: Estimates of Directed Homophily Using Disaggregated Data

	Dep. Var.: Male Specialist				
	(1)	(2)	(3)	(4)	(5)
Male Doctor	0.045*** (55.8)	0.039*** (52.8)	0.038*** (50.4)	0.035*** (46.9)	0.042*** (51.5)
Male Patient				0.021*** (81.8)	0.037*** (56.9)
Male Doctor x Male Patient					-0.019*** (-27.4)
Specialty (both)	No	Yes	Yes	Yes	Yes
Experience (both)	No	No	Yes	Yes	Yes
Obs. (Triples)	10,545,049	10,545,049	10,127,806	10,127,806	10,127,806
Clusters (Doctors)	385,104	385,104	382,924	382,924	382,924
R Sqr.	0.00242	0.0689	0.0989	0.0997	0.0998

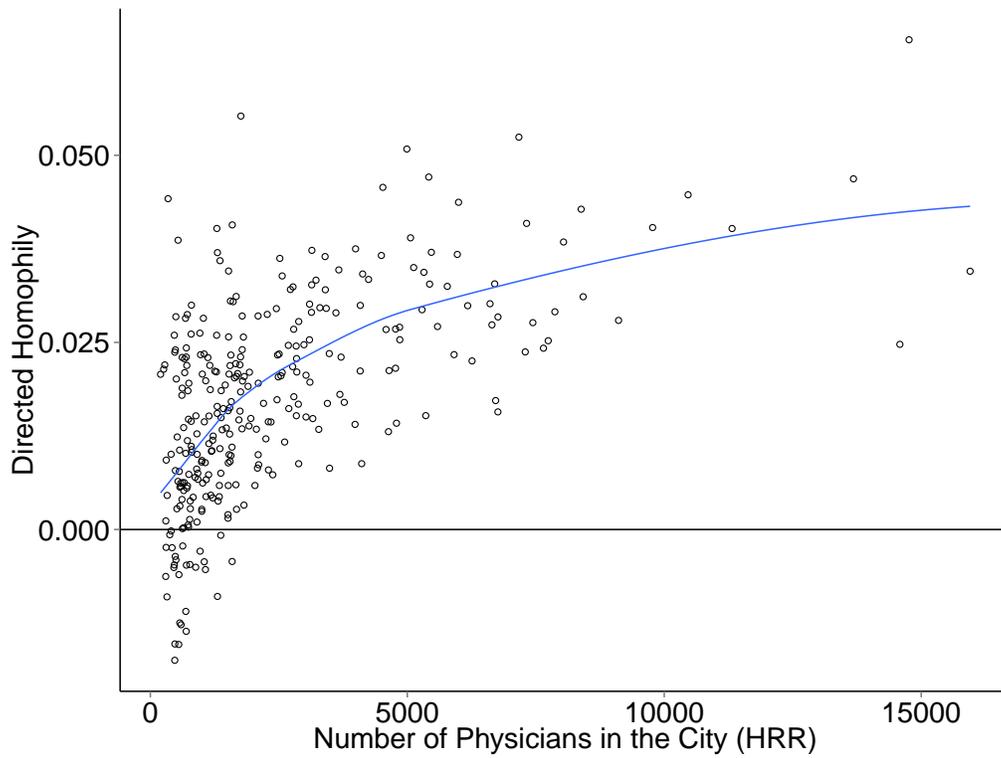
Notes: t statistics in parentheses. Estimates of directed homophily using one observation for each unique triple of a doctor, a specialist and a referred patient. The sample consists of all such triples for 2012, for a sample of 20% of Medicare patients.

Table A5: Homophily Estimates for Different Age Groups

	Percent of Referrals to Male Specialists		
	Young	Old	All
Male Doctor	0.038 (0.0011)	0.044 (0.0015)	0.040 (0.00090)
Male Patients (pct)	0.028 (0.0024)	0.031 (0.0026)	0.029 (0.0018)
Constant	0.79 (0.0078)	0.81 (0.0040)	0.80 (0.0032)
Specialty (Doctor)	Yes	Yes	Yes
Experience (Doctor)	Yes	Yes	Yes
Obs. (Doctors)	200,670	184,315	384,985
Rank	57	57	57
Mean Dep. Var.	0.82	0.83	0.82
R^2	0.035	0.041	0.039

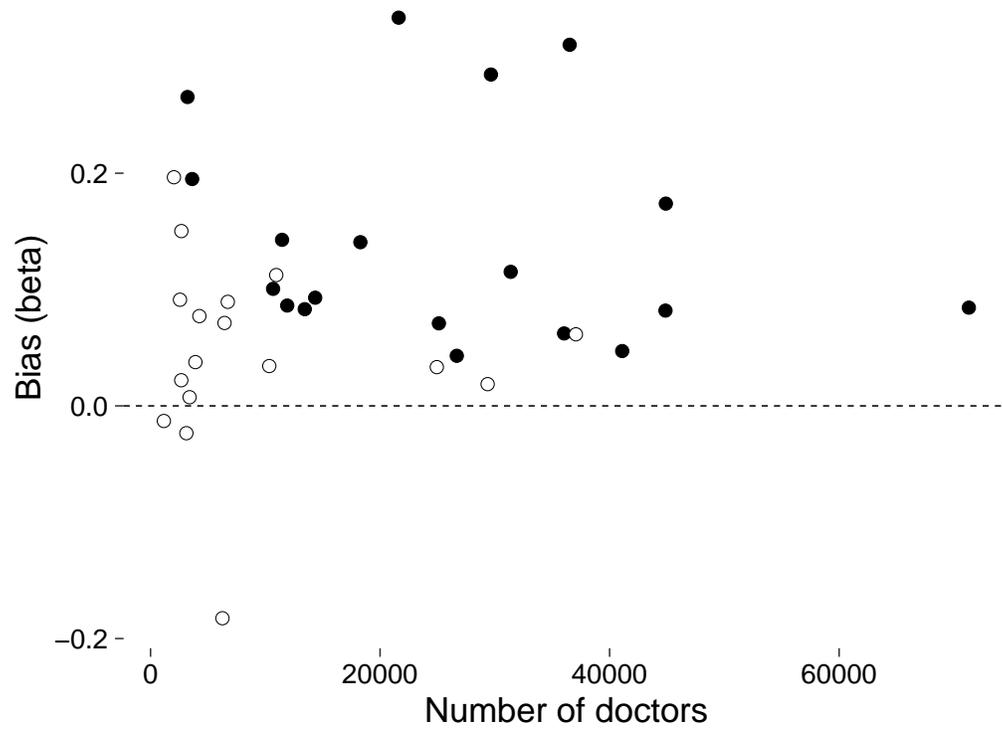
Notes: standard errors in parentheses. OLS estimates of (2) are shown for three subgroups: young doctors (below median experience of 24 years, Column 1); old doctors (above median experience, Column 2); and all doctors together (Column 3). Despite the similar opportunity pools they face, older doctors exhibit stronger average directed homophily than younger ones.

Figure A6: Homophily and Market Size



Notes: Homophily estimates of (2), estimated separately for each local physician market (Dartmouth Hospital Referral Region), over the overall number of physicians in the market (men and women). The line is local regression (LOESS) fit. Beyond the mechanical effect of a reduction in variance of the estimates with sample size, estimated homophily is also greater for larger markets.

Figure A7: Conditional-Logit Estimates of Gender Bias, by Specialty



Notes: Estimates of β , the gender bias, from equation (5) with the sample in Table 5, separately for each medical specialty of the receiving physician. Black circles denote estimates that are significantly different than zero ($p < 0.05$).

Table A6: Homophily Estimates with Weighted Links

	Percent Referrals to Male Specialists, by:			
	(1)	(2)	(3)	(4)
	Links	Patients	Claims	Dollars
Male Doctor	0.038 (43.2)	0.040 (44.8)	0.040 (42.7)	0.040 (41.4)
Percent Male Patients	0.029 (16.6)	0.029 (16.5)	0.029 (16.1)	0.029 (15.4)
Cons.	0.80 (262.2)	0.80 (254.3)	0.80 (243.8)	0.81 (243.9)
Specialty (Doctor)	Yes	Yes	Yes	Yes
Experience (Doctor)	Yes	Yes	Yes	Yes
Obs. (Doctors)	384,985	384,985	384,985	383,054
R^2	0.0384	0.0394	0.0360	0.0368

Notes: t statistics in parentheses. OLS estimates of (2) using different definitions of link weights: The first column shows results for unweighted links. Columns 2–4 show results for different weights: number of patients, number of claims, and Dollar value of services.

Table A7: Conditional-Logit Estimates: Referral Probability, with Interaction Terms

Doctor and Specialist:	Doctor Referred to Specialist			
	(1)	(2)	(3)	(4)
Same Gender	0.0841 (35.81)	0.0662 (16.49)	0.104 (29.65)	0.0758 (12.88)
Male Specialist	0.175 (73.70)	0.175 (73.26)	0.165 (46.03)	0.164 (45.69)
Same Hospital	3.114 (720.15)	3.072 (579.18)	2.941 (540.97)	2.887 (414.00)
Same Hospital x Same Gender		0.0598 (13.65)		0.0770 (12.49)
Same Group	1.346 (178.27)	1.372 (151.85)	1.320 (135.26)	1.344 (111.58)
Same Group x Same Gender		-0.0386 (-5.24)		-0.0354 (-3.43)
Same Zipcode	1.074 (219.53)	1.065 (163.68)	1.054 (164.04)	1.047 (118.78)
Same Zipcode x Same Gender		0.0130 (2.08)		0.0104 (1.20)
Similar Experience	0.128 (131.66)	0.120 (75.64)	0.131 (92.95)	0.123 (52.37)
Similar Experience x Same Gender		0.0117 (6.28)		0.0110 (3.99)
Same Med. School			0.209 (49.96)	0.206 (28.35)
Same Med. School x Same Gender				0.00447 (0.54)
Specialist Experience	Yes	Yes	Yes	Yes
Obs. (Dyads)	14,555,821	14,555,821	6,712,241	6,712,241
Clusters (Doctors)	367,370	367,370	242,579	242,579
Pseudo R Sqr.	0.361	0.361	0.347	0.347

Notes: Results of conditional logit estimates of (5) for 2012, including interaction terms (denoted by ×). See Table 6 notes for variable definitions and details.