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

Dan Zeltzer*

In this paper, I assess the extent to which gender gaps in earnings may be driven by physicians' preference for working with specialists of the same gender. Analyzing administrative data on 100 million Medicare patient referrals, I provide robust evidence that physicians refer more to others of their same gender (i.e., referrals exhibit gender *homophily*). I show that homophily in referrals is predominantly driven by physicians' decisions, rather than by endogenous sorting of physicians or patients. As 75% of referring physicians are men, my estimates suggest that gender homophily in referrals makes, all else being equal, demand for female physicians 5% lower than demand for male physicians, thus contributing to the persistence of gender inequality. Overall, my results point to the positive externality associated with increased female participation in medicine, and perhaps in other contexts where networking is important.
JEL Codes: I11, J16, L14.

*Tel Aviv University, School of Economics, Tel Aviv, Israel 6997801 (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, University of Pennsylvania, University of Texas at Austin, University of Maryland, the EIEF, the IDC, Ben-Gurion University, University of Haifa, The Hebrew University, Tel Aviv University, and the AEA Meeting in Chicago. 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

Why women are still underrepresented in top career position and earn less than their male counterparts remains unclear. 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 Altonji and Blank, 1999; Blau and Kahn, 2000; Niederle and Vesterlund, 2007; Bertrand, 2011; Azmat and Petrongolo, 2014). In this paper, I study in detail one such channel: the extent to which gender gaps in earnings may be driven by biased professional networks.

It is well documented that individuals tend to develop relationships with others similar to themselves. I study in detail this tendency, known as *homophily*, and quantify its contribution to the gender earnings gap in the medical profession. I exploit data on 100 million patient referrals among half a million U.S. physicians from Medicare, a unique context in which referral relationships and payments are jointly observed. This paper makes three main contributions. First, it shows that doctors refer more to specialists of their same gender.[1] Second, it shows that this gender bias in referrals is predominantly driven by doctors' decisions rather than by endogenous sorting of physicians or patients. Third, it shows 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.

The main identification challenge is separating gender bias in referrals from other differences between the genders. For example, male physicians may receive more referrals because on average they are more experienced than their female counterparts (as female entry to medicine is rather recent), or because men work longer hours. To test whether doctors refer to specialists of their same gender disproportionally, I define a new homophily measure, *directed homophily*, that compares the fraction of referrals made to male specialists between male and female doctors. I show that if referral rates were fully explained by differences between male and female specialists, then, all else being equal, referral decisions should be independent of the gender of the doctor making them. Specifically, there should be no directed homophily.[2]

A second concern is that different doctors choose specialists from different pools. Thus, if physicians endogenously sort by gender into market segments (e.g. hospitals, or medical

[1]Throughout this paper the terms *doctor* and *specialist* are used to denote the role of a physicians 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.

[2]Oft-used homophily measures disregard this concern, 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).

specialties), doctors could be more exposed to—and thus more likely to refer to—specialists of their own gender. Even absent sorting, quantifying the underlying gender bias from 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 choice sets, I use choice-based sampling (See McFadden, 1984; Manski and Lerman, 1977).

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. Furthermore, controlling for the gender of patients shows that results are not explained by patients' preference for physicians of their same gender.

Most of the bias in referrals is due to doctors' choices. Estimates of the discrete choice model suggest that doctors' preferences exhibit significant gender bias: 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.

Given that most referrals are 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, all else being equal, demand for female specialists 5% lower than demand for male specialists. This gender disparity in demand would vanish if doctors' choices were unbiased, or, alternatively, if their population were gender-balanced. 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]

Previous research has shown that gender gaps in earnings exist among highly skilled occupations, such as 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 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, where reimbursement rates are fixed and therefore pay gaps are not due to differences in compensation, 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.

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 where much of the work depends on referrals, such as most surgical specialties, an area in which there are currently still 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 even be 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 three existing lines of work: the

[3] Specifically, 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.

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; Kossinets and Watts, 2006), and has been shown to contribute to network segregation and limit the flow of information, both in theory (Bramoullé et al., 2012a; Golub and Jackson, 2012), and in practice (Currarini, Jackson, and Pin, 2009; Himelboim, 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 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.[4] 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 account 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 its impact on pay disparities.[5] For each encounter of a patient with a physician, the data contain the follow-

[4]Medicare is the federal health insurance program for people who are 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).

[5]For 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 2016.

ing: 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.[6] To protect the privacy of patients, no statistics are reported for demographic cells based on fewer than 11 individual patients. Thanks to the large sample size, such cells are rarely encountered.

Data are combined with additional data on physician gender and other characteristics from Physician Compare, a public CMS database that provides information on physicians and other health care professionals who provide Medicare services.[7] 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 more descriptive information see the supplementary material). These data are further combined with beneficiary sex 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, although specialties related to elderly patients are over represented. By volume, Medicare billed physician 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.[8] 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 probability 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

[6]Claims 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 2016. About a third of the remaining claims record a referring physician provider.

[7]Physician Compare Database, `https://data.medicare.gov/data/physician-compare` Accessed December 2016.

[8]2012 National Health Expenditure Accounts (NHEA).

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.

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 difference in the number of colleagues men and women physicians work with. Conditional on making any referrals, doctors refer to 16 specialists an 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 differences in specialization, because the average number of working referral relationships each physician has varies greatly by medical specialty, and men and women are not equally represented across specialties (see supplementary materials for more details). It is therefore 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 constraint in that way. However, evidence suggests that referrals are still an important determinant of care trajectory (Johnson, 2011; Barnett et al., 2011, 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.[9] 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 in 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

[9] "Dartmouth Atlas of Healthcare", `http://www.dartmouthatlas.org/tools/downloads.aspx?tab=39`. Accessed December 2016

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

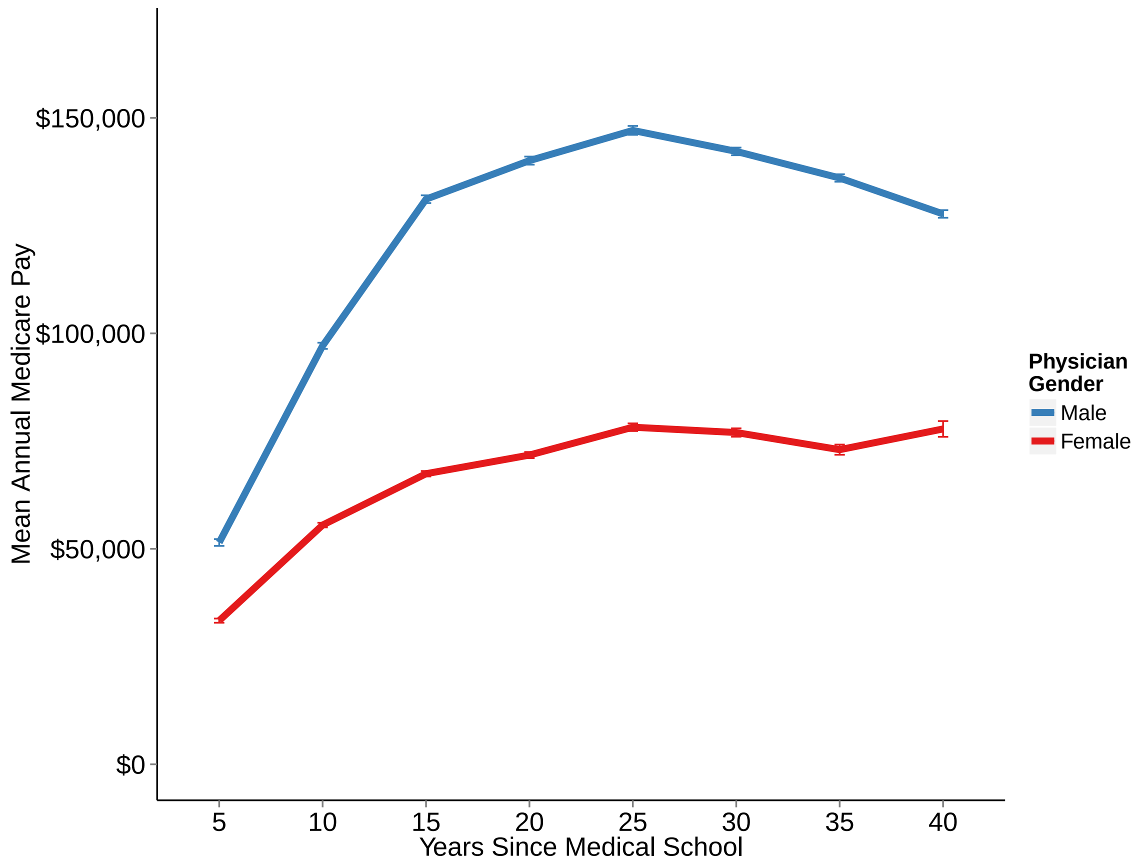

Notes: Source: 20% sample of Medicare physician claims for 2012. Mean Annual Medicare Pay is total annual payments (by all payers) to physicians for Medicare services, multiplied by 5 to adjust for sampling. Years are since medical school graduation (bin maximum, e.g. 10 stands for 6-10).

Table 2: The Gender Pay Gap for Medicare Physicians

	Dependent variable:					
	Log(Annual Pay)					
	(1)	(2)	(3)	(4)	(5)	(6)
Male Physician	0.668	0.654	0.468	0.361	0.337	0.340
	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)	(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.

the gross earnings gap by estimating a standard (log) annual pay equation:

$$log(Pay_k) = \beta \mathbb{1}_{g_k=M} + \delta X_k + \varepsilon_k \tag{1}$$

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), a dummy for city (HRR), a dummy for the 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 obstetrician-gynecologists are women, but less than 6% of active orthopedic surgeons are women (Figure A.1). Consistent with previous works (e.g., Bertrand, Goldin, and Katz, 2010), 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) to estimate a median gap ranging between 16% and 25% (18–30 log points) among U.S. physicians that was quite persistent throughout the period of 1987 to 2010. Using Physician Surveys administered between 1998 and 2005, Weeks, Wallace, and Wallace (2009) find women earn about a third less than men. My estimates of the gender pay gap in Medicare are also on par with pay gaps in other highly skilled occupations: Bertrand, Goldin, and Katz (2010), using data from MBA graduates working in the financial and corporate sectors, fund a gross gap of almost 60 log points 10 to 16 years after graduation, and Azmat and Ferrer (2016) find large gaps in hours billed and new client revenue between male and female lawyers.

The earnings gap in medicine, as in previously documented gaps, partly reflects known differences between the genders, including differences in labor supply. However, much of it remains unexplained. The rest of this paper focuses on studying in detail and quantifying one additional contributing channel: how gender-biases in referral relationships among physicians generate differences in opportunities men and women experience based on their gender.

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 disproportionally 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 preferences 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.[10] (Throughout, I use lowercase and uppercase to index doctors and specialists, respectively.) I define directed homophily as follows:

Table 3: Overall Directed Homophily (DH) in Medicare

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

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

[10]Let G_g be the average fraction of referrals doctors of gender $g \in \{m, f\}$ send to specialists of gender $G \in \{M, F\}$.[11]. 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 g send to specialists of gender G.

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.[12] 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 such measure always has the same sign as directed homophily. Preserving link direction is important (see supplementary material for an example). Directed homophily can also be redefined to admit weighted links. Weights reveal whether same-gender referrals are not only more likely, but also more voluminous.[13]

Directed homophily is not driven by baseline imbalance in the gender distribution of doctors or specialists: if most specialists are men then both male and female doctors are expected to refer more to male specialists, but not differentially so.

It is useful to compare directed homophily with a different homophily measure: *inbreeding homophily.*[14] This measure uses population shares as the baseline:

Definition 2 (Inbreeding Homophily)**.** Male 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 $M_f < M$).

---

[12] Referrals exhibit *directed heterophily* if they are biased towards the other gender, that is, $DH < 0$.

[13] To adapt directed homophily to weighted links, simply redefine 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.

[14] This measure, attributed to Coleman (1958), has long been used in sociology (see McPherson, Smith-Lovin, and Cook, 2001; Thelwall, 2009), and more recently in economics. See Currarini, Jackson, and Pin (2009); Bramoullé et al. (2012b); Currarini and Vega-Redondo (2013) (normalized or approximated variants are often used). Golub and Jackson (2012) define a different measure, *spectral homophily*: the second-largest eigenvalue of a matrix that captures relative densities of links between various pairs of groups; this measure captures a notion of segregation of the network: how "breakable" it is to two groups with more links within them and less between them. Neither spectral homophily or its simpler estimate, *degree-weighted homophily*, imply directed homophily nor are they implied by it.

Note that inbreeding homophily by both genders immediately implies directed homophily, while the reverse is not true, e.g. if $M_m > M_f > M$.

Unlike inbreeding homophily, directed homophily does not use population gender shares as a benchmark. Therefore, directed homophily is insensitive to differences between the genders in the propensity to send or receive referrals. For example, if both genders referred more patients to male specialists only because male specialists were more qualified 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 specialists. Put differently, directed homophily measures relative, not absolute, gender differences in referrals. With unobserved heterogeneity, absolute differences are generally not identified.

Table 4: Medicare Referrals by Gender

		A. Referrals			B. Percent of Outgoing				C. Percent of Incoming	
		To			To				To	
		F	M		F	M	Total		F	M
From	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.

3.2 Patterns of Homophily in Physician Referrals

Overall, Medicare referrals exhibit significant directed homophily (Table 4). Before accounting for any differences in characteristics, 84.77% of the referrals made by male doctors are to male specialists, compared to only 80.27% of the referrals made by female doctors. That is, compared to female doctors, male doctors refer more to male specialists, and vice versa. 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).

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

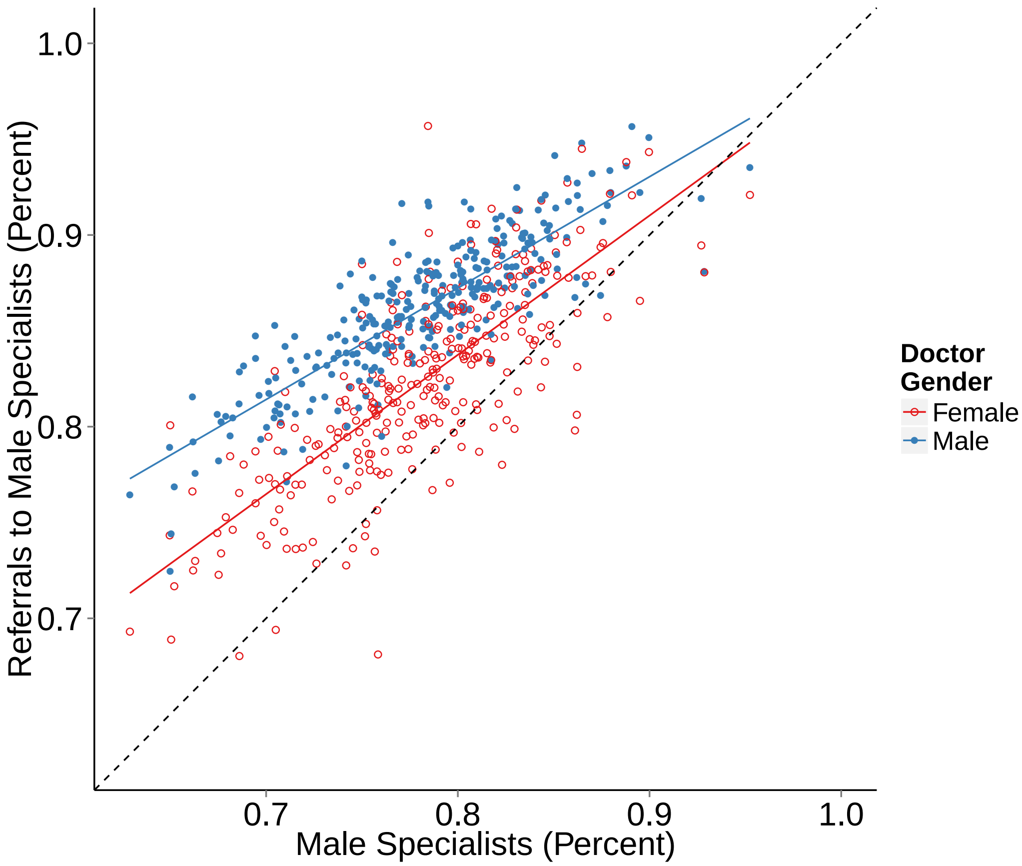

Notes: For each local physician market (Dartmouth Hospital Referral Region), average fractions of referrals from male and from female doctors to male specialists are plotted over the fraction of male specialists in the market. Each of these 306 local U.S. markets is thus represented by two vertically 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.

Directed homophily can be estimated separately 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:

$$M_j = \alpha_1 + \beta_1 g_j + \delta_1 X_j + \varepsilon_j, \tag{2}$$

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

A potential explanation for observed homophily, rather than bias in physicians' preferences, is that physicians sort into market segment by gender, which makes the pool of specialists biased towards their own gender. Some control for sorting can be obtained from a variant of (2):

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

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 preferences.

Another potential 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 preferences 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 mix of patients hardly changes estimated directed homophily, suggesting homophily in physician referrals is not driven by homophily on behalf of patients. Table 5 shows estimates for directed homophily obtained from individual physician data. There is a 4.3 percentage points difference between their referral rate and that of female doctors of similar specialty and experience (Column 2). Estimated homophily 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 A.1). Note that controlling for the gender of patients accounts for patients' gender-preferences that are correlated with their gender, but not for residual gender preferences that are uncorrelated with patient gender. In light of how little of homophily is explained by the former, it would be surprising were much of it explained by the latter.[16]

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

[16]The use of patient fixed effects analogous to the hospital fixed effects in (3) is limited by the fact that patients seldom obtain referrals from different doctors of the same specialty. Restricting attention to patients who have been referred by two such doctors of opposite genders yields a small sample and one

Table 5: Estimates of Directed Homophily

	Percent of Referrals to Male Specialists					
	OLS				FE	
	(1)	(2)	(3)	(4)	(5)	(6)
Male Doctor	0.053	0.043	0.040	0.040	0.029	0.030
	(62.7)	(49.1)	(44.8)	(44.0)	(30.5)	(32.6)
Percent Male Patients			0.029	0.028	0.031	0.043
			(16.5)	(14.7)	(16.1)	(23.4)
Cons.	0.79	0.81	0.80	0.81	0.82	0.78
	(1027.6)	(263.8)	(254.3)	(256.9)	(249.4)	(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 sub sample 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.

Homophily estimates are virtually unchanged when links are weighted. Appendix Table A.4 shows estimation results for different measures of referral volume: number of patients, number of claims, or overall dollar value of services. I also find that older doctors (with above-median experience) exhibit greater directed homophily than younger ones (Table A.3).

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 the patient gender 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 doctors may prefer to work with specialists of their same gender. However, directed homophily varies with the fraction of male specialists available (note that the fitted curves in Figure 2 are not parallel). It does not directly reflect the magnitude of an underlying preference bias. 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 preferences were biased. Therefore, identifying and quantifying the underlying bias in preferences 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 Preferences

This section examines the relationship between the observed homophily in referrals and the underlying gender bias in doctors' preferences and estimates the latter directly. Analyzing a discrete choice model where doctors choose specialists, I show that homophily decomposes to gender-biased preferences within market segments and sorting across such segments. I then use this model to estimate the gender bias in doctors' preferences. The bias in preferences is

possibly selected on patient preferences.

identified by comparing the average characteristics between chosen and unchosen specialists. I find that faced with a choice between otherwise similar male and female specialists, doctors are 10% more likely to refer to the specialist of their same gender.

4.1 A Model: Gender-Biased Referrals versus 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:

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

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:

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

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 an 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:

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

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 $\omega = e^{-\beta} \in (0, 1]$.[17] 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). Conversely, with unbiased preferences ($\beta = 0$) directed homophily is zero, as referral rates to men are common to doctors of both genders:

$$P(M_m) = M = P(M_f). \tag{7}$$

Clearly, 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

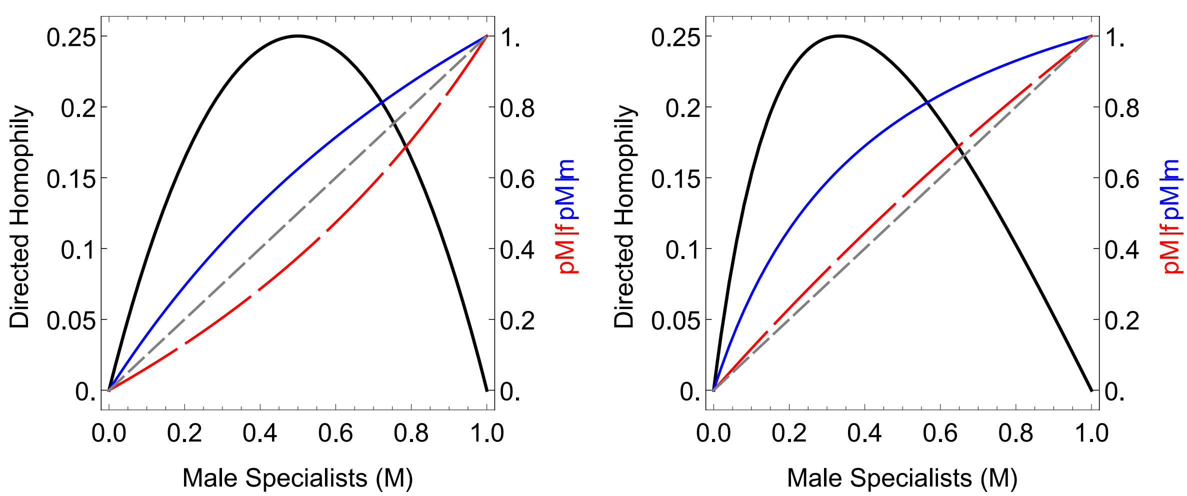

(a) Homogeneous Case ($\omega = 0.6, \eta = 1$) (b) Heterogeneous Case ($\omega = 0.6, \eta = 0.5$)

Notes: Probability of referrals to male specialists by male and female doctors and the difference: Directed Homophily for different fractions of male specialists, M, with gender-biased preferences ($\omega = e^{-\beta} = 0.6$). Case (a) shows gender-biased preferences in a homogeneous specialist population ($\eta = 1$): Male specialists receive more referrals than their fraction in the population from males, and less than this fraction from females. Case (b) combines gender-biased preferences with heterogeneity ($\eta = .5$): male specialists receive more referrals than their fraction in the population from both male and female doctors, but more from male than from female doctors.

An important implication of (6) is that observed homophily varies with the gender

[17] see appendix for derivation of this and other results

composition of the pool of available specialists: when the pool is more gender-balanced, observed homophily is greater (as illustrated in Figure 3a, and as seen before in Figure 2). With balanced pools, doctors' choices more strongly reflect their preferences. Conversely, when most specialists are of one gender, there is less room for choice and thus homophily is weaker. In the extreme cases where $M \in \{0, 1\}$, there is no homophily even if preferences are biased. The preference bias β is therefore a more portable parameter than directed homophily, which in part reflects differences among doctors in the mix of available specialists.

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:

$$P(M_m) = \frac{M}{M + \omega\eta(1 - M)} \geq \frac{\omega M}{\omega M + \eta(1 - M)} = P(M_f). \tag{8}$$

Regardless of gender-biased preferences, if $\eta < 1$ male specialists attract a disproportionally 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 (8) 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.

With heterogeneity ($\eta \neq 1$), directed homophily is a better measure of homophily than inbreeding homophily because it does not use M as the benchmark, but rather compares referrals of both genders against each other. For simplicity, for the rest of this section again assume homogeneity.

Homophily due to Sorting by Gender into Market Segments Apart from preferences, physicians 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 instead of a single market 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 $\text{Cov}(m^c, M^c) > 0$.[19]

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

[19]This definition extends to the more general case where K_j is specific to each doctor as: $\text{Cov}(m^j, M^{K_j}) > 0$, where $m^j = \mathbb{1}_{g_j = m}$ and M^{K_j} is the fraction of male in K_j.

Proposition 2 (Sorting-Based Homophily). *With sorting, referrals exhibit homophily when pooled together 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 preferences are unbiased ($\beta = 0$) so there is no homophily within each market.

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. Its equivalent for directed homophily is in the Appendix.

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

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

That is, the homophily observed when all markets are pooled together 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. Note that sorting could also dampen homophily, rather than augment it: If $\mathrm{Cov}[m^c, M^c] < 0$ then even if preferences are biased, overall homophily could be zero.

The proof of Proposition 3 only uses Bayes' rule to relate aggregate and market-specific referral probabilities. So, while it is natural to specify the probabilities $P(M_m|c)$ as done in the case of a single market discussed above, the proof of Proposition 5 does not rely on a specific parameterization of these probabilities: it only requires relevant moments to exist.

A corollary of Proposition 3 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), Proposition 3 shows that observed homophily in each market is a combination of preferences and sorting into unobserved market segments. Accounting for sorting is therefore required to identify preference bias. Inasmuch as sorting is accounted for, the directed homophily estimated in Section 3 provides suggestive evidence for the presence of gender bias in preferences. 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, I defined a homophily measure that identifies gender bias by comparing referrals between genders (thus "differencing out" unobserved differences between the genders), and by narrowing this comparison to within smaller market segments (thus accounted for sorting). In this section, I take a more direct approach, and estimate the parameters of the specialist choice model specified in equation (5). The main parameter of interest, the gender-bias in preferences, 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. For example, if doctors mostly refer within-hospitals, omitting hospital affiliation is only a problem if both doctors and specialists sort by gender into hospitals, i.e., if there is a correlation across hospitals between the gender mix of affiliated 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 below, 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 a conditional logit model 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:

$$log\frac{p_{jk}}{p_{jk'}} = \beta(\mathbb{1}_{g_j=g_k} - \mathbb{1}_{g_j=g_{k'}}) + \delta(X_{jk} - X_{jk'}). \tag{9}$$

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.

Since the opportunity pools 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, and controls for multiple other characteristics are included (e.g., distance) to capture other determinants of the decision. Rather, it only assumes that specialists from different markets or from different medical specialties are not substitutes. The model parameters, estimated using variation within city and specialty, capture the actual substitutability within those cells. Estimates are consistent under this sampling scheme (see McFadden, 1984).

4.3 Estimation 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 6 describes the sample used for estimating preference bias. It compares the average characteristics of specialists that are chosen with two randomly samples 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 went to the same medical school.

Table 6: Average Characteristics of Chosen versus Unchosen Specialists

	Doctor Referred to Specialist	
Doctor and 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. + 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 7, 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. Estimated separately by specialty, preference bias is stronger in all large-enough specialties (Figure A.5). The estimated bias is not positive only for small specialties, where doctors are likely too restricted in their choices to be able to express any gender preference. Results from including interaction terms in the estimated model suggest that the preference bias is somewhat stronger within hospitals, and somewhat weaker within groups, albeit only slightly (Table A.5). There is no additional increase in the probability of referral between doctors and specialists that graduated from the same medical school and are of the same gender.

The estimated gender bias is comparable with the previously estimated homophily (Table 5). 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. In comparison, the counterfactual

[20]Estimates represent odds ratios, but due to sparsity, estimates close to zero approximately equal the percentage increase in probability. That is, using the notation of (4), around zero $\beta \approx \dfrac{p_{jk}|g_j = g_k}{p_{jk'}|g_j \neq g_{k'}} - 1$.

Table 7: Conditional-Logit Estimates: Referral Probability

	Doctor Referred to Specialist					
Doctor and Specialist:	(1)	(2)	(3)	(4)	(5)	(6)
Same Gender	0.104	0.0983	0.0841	0.105	0.104	0.0844
	(55.27)	(41.96)	(35.81)	(29.79)	(29.65)	(27.71)
Male Specialist	0.197	0.169	0.175	0.165	0.165	0.194
	(103.46)	(71.43)	(73.70)	(46.09)	(46.03)	(63.15)
Same Hospital		3.116	3.114	2.945	2.941	2.803
		(721.31)	(720.15)	(541.87)	(540.97)	(568.47)
Same Practice Group		1.346	1.346	1.320	1.320	1.652
		(178.27)	(178.27)	(135.27)	(135.26)	(142.36)
Similar Experience			0.128	0.132	0.131	0.136
			(131.66)	(93.47)	(92.95)	(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. Schooling information is only partly available, thus Column (4) estimates the same specification as (3) with non-missing school data.

directed homophily with equal fractions of male and female specialists is 5 percentage points. That is, holding the underlying gender bias constant, homophily increases when choice sets are more gender balanced.

In addition to gender homophily, referrals also exhibit homophily on other dimensions: doctors prefer specialists of similar experience and specialists that attended the same medical school. A doctor and a specialist a decade closer in age have a 13% greater probability of referrals between them. The other dimension of affinity, having gone to the same medical school, is also a strong determinant of referrals, with doctors being 20% more likely to refer to same-school graduates. Since medical school data are partial, estimates with and without inclusion of same-school dummies are presented; they do not differ much.

Measuring the estimated role of gender in choices against the role of other attributes implies a social distance between the genders. There is a comparable effect on the likelihood of referrals for being of different gender and for having an age difference of ten years. Note 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.

In Appendix B, I further study the dynamics of referral relationships. I find that same-gender physicians are also more likely to maintain referral relationships over time. 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. Further developing the model introduced in Section 4 I show that when homophily is driven by gender-biased preferences, it leads to a disparity in demand between the genders. Specifically, the more referrals that are made by male doctors, the lower female specialists' demand. I directly test this predicted relationship using longitudinal data on Medicare payments for each specialist over 2008–2012. I find that female specialists indeed earn less—and male specialists earn more—when a greater fraction of primary care claims in their market is handled by male doctors. This contribution of biased referrals to the earnings gap is separate from possible differences between the genders in labor supply. Counterfactual calculations show that biased referrals explain 14% of the current gender gap in physician specialist earnings. Counterfactuals highlight that the resulting gap depends on the gender distribution of doctors and specialists. One way to eliminate the contribution of biased referrals on earnings is to balance the gender of referring doctors.

5.1 Homophily's Consequences for Gender-Disparities in Specialist Demand: Theory

Intuitively, when preferences are gender-biased, specialist receive fewer referrals when fewer doctors share their gender. The following proposition shows that the gender of fellow specialists matters too, in a more nuanced way. Whether same-gender specialists substitute or complement each other depends on the gender distribution 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 doctors are of his (or her) gender.*
- *(b) Same-gender specialists substitute for each other when most doctors are of their same gender, and complement each other when most doctors are of the opposite gender.*

The proof of Proposition 4 notes that demand for male specialist—the average number of referrals received (denoted by D^M)—is a weighted average of doctors' respective probability of referring to a male:

$$D^M = mP(M_m) + (1-m)P(M_f) \tag{10}$$

where m is the fraction of male doctors in the market (superscript c is omitted for clarity as all magnitudes are within markets), and assumes $|J| = |K|$ for tractability (see Appendix for general results). Substituting (6) into (10) and differentiating by m and M yields:

$$\frac{\partial D^M}{\partial m} = \overbrace{P(M_m) - P(M_f)}^{\text{directed homophily}} \tag{11}$$

$$\frac{\partial D^M}{\partial M} = (1-m)\overbrace{\frac{w(1-w)}{(1-M(1-w))^2}}^{\text{Complements (+)}} + m\overbrace{\frac{-(1-w)}{(M+w(1-M))^2}}^{\text{Substitutes (–)}}. \tag{12}$$

Intuitively, equation (11) shows that when doctors are gender-biased, the demand for male specialists increases in the fraction of male doctors. This demand also depends on the gender mix of specialists. Male specialists are substitutes if most doctors are male and complements if most doctors are female, as seen by observing the terms of (12). This relationship is illustrated in Figure 4, which depicts the average demand for a male specialist as a function of the gender compositions of doctors, for the current gender bias and for a typical specialist gender mix. As females are still the minority of both doctors and specialists in most markets (corresponding to the darker area of the surface), they suffer from a lower demand due to both these effects: fewer doctors favor them, and male substitutes are easily found. Figure A.2 depicts the relationship for different values of M and m.

5.2 Testing for Homophily's Impact: Empirical Strategy

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 link between such bias and the gender gap in 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

Figure 4: Demand for Specialists and the Gender Mix of Doctors

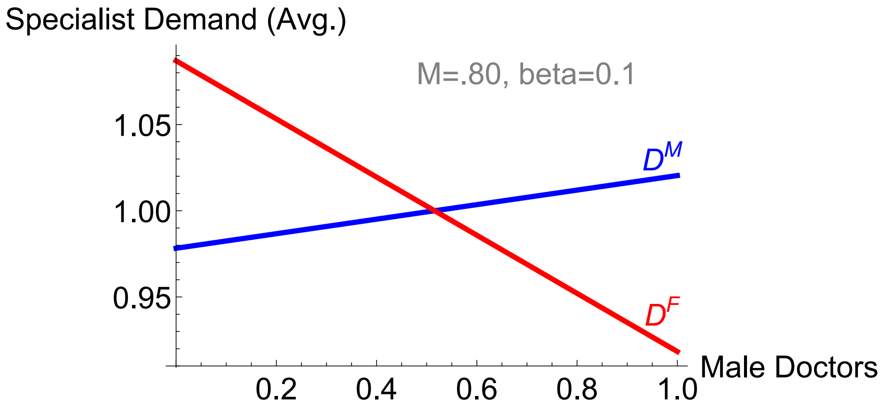

Notes: Counterfactual demand levels, by specialist gender, as a function of m, the fraction of male doctors in the market, for the estimated gender bias $\beta = 0.1$ in a market with $M = 80\%$ male specialists.

handle most outgoing referrals, and all other medical specialists. I estimate:

$$log(Pay_{k,t}) = (\beta_M \mathbb{1}_{g_k=M} + \beta_F \mathbb{1}_{g_k=F}) m_{c(k,t),t} + \gamma_t + \alpha_k + \varepsilon_{k,t} \tag{13}$$

for all non-primary-care specialists, denoted by k, and months, denoted by t. The dependent variable $Pay_{k,t}$ 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.[21]

I estimate (13) using a monthly panel of individual-physician pay for the period 2008–2012. That is, I calculate for each market and month the fraction of primary-care claims

[21] To control for patient homophily, a term $(\delta_M \mathbb{1}_{g_k=M} + \delta_F \mathbb{1}_{g_k=F})\mu_{c(k,t),t}$ is included, where μ is the 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.

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.

5.3 Results: Homophily's Impact on Specialists' Earnings by Gender

When more referrals are handled by male primary-care physicians, demand for male specialists increases, while demand for female specialists decreases. 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 hardly change when 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.

Table 8: Male Fraction of Primary Care and Specialist Workload

	(1) log(Monthly Pay)	(2) log(Monthly Pay)
Female specialist x pct. male PCP (HRR)	-0.26	-0.27
	(0.054)	(0.054)
Male specialist x pct. male PCP (HRR)	0.49	0.47
	(0.029)	(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 (13) 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 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.

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 patient through word of mouth. Therefore, the overall impact of gender-biased referrals could be greater.

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. Historically, women were significantly underrepresented in medicine, and the question remains as to whether historically, homophily slowed their entry. The absence of women from medicine has since dramatically changed: 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 to the possibility that homophily is an impediment to female entry, both in general and into particular specialties. If true, it would imply an even greater contribution of biased referrals to the gender gap.

One potential caveat is that specialist workload outside Medicare is unobserved here. However, even if women were to perfectly make up for their missing referrals by working more elsewhere, which is doubtful given that earnings gaps have been documented among physicians in other settings, results still imply that female specialists encounter a restricted demand in Medicare solely because of their gender.

The magnitude of the effect of referrals on gender pay disparities is fairly large: considering the counterfactual scenario where females handle exactly half of outgoing primary-care referrals instead of their current share. In such case, the pay gap among non-primary-care specialists would decrease by an estimated $(.50 - .35) \times (0.47 + 0.27) = 11\%$. However, such back-of-the-envelope calculation does not hold constant the overall volume of patients and therefore should be taken as suggestive only. The next section presents more disciplined counterfactuals.

5.4 Counterfactuals

To quantify the impact of biased referrals on the earnings gap, I use the model developed and estimated in Section 4 to calculate the demand for male and female specialists with different levels of bias and for different gender mixes of the doctor and specialist populations. I focus on two sets of counterfactuals. First, I calculate the impact of reducing the bias on the average demand by specialists of each gender, using the current male fraction of doctors and specialists in the United States. Second, I calculate the impact of balancing the gender of doctors, keeping the current bias. Even by conservative estimates, each year female specialists lose thousands of dollars' worth of work to their male colleagues, due to a combination of biased preferences and most referrals being made by men.

Figure 5 captures the expected 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. In this figure, average demand is normalized to one. 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 that demand for male ones. This bias amounts to 14% 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. Note the asymmetry, that is due to the fact females are 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 no-work spells (cf., Table 2).

The impact of a given bias in preferences on the earnings gap depends on the gender composition of doctor and specialist populations. Therefore, a second way to assess the impact of such bias on gender earnings disparity is to hold the bias constant at its estimated level and use the model to predict counterfactual average demand with different fractions of male doctors and specialists. Table 9 shows such counterfactual earnings gaps associated with different fractions of male doctors (m) and specialists (M). At a gender composition similar to the current one, $M = m = 0.8$, the estimated bias of $\hat{\beta} = 0.1$ is associated 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

Figure 5: Counterfactual Workload Gap for Different Levels of Preference-Bias

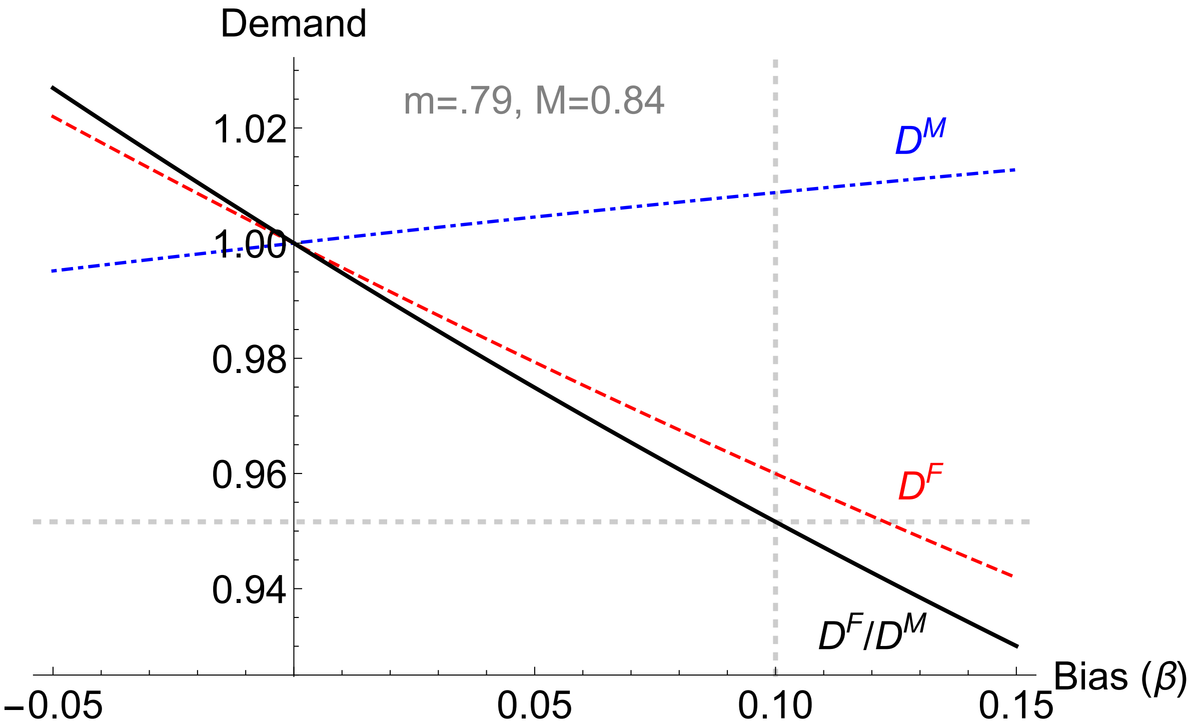

Notes: Average demand for female and male specialists are predicted using the model, given current fractions of male doctors and specialists (m, and M). The thick line gives the fraction of the gap contributed solely due to homophily. Eliminating the bias would reduce the gap by about 5% (in popular terms, restoring 5 cents per dollar to women). See Appendix for the same calculation with different values of m and M.

Table 9: 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. The formula is given below. 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.

only affects the composition of demand, not its level.[22]

The estimated bias in referrals explains a substantial fraction of the current gap in workload between male and female specialists. This gap is not due to difference between the genders, but rather due to differences in the demand specialists of opposite genders face, based on how doctors refer to them. Particularly, while it is conceivable that differences still exist between male and female physician specialists in their preferences for flexibility or in their overall desired workload, such differences fail to explain the documented bias in choices and the dependency of actual workload on the gender of nearby primary care doctors. The totality of evidence therefore suggests that a significant portion of the current gap is due to differences in demand, not just supply.

6 Conclusion

I examine 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 measure homophily by comparing the outgoing referrals rates to male specialists between male and female doctors. 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. To quantify the contribution of doctor's preferences to the observed homophily, net of potential sorting, I estimate a discrete-choice model where doctors choose specialists from heterogeneous, local pools. I use the same model to quantify

[22] Table 9 is based on the following calculation:

$$Gap(m, M; \beta) = \frac{D^F}{D^M} \approx D^F - D^M = \frac{1}{1-M}(mP(F_m) + (1-m)P(F_f)) - \frac{1}{M}(mP(M_m) + (1-m)P(M_f)). \quad (14)$$

Substituting $P(Gg)$ for $g \in \{m, f\}$ and $G \in \{M, F\}$, given in (6), and simplifying yields:

$$Gap(m, M; \beta) = -\frac{(1-w)((M-m) + w(1-(M+m)))}{(1-M(1-w))(w+M(1-w))} \quad (15)$$

where $w = e^{-\beta}$. Note that this assumes no additional systematic bias against female specialists on behalf of doctors of both genders, a bias that is not separately identified from differences in labor supply. However, even if such bias exists, it is still true that increasing the fraction of female doctors would improve the lot of female specialists relative to their male counterparts.

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. I find that homophily is predominantly driven by gender-biased doctors' choices, not sorting. Because most referring doctors are currently men, biased referrals generate demand-driven gender differences in workload, which contribute substantially to the gender earnings gap in Medicare payments. This channel is separate from known supply-side channels that explain other parts of the gender disparity in earnings.

The empirical evidence suggests that a positive externality is associated with increased female participation in medicine, and perhaps in other contexts where networking is important. More generally, it suggests that homophily—a widely documented tendency in many social networks—can bias professional networks and provide a key to the persistence of gender inequality.

Appendices

A Proofs

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

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

The probability $P(M_f)$ is similarly derived. For (8):

$$\begin{aligned} P(M_m) &= \frac{\sum_{k:g_k=M} e^{\beta \mathbb{1}_{g_j=g_k} + \delta X_k}}{\sum_k e^{\beta \mathbb{1}_{g_j=g_k} + \delta X_k}} = \frac{\sum_{k:g_k=M} e^{\beta + \delta X_k}}{\sum_{k:g_k=M} e^{\beta + \delta X_k} + \sum_{k:g_k\neq M} e^{\delta X_k}} \\ &\xrightarrow{P} \frac{M\eta_M e^{\beta}}{M\eta_M e^{\beta} + (1-M)\eta_F} = \frac{Me^{\beta}}{Me^{\beta} + \eta(1-M)} \end{aligned}$$

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

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:

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

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. □

Alternatively, for the more general definition of sorting, based on covariance at the doctor, rather than market, level—$\mathrm{Cov}[m_j, M^j] > 0$, the proof follows immediately from Proposition 1: with unbiased preferences $P(M_m) = E[M^j|g_j = m] > M$, by $\mathrm{Cov}[m_j, M^j] > 0$. Q.E.D. Note that this definition is indeed more general, as by covariance decomposition, $\mathrm{Cov}[m_j, M^j] = \mathrm{Cov}[m^c, M^c]$ under separate markets with common $K_j = K^c$ in each.

Proof. (Proposition 3)

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

□

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):

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

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:

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

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:

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

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. □

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

$$P(M_m) - P(M_f) = \mathrm{E}\Big[\frac{m^c}{m}P(M|m,c) - \frac{1-m^c}{1-m}P(M|f,c)\Big] + \frac{1}{m(1-m)}\mathrm{Cov}[m^c, M^c] \quad (16)$$

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

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

Hence

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

rearranging yields the result. □

B Homophily Dynamics

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:

$$P(Y_{jk,t+1} = 1 | Y_{jk,t} = 1, g, X) = \frac{e^{\eta_{jkt}}}{1 + e^{\eta_{jk't}}} \tag{17}$$

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 $\eta_{jkt} := \alpha_j + \beta \mathbb{1}_{g_j = g_k} + \delta X_{jkt}$. That is, (17) 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 A.2 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.

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*.

Altonji, Joseph G and Rebecca M Blank (1999), "Race and gender in the labor market." *Handbook of labor economics*, 3, 3143–3259.

Antoninis, Manos (2006), "The wage effects from the use of personal contacts as hiring channels." *Journal of Economic Behavior & Organization*, 59, 133–146.

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, 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, 127–151.

Azmat, Ghazala and Rosa Ferrer (2016), "Gender gaps in performance: Evidence from young lawyers." *Journal of Political Economy*.

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.

Bagues, Manuel F and Berta Esteve-Volart (2010), "Can gender parity break the glass ceiling? evidence from a repeated randomized experiment." *The Review of Economic Studies*, 77, 1301–1328.

Baker, Laurence C (1996), "Differences in earnings between male and female physicians." *New England Journal of Medicine*, 334, 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, 506–512.

Barnett, Michael L, Bruce E Landon, A James O'Malley, Nancy L Keating, and Nicholas A Christakis (2011), "Mapping physician networks with self-reported and administrative data." *Health services research*, 46, 1592–1609.

Bertrand, Marianne (2011), "New perspectives on gender." *Handbook of labor economics*, 4, 1543–1590.

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.

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, 991–1013.

Blau, Francine D and Lawrence M Kahn (2000), "Gender differences in pay." *The Journal of Economic Perspectives*, 75–99.

Blau, Francine D and Lawrence M Kahn (2006), "The us gender pay gap in the 1990s: Slowing convergence." *Industrial & Labor Relations Review*, 60, 45–66.

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, 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, 1754–1786.

Burks, Stephen V, Bo Cowgill, Mitchell Hoffman, and Michael Gene Housman (2014), "The value of hiring through employee referrals." *Qaurterly 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, 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, 1861–1862.

Coleman, James (1958), "Relational analysis: the study of social organizations with survey methods." *Human organization*, 17, 28–36.

Currarini, Sergio, Matthew O Jackson, and Paolo Pin (2009), "An economic model of friendship: Homophily, minorities, and segregation." *Econometrica*, 77, 1003–1045.

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*.

Eckstein, Zvi and Osnat Lifshitz (2011), "Dynamic female labor supply." *Econometrica*, 79, 1675–1726.

Esteves-Sorenson, Constança, Jason Snyder, et al. (2012), "The gender earnings gap for physicians and its increase over time." *Economics Letters*, 116, 37–41.

Ghani, Ejaz, William R Kerr, and Christopher Stanton (2014), "Diasporas and outsourcing: evidence from odesk and india." *Management Science*, 60, 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, 20–35.

Goldin, Claudia (2013), “A pollution theory of discrimination: Male and female differences in occupations and earnings.” In *Human Capital in History: The American Record*, 313–348, University of Chicago Press.

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, 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.” Technical report, 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, 659–695.

Himelboim, 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, 40–60.

Ho, Kate and Ariel Pakes (2014), “Physician payment reform and hospital referrals.” *The American Economic Review*, 104, 200–205.

Hochberg, Yael V, Alexander Ljungqvist, and Yang Lu (2007), “Whom you know matters: Venture capital networks and investment performance.” *The Journal of Finance*, 62, 251–301.

Johnson, Erin M (2011), “Ability, learning and the career path of cardiac specialists.” *Cambridge, Massachusetts Institute of Technology.*

Kossinets, Gueorgi and Duncan J Watts (2006), “Empirical analysis of an evolving social network.” *Science*, 311, 88–90.

Kugler, Adriana D (2003), “Employee referrals and efficiency wages.” *Labour economics*, 10, 531–556.

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, 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.

Mulligan, Casey B and Yona Rubinstein (2008), "Selection, investment, and women's relative wages over time." *The Quarterly Journal of Economics*, 123, 1061–1110.

Neumark, David, Roy J Bank, and Kyle D Van Nort (1996), "Sex discrimination in restaurant hiring: An audit study." *The Quarterly Journal of Economics*, 111, 915–941.

Niederle, Muriel and Lise Vesterlund (2007), "Do women shy away from competition? do men compete too much?" *The Quarterly Journal of Economics*, 122, 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, 1132–1157.

Reyes, Jessica Wolpaw (2006), "Do female physicians capture their scarcity value? the case of ob/gyns." Technical report, NBER Working 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, 1748–1750.

Thelwall, Mike (2009), "Homophily in myspace." *Journal of the American Society for Information Science and Technology*, 60, 219–231.

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, 557–566.

Zimmerman, Seth D (2014), "The returns to college admission for academically marginal students." *Journal of Labor Economics*, 32, 711–754.

Figure A.1: Male Fraction of Physicians in Common Medical Specialties

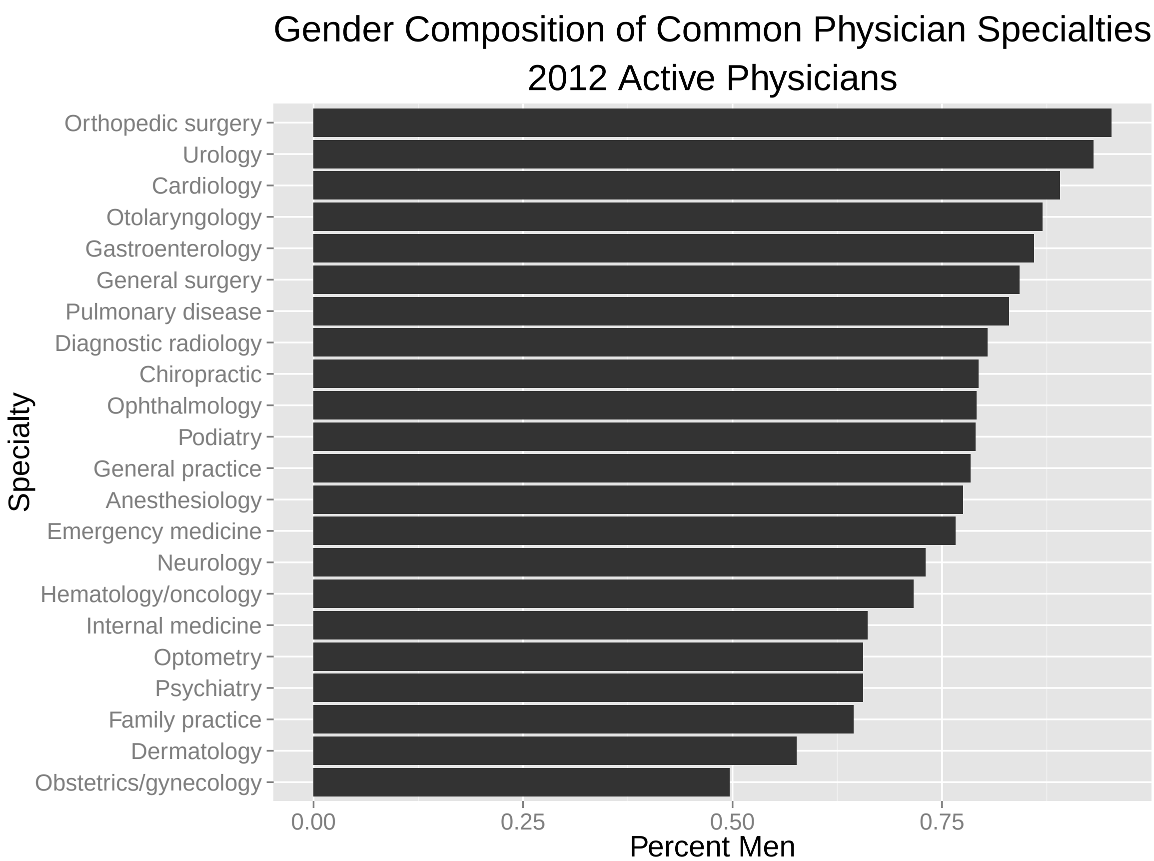

Notes: Percent of active physicians (with any claims) who are male, for the most common specialties by overall number of physicians. Columns are sorted so specialties with the greatest male shares are at the top.

Figure A.2: Average Specialist Demand with Gender-Biased Preferences

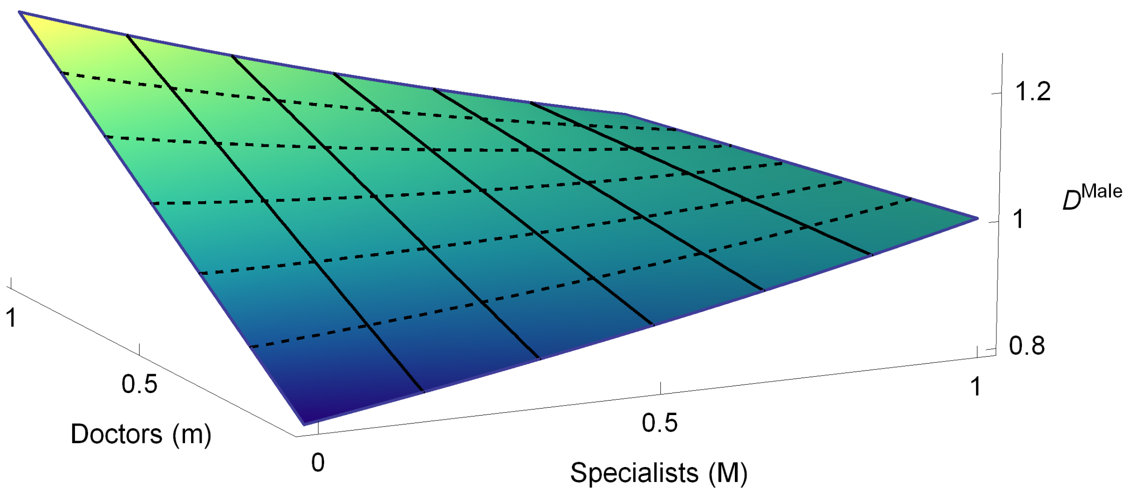

(a) Demand for Male Specialists over the Fractions of Male Doctors and Male Specialists

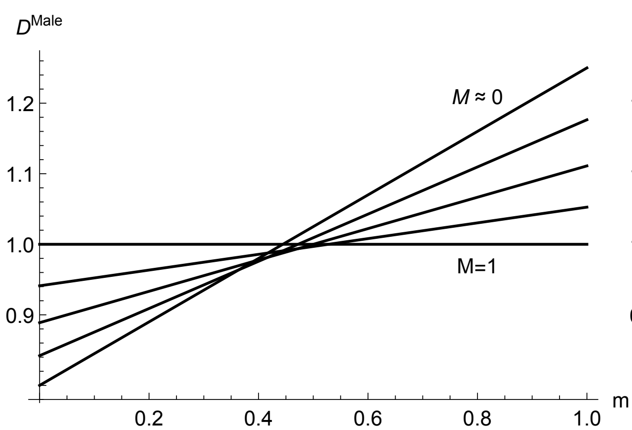

(b) Demand for Male Specialists and Male Doctors

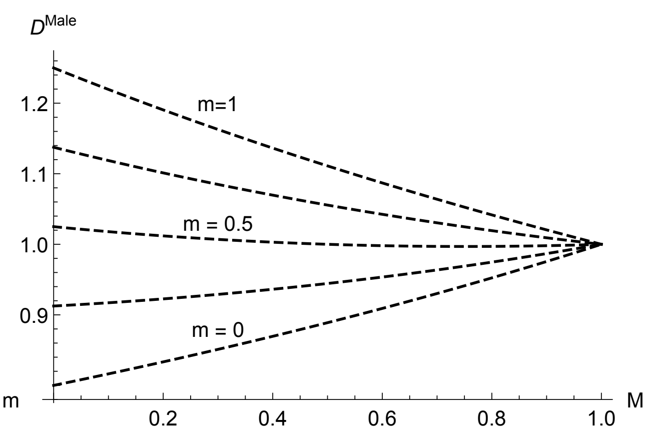

(c) Demand for Male Specialists and Male Specialists

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.

Table A.1: Estimates of Directed Homophily Using Disaggregated Data

	Dep. Var.: Male Specialist				
	(1)	(2)	(3)	(4)	(5)
Male Doctor	0.045***	0.039***	0.038***	0.035***	0.042***
	(55.8)	(52.8)	(50.4)	(46.9)	(51.5)
Male Patient				0.021***	0.037***
				(81.8)	(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 A.2: 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.

Supplementary Appendices

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 9, 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 $\beta = 0$:

$$Gap = D^F - D^M \approx (2m - 1)\beta + O(\beta^2) \tag{18}$$

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 verse.

In fact, the gap mostly depends on the gender distribution upstream even for relatively high levels of bias (Figure A.8). However, for extremely high levels of gender bias, both upstream and downstream majorities matter:

$$\lim_{\beta \to \infty} Gap = \frac{m - M}{M(1 - M)} \tag{19}$$

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.[23]

[23] I thank Alexander Frankel for bringing this case to my attention.

Figure A.3: Homophily and Link Direction

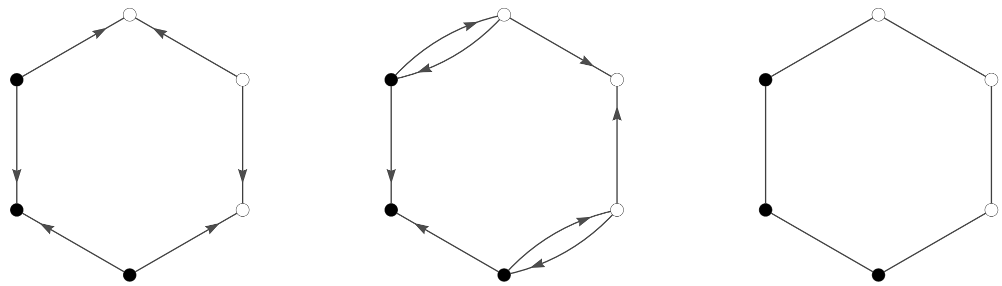

(a) Homophily (b) No Homophily (c) Undirected

Link direction is important in defining homophily: the network (a) exhibits homophily while (b) does not, a difference concealed in their undirected counterpart (c).

Table A.3: Homophily Estimates for Different Age Groups

	Percent of Referrals to Male Specialists		
	Young	Old	All
Male Doctor	0.038	0.044	0.040
	(0.0011)	(0.0015)	(0.00090)
Male Patients (pct)	0.028	0.031	0.029
	(0.0024)	(0.0026)	(0.0018)
Constant	0.79	0.81	0.80
	(0.0078)	(0.0040)	(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 A.4: Homophily and Market Size

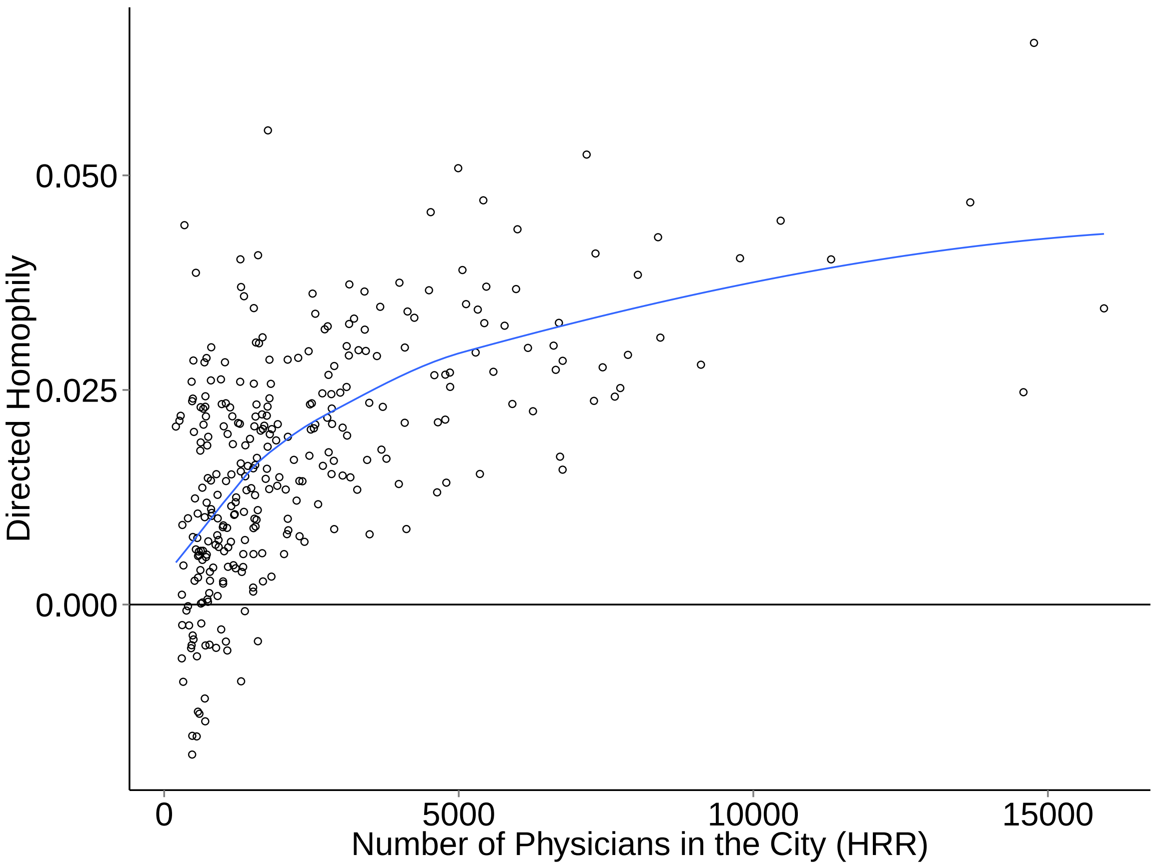

Notes: Homophily estimates of (2), estimated separately for each local physician market (Dartmouth Hospital Referral Region), over the overall number of physicians in the market (men and women). The line is local regression (LOESS) fit. Beyond the mechanical effect of a reduction in variance of the estimates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

Figure A.5: Conditional-Logit Estimates of Gender Bias, by Specialty

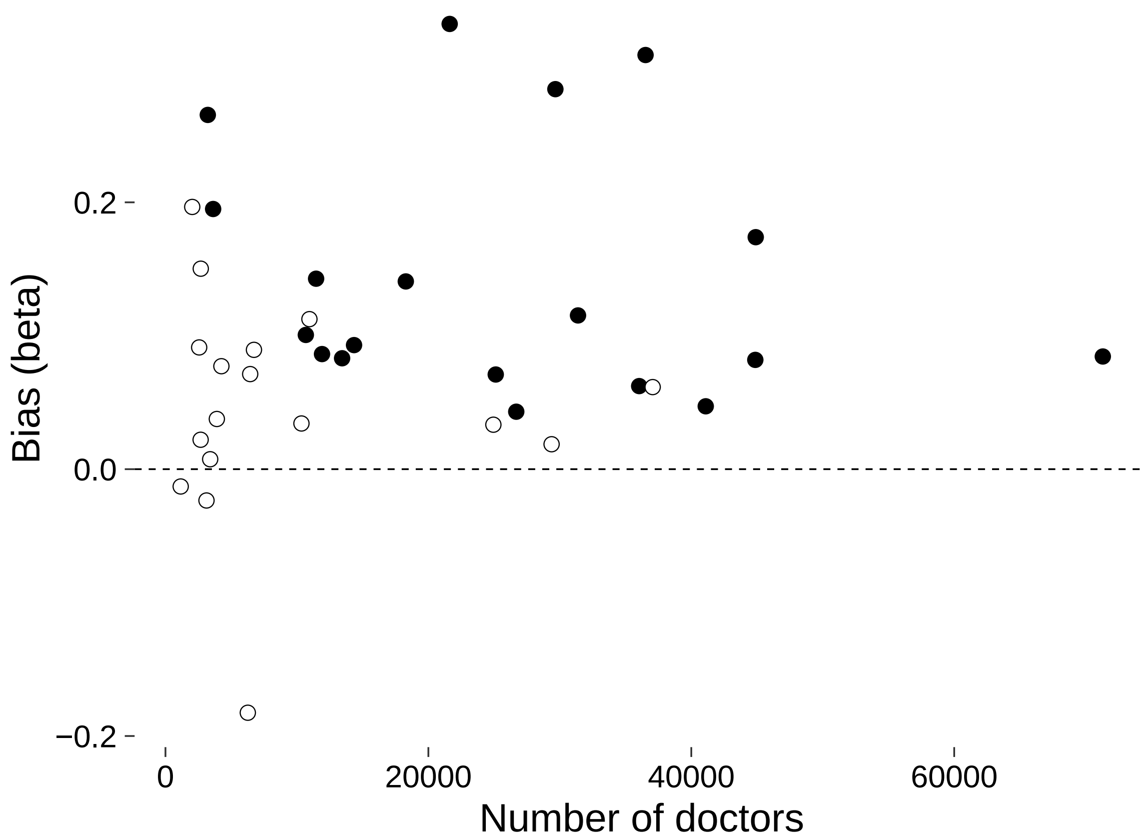

Notes: Estimates of β, the gender bias, from equation (5) with the sample in Table 6, separately for each medical specialty of the receiving physician. Black circles denote estimates that are significantly different than zero ($p < 0.05$).

Figure A.6: Average Number of Referral Relationships by Medical Specialty

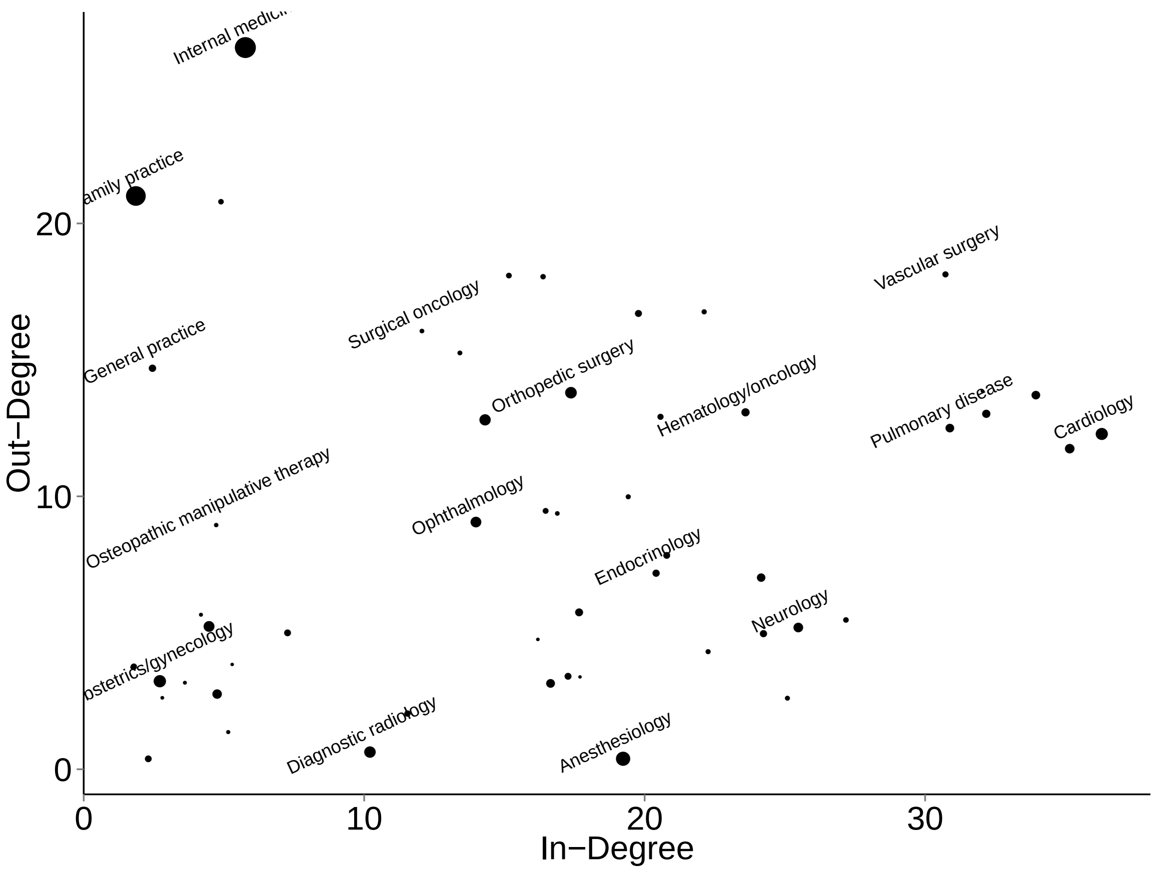

Notes: Degree-heterogeneity is to be expected because doctors in different specialties play different roles in routing patients: some mostly diagnose and refer out, others mostly receive referrals and treat. The figure shows degree distribution by specialty for 2012 referrals: Out-degree is the average 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 A.6 for the data.

Table A.4: Homophily Estimates with Weighted Links

	Percent Referrals to Male Specialists, by:			
	(1)	(2)	(3)	(4)
	Links	Patients	Claims	Dollars
Male Doctor	0.038	0.040	0.040	0.040
	(43.2)	(44.8)	(42.7)	(41.4)
Percent Male Patients	0.029	0.029	0.029	0.029
	(16.6)	(16.5)	(16.1)	(15.4)
Cons.	0.80	0.80	0.80	0.81
	(262.2)	(254.3)	(243.8)	(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.

Figure A.7: Overall Demand as a Function of Gender ($m = M$)

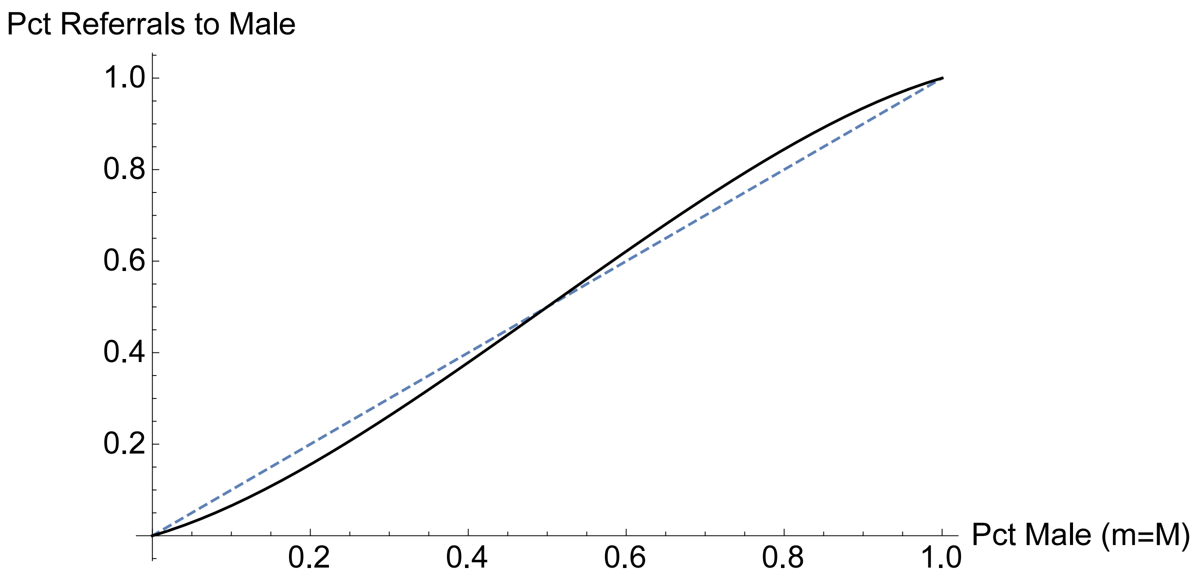

Notes: The figure plots the overall share of referrals going to males, as a function of the male population fractions, in the special case where $m = M$, for gender-biased preferences ($\beta > 0$). The majority gender receives more than its share; the minority gender receives less than its share.

Table A.5: Conditional-Logit Estimates: Referral Probability, with Interaction Terms

	Doctor Referred to Specialist			
Doctor and Specialist:	(1)	(2)	(3)	(4)
Same Gender	0.0841	0.0662	0.104	0.0758
	(35.81)	(16.49)	(29.65)	(12.88)
Male Specialist	0.175	0.175	0.165	0.164
	(73.70)	(73.26)	(46.03)	(45.69)
Same Hospital	3.114	3.072	2.941	2.887
	(720.15)	(579.18)	(540.97)	(414.00)
Same Hospital x Same Gender		0.0598		0.0770
		(13.65)		(12.49)
Same Group	1.346	1.372	1.320	1.344
	(178.27)	(151.85)	(135.26)	(111.58)
Same Group x Same Gender		-0.0386		-0.0354
		(-5.24)		(-3.43)
Same Zipcode	1.074	1.065	1.054	1.047
	(219.53)	(163.68)	(164.04)	(118.78)
Same Zipcode x Same Gender		0.0130		0.0104
		(2.08)		(1.20)
Similar Experience	0.128	0.120	0.131	0.123
	(131.66)	(75.64)	(92.95)	(52.37)
Similar Experience x Same Gender		0.0117		0.0110
		(6.28)		(3.99)
Same Med. School			0.209	0.206
			(49.96)	(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 $\times$). See Table 7 notes for variable definitions and details.

Table A.6: 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.

Figure A.8: The Gender Earnings Gap With Different Levels of Bias

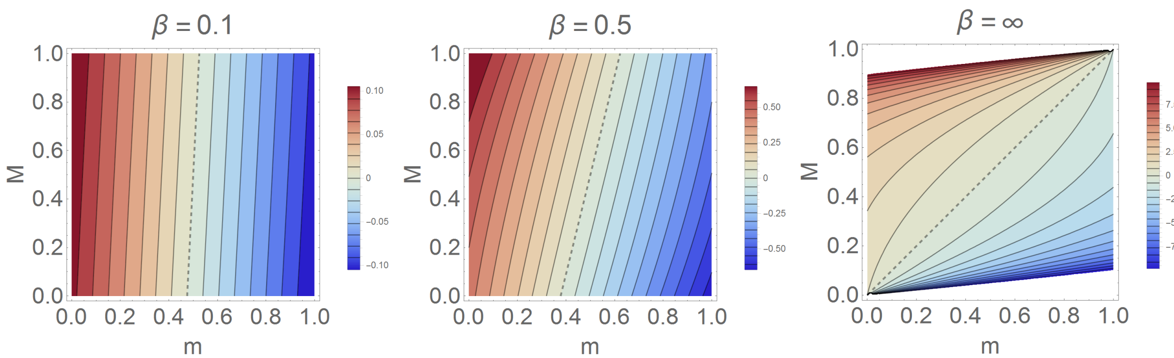

(a) (b) (c)

Colored contour plots of the gender earnings gap, $D^F - D^M$ (Equation 15) with different levels of bias β, for different fractions of males upstream m and downstream M. Blue (right) and red (left) darker shades reflect higher 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.