Do Job Networks Disadvantage Women?

Evidence from a Recruitment Experiment in Malawi [*]

Lori Beaman[†], Niall Keleher [‡], and Jeremy Magruder[§]

February 2013

Abstract

This paper uses a field experiment in Malawi to show that highly skilled women are systematically disadvantaged by referral-based hiring, highlighting another channel behind gender disparities in the labor market. This happens both because most men recommend other men, and because women refer fewer qualified candidates. This result cannot be explained by gender-segregation in networks. We develop a theoretical model of referral choice and exploit random variation in referral contract terms to find that these biases result from social incentives rather than performance expectations. We also document that the screening potential of networks is maximized when men refer men.

1 Introduction

While the gender gap in labor force participation has declined sharply in the last 30 years, women continue to earn less than men in countries around the world (World Bank Group and others, 2011). In Malawi, women are significantly under-represented in the formal sector (World Bank Group and others, 2010) as is common in many developing countries (Bell and Reich, 1988). A large portion of the literature in economics has focused on labor market discrimination (taste-based or statistical) or differences in human capital accumulation as reasons for the

[†]Northwestern University. Email: l-beaman@northwestern.edu
[‡]Innovations for Poverty Action. Email: nkeleher@poverty-action.org
[§]University of California, Berkeley. Email: jmagruder@berkeley.edu

gender gap in earnings (Altonji and Blank, 1999).[1] Another possibility is that hiring processes themselves disadvantage women. We conduct a field experiment recruiting employees for a job in which men and women regularly compete in order to ask whether the use of referrals inherently disadvantages women in the labor market.

A large fraction of jobs - up to 50% - are attained through informal channels, including employee referrals (Bewley, 1999; Ioannides and Loury, 2004). While there is relatively little empirical evidence on the distributional consequences of referral systems, the potential of social network-based hiring to create inequality between groups has been described theoretically (Calvo-Armengol and Jackson, 2004).[2] Ioannides and Loury (2004) document stylized facts that women are less likely to report being hired through a referral and that unemployed women are less likely than unemployed men to report using family and friends as a means of search[3].

Of course, these stylized facts alone do not show that women are disadvantaged by the use of networks in the labor market: women may work in occupations where networks are less relevant, or they may be less likely to report network help for the same hiring procedure. Moreover, if individuals are able and willing to screen on hard-to-observe dimensions for their employers (Montgomery, 1991; Beaman and Magruder, 2012), then referral networks may be advantageous for disadvantaged groups including women. Female applicants could have, on average, weaker easy-to-observe characteristics - like job experience - but network screening

---

[1]Additional explanations include the role of technology (Goldin and Katz, 2002), deregulation and globalization (Black and Strahan, 2001; Black and Brainerd, 2004), and differences in psychological attributes and preferences such as risk preferences, attitudes towards competition, other-regarding preferences, and negotiation (Niederle and Vesterlund, 2007; Bertrand, 2011).

[2]For the Calvo-Armengol and Jackson (2004) mechanism to be a relevant source of long-run inequality between men and women, job networks would need to be characterized by gender homophily. A large literature in Sociology (reviewed in McPherson, Smith-Lovin, and Cook, 2001) suggests that gender homophily in networks begins at early ages and is particularly strong in workforce networks. Mortensen and Vishwanath (1994) also show theoretically that network-based job information dissemination can disadvantage women, even if men and women are are equally productive but men have a higher contact probability.

[3]Moreover, occupational segregation is commonly cited as a source of income disparity across gender (Blau and Kahn, 2000; Arbache, Kolev, and Filipiak, 2010). The use of employee referrals may be one of the mechanisms creating this segregation (Fernandez and Sosa, 2005; Tassier and Menczer, 2008).

may succeed in identifying the women who have strong hard-to-observe but productive characteristics. We may also anticipate that informal information flows are particularly important for reaching women who are less likely to be employed in the formal sector. Therefore, whether women are made worse off by the use of employee referrals remains an open question.

We used a competitive recruitment drive conducted by a research organization in Malawi, Innovations for Poverty Action (IPA-Malawi), as an opportunity to generate a database of male and female applicants[4] to test how job referrals affect the recruitment of men and women in an experimental setting. IPA-Malawi has historically struggled to hire female enumerators and was interested in exploring whether referrals could uncover an otherwise untapped pool of qualified female applicants specifically and qualified applicants in general[5]. The position was advertised using the traditional method of posting flyers. Initial applicants attended a half-day interview process which included a written exam and a mock interview, where the candidate surveyed an actor playing the role of a typical respondent. At the conclusion of the application process, candidates were asked to refer a friend or relative to apply for the position and were offered a finder's fee. The referral process was cross-randomized along two dimensions: candidates were either told that they must refer a woman, that they must refer a man, or that they may refer anyone, and their finder's fee was randomly selected to be a fixed fee of either 1000 or 1500 Malawi Kwacha (MWK; $1=153 MWK) or a performance incentive (a guaranteed 500 MWK with the potential to earn an additional 1300 MWK, for a total of 1800 MWK, if the referral attained a certain threshold). Applicants who perform above the median qualify for future positions, are maintained in an IPA database, and called as positions open.

---

[4]Literally, binders full of men and women.

[5]Often, the gender of the enumerator is important. For example, IPA-Malawi and many other survey firms prefer to use female enumerators when surveying women about sensitive questions, such as family planning practices. Therefore, IPA wanted to recruit both men and women, and historically had found that qualified women were particularly difficult to attract. Informal interviews with qualified female applicants suggest that one reason qualified female applicants were hard to find was that there are gender differences in willingness to travel regularly and for several weeks at a time in Malawi, which is necessary to work as a survey enumerator.

We find that qualified female candidates are strongly disadvantaged by the use of social networks in the hiring process. Among the conventional applicants (CAs) who were allowed to choose either gender for a referral, only 30% of referrals are women. This is significantly lower than the fraction of women who apply through traditional recruitment channels (38%). The low number of women referred is driven largely by male candidates. When given the choice, men systematically refer men: 77% of men's referrals are other men. Women refer women at approximately the same rate at which they apply through the traditional recruitment method. However, women systematically refer people who are less likely to qualify: a female candidate is nearly 20 percentage points less likely to refer someone who qualifies than a male candidate. These two effects combine to create a scenario where very few people ultimately refer qualified women when given a choice over which gender to refer: only 13% of men and 17% of women refer qualified women, which compares to 43% of men and 22% of women who refer qualified men. This disadvantage for women, however, does not appear due to men (or women) being unconnected to women that they deem suitable: men and women make referrals at identical rates when required to refer women as when they are required to refer men under all contracts.

Because men and women are capable of identifying suitable female referrals, this disadvantage for women appears due to some aspect of the referral choice problem. We propose a simple model of who individuals choose to refer under different types of referral contracts. The model provides a guide to interpreting our experimental variation and suggests empirical tests to provide evidence on the underlying reasons women are disadvantaged. In the model, individuals receive a social payment from referring a particular network member, in addition to any payment provided by the firm. The social payment may capture altruism, an actual financial transfer, or reduced future transfers if the two are in a repeated risk sharing arrangement (Beaman and Magruder, 2012). They also receive a noisy signal of each network member's

4

quality. We allow there to be several key gender differences in this decision process, which could generate the main empirical finding that qualified women receive different opportunities from their networks. For any type of conventional applicant (CA), networks of men and women may differ in: (i) the magnitude of social payments; (ii) the quality of friends who provide the highest social payments; (iii) the extent of the tradeoffs between social payments and observable quality (this captures for example if high quality women are scarce relative to high quality men in some networks); and (iv) the accuracy of the signal of observed quality relative to actual quality.

Because of random variation in the structure of the finder's fee, we are able to observe the characteristics of two nodes in CAs' networks: the person chosen when CAs are motivated in their choice only by network incentives, and the person selected when firms provide an additional incentive to find a person who is high ability. This facet facilitates several tests of the various sources of heterogeneity. First, the model suggests that referrals selected under fixed fee payments should reflect the people that give the highest social payments (net of recruitment costs) to CAs. From this framework, we identify that men systematically receive the highest social payments from other men. We also identify that both the men and women who are "preferred" by men in this way are of similar ability, which is slightly below the average CA's ability. Women are not systematically closer to one gender over the other, but women CAs receive the largest social payments from individuals who are significantly less likely to qualify than the average CA[6]. Overall, social incentives among both men and women's networks make it harder for qualified women to get job opportunities[7].

Other network characteristics may lead to disadvantages for qualified women if firms

---

[6]This result comes from women who are allowed to refer either gender. From restricted-gender treatments, we also observe that the men who are closest to women outperform the women who are closest to other women.

[7]We discuss the possibility that CAs' referral choices are affected by the fact that CAs are competing with their referrals for the job in section 7.1, and the possibility that these differences could be suggested by (biased or unbiased) expectations of ability or reputational concerns in section 8.

add additional incentives for workers to identify highly-skilled referrals[8]. Given our results on social payments and the quality of people referred under fixed fees, we derive that steeper tradeoffs or worse information could also put women at a disadvantage, and that either of these characteristics would manifest in a smaller change in referral qualification rates induced by performance pay. We note that this is additionally an important criterion for employers who use networks to screen: any gender disparities in performance premia will have direct implications for employer incentives to encourage or discourage women's disadvantage.

We find that men exhibit a large performance premium when referring men, but no performance premium when referring women: this factor allows us to conclude that men's networks of women have either worse information or greater social tradeoffs than men's networks of men. These two concepts are highly related, which the model highlights, and accordingly difficult to disentangle. Nevertheless, several pieces of evidence point towards men having poor information about women. Women CAs in performance pay treatments are not more likely to make referrals who qualify when referring men or women, but performance incentives do change women's referrals of both men and women in some productive dimensions.

This paper suggests that an additional factor leading to the observed gender gap in women's earnings is the way employees are recruited. As we discuss below, this experiment suggests that gender differences in network characteristics create several profitable motivations for firms to permit this disadvantage, even in the absence of any intended (taste-based or statistical) discrimination on the part of the firm. The results also highlight that the screening potential of networks must be bought: firms must be willing to generate incentives to refer only the best people within their employees' networks in order to find high quality referrals, as in Beaman and Magruder (2012).

---

[8]In other contexts, firms incentive workers both through actual cash incentives for making a high-performing referral and through reputational considerations.

6

The paper is organized as follows. The experimental design and data are described in section 2. The main results are discussed in section 3. The theoretical framework which highlights potential mechanisms for why women are disadvantaged by referral hires is elaborated in sections 4 and 5. The empirical evidence for each of the three hypotheses are presented in sections 6 and 7, followed by discussion of an alternative explanation in 8 and then the conclusion.

2 Experimental Design

2.1 Setting and Overview

Women in Africa are more likely to be in the informal sector and the proportion of women with formal employment is less than half that of men. Malawi is not an exception to this trend. A recent survey of Malawian households suggests that less than one-third of women participate in the formal labor force, while nearly 58% of men do so (World Bank Group and others, 2010). Among urban women, 38.2% had not been employed in the preceding twelve months; this rate is more than double that found among urban men (18.6%) (National Statistics Office (NSO) and ICF Macro, 2011).

IPA-Malawi hires enumerators to conduct interviews of farmers, business owners, and households in rural and urban Malawi. In the 12 months following the recruitment drive (our experiment), IPA-Malawi projected hiring a minimum of 200 enumerators for its survey activities. IPA-Malawi had an explicit motivation to hire more female enumerators than their usual recruitment methods allow. Typically, only 15 to 20 percent of enumerators hired by IPA-Malawi are women, and some survey tasks require same-gendered enumerators (for example, same-gendered enumerators are sometimes important for sensitive questionnaires). For

this experiment, we introduced incentives for job applicants to make referrals during IPA's recruitment sessions in the two main Malawian cities, Blantyre and Lilongwe.

In this experiment, IPA posted fliers indicating a hiring drive at a number of visible places in urban areas. The posters included information on the minimum requirements for IPA enumerators, the dates and times of the recruitment sessions, and a solicitation to bring a CV and certificate of secondary school completion (MSCE). Minimum requirements to be hired for an enumerator position are: a secondary certificate, fluency in the local language (Chichewa), and English reading and oral comprehension. Candidates with data collection experience, good math skills, and basic computer skills are given preferential review. Participants then attended an interview session, where they submitted their CV and were registered with a unique applicant number. Participants were limited to those individuals who had never worked for IPA. Each day, two sessions were conducted by IPA staff. At the start of each session, participants were introduced to IPA and the role of an enumerator was described.

2.2 Quality Assessment

The screening session included a written test similar to the one IPA had previously used, and a practical test which served as a condensed version of the training that IPA had previously used to select enumerators.[9] Participants were given one of two distinct written tests[10]. Each test consisted of several math problems, ravens matrices, English skills assessment, job comprehension component, and a computer skills assessment. Our screening session integrated a practical test to obtain information on otherwise hard-to-observe qualities that are important for the work of an enumerator.

---

[9]The standard IPA-Malawi screen session includes a written test similar to what was used in the experiment. Instead of the practical test used in the experiment, applicants deemed to be qualified from the written test and CV would be invited for a survey-specific training of enumerators. After a multi-day training for that survey, a subset of the candidates who were trained are typically selected to work on that survey.

[10]The two tests were distributed at random to limit cheating.

For the practical test, the participant played the role of the enumerator for a computer assisted personal interview.[11] An experienced IPA enumerator played a scripted role of the interview respondent. The respondent scripts included implausible or inconsistent answers (i.e. age, household size, household acreage) to survey questions. These false answers were used as checks on the participant's ability to pay attention to detail and verify inaccuracies in responses. When the participant pressed the respondent for a correction, the respondent gave a plausible answer. Among the respondents, two sets of implausible answers were used in order to limit any ability to predict the practical test.

Scores were calculated for all participants on a 0-to-100 scale. The total score was a combination of the CV score, written test score and practical test score.

2.3 Referral Instructions

The setting offered an opportunity to test several potential channels through which a firm can influence the type and quality of applicants generated through a referral process. Prior to leaving the recruitment session, participants had a one-on-one conversation with the recruitment manager. During this conversation, a letter was provided to the applicant inviting the applicant to identify another individual to refer to IPA for consideration as an enumerator. The message provided to the participant was the crux of this experiment: we randomly varied the content of the letters.

Each letter included an instruction about the gender requirement, if any, of the referral who could be invited to attend a future recruitment session. The letter instructed the original participants that their referral had to be male, had to be female, or could be anyone. The referral needed to be someone who had not worked for or been tested by IPA in the past.

---

[11]All participants were required to go through a short self-administered training with a computer-assisted personal interviewing (CAPI) software in order to ensure a consistent level of familiarity with the computer program. Once finished with the self-administered CAPI training, participants moved to the practical test.

The letter also said that the referral should be highly qualified for the enumerator position and given a suggestive guide about what this would entail. Namely, the letters stated that a strong enumerator should have a secondary school certificate, fluency in Chichewa, excellent comprehension of English, data collection experience, and good math and computer skills. The CA was told that the referral should perform strongly on the written and practical assessments completed by the CA.

Conventional applicants were also randomly assigned into one of three pay categories (cross randomized with the gender treatments): a fixed fee of 1000 Malawi Kwacha, a fixed fee of 1500 MWK, or a performance incentive of 500 MWK if their referral does not qualify or 1800 MWK if their referral does qualify. All treatments were fully blind from the perspective of the evaluators. All CAs were eligible to receive payment (fixed fee or base pay, if in the incentive group) if their referral attended and completed a recruitment session. Referrals typically participated in recruitment sessions 3 to 4 days after the conventional applicant's session. The screening session, including the written and practical test components, were the same as for conventional applicants.

Each week, a list of qualified applicants was posted at the recruitment venue, and qualified applicants were told that they would be considered for future job opportunities with IPA-Malawi. Any original applicant who qualified for a payment was informed and given payment in a sealed envelope.[12]

Appendix Table A1 displays summary statistics for the sample of CAs, for men and women separately. It also shows that the randomization led to balance along most characteristics. The p value for the joint test of all the treatment variables, and their interactions,

---

[12]To maintain a quick turn-around in notifying applicants of qualifying, real-time test-scoring and data entry was necessary. This led to a few misentered values which slightly affected the identities of qualifying people. In this paper, we use corrected scores and qualifying dummies which do not reflect these typos in all main analysis, though results are robust to using the actual qualification status.

is displayed in column (2) for male CAs and column (5) for female CAs. Among male CAs, only the number of feedback points for male CAs is significant 5% level (though the Practical Component Z-score is almost significant at the 10% level for both men and women CAs). For women CAs, there is a baseline difference in test scores at the 10% level. This is driven by women CAs who were in the male-fixed fee treatments performing slightly worse on average than other women CAs in either unrestricted or women-only fixed fee treatments.

3   Are Qualified Women Disadvantaged?

Figure 1 plots kernel densities of CA overall test score separately for men and women, and confirms that men and women who respond to the traditional recruitment method on average have similar distributions of test scores. There is some evidence that male CAs outperform female CAs on the assessment, which can be seen in the small rightward shift in men's performance across the distribution of the referral test scores. Panel A of Table 1 confirms that this difference is statistically significant. However, there is much more variation within CA gender than there is between CA genders, and nearly all of the support of men's and women's test scores is common. As such, men and women are in true competition for these jobs. Nonetheless, we may be concerned over whether the distribution of quality of potential referrals is different in networks of men and women. In section 5 we will develop and test a model to evaluate whether there are gender differences in the quality of potential referrals.

Panels A through C of Table 2 document the primary result of this paper. While 38% of applicants themselves were women (and 39% of applicants who could refer either gender, panel B), only 30% of referrals are women when we allow CAs to choose which gender to refer (difference significant at the 5% level). This difference in application rates happens entirely because men systematically do not refer women when given the choice: women refer women at

11

approximately the rate by which women apply themselves (43% of the time), while men refer women only 23% of the time. The difference between male and female CAs is significant at the 1% level, as shown in column (4). Moreover, these difference persist across the range of CA performance: Figure 2 presents local polynomial regressions of the gender choice of referral on CA overall test score, disaggregated by men and women[13]. Across the distribution of potential test scores, CA women are more likely to refer women than CA men, with particularly large differences at the top and bottom of the distribution of CA test scores.

However, qualified women can also be disadvantaged if unqualified people are being referred more than qualified people regardless of any gender preference. Table 1 also shows that there is a large gender difference in the qualification rate of referrals: while men make references who are about as likely to qualify as CAs are on average, women make references who are eighteen percentage points less likely to qualify (38% versus 56%) when given an unrestricted choice of genders. Rows 3 and 4 reveal that these two results together combine to create a scenario where very few CAs - either men or women - refer qualified women, as only 13% of men and 17% of women refer qualified women (in contrast to 43% of men and 22% of women who refer qualified men). Women's references are much less likely to qualify irrespective of the gender of the referral. When women choose to refer men, those men qualify at a 38% rate while the women they choose to refer qualify 39% of the time. Men's male referrals have a 55% qualification rate while men's female references have a 57% qualification rate. In Figure 3, we again verify that this difference persists across the range of CA test scores. In this case, the qualification difference is most notable at the top of the distribution, as male CAs make referrals who are more likely to qualify in a way which increases monotonically with their test scores, while women's referral quality faces an inverted-u shape, so that the most-skilled and

---

[13]In both cases, the sample is restricted to CAs who have the choice of which gender to refer.

least-skilled women make referrals who are similarly unlikely to qualify.

These two differences together put qualified women at a substantial disadvantage: most men seem to respond to an unrestricted referral situation by identifying men, while most women seem to respond to such a situation by referring unqualified people of either gender. This is consistent with the finding from observational data from a call-center in Fernandez and Sosa (2005)[14] and supports the large literature in sociology arguing that informal referral processes are one of the drivers of segregation of jobs (Doeringer and Piore, 1971; Mouw, 2006; Rubineau and Fernandez, 2010). Overall, we conclude that the use of referral systems strongly disadvantages qualified women in this context.

4   Are Men and Women Connected?

One explanation for why men refer so few women is that it may not be a choice: men may simply not be connected to women. Indeed, one proposed cause of gender segregation in the labor market is segregated social networks (Tassier and Menczer, 2008). Based on this explanation, referrals serve to perpetuate job segregation due to the limited overlap of groups from which referrals are drawn.

Our view is that a sensible definition of connectedness would reflect contract terms: clearly, any of our male CAs would be successful at finding a female referral at a sufficiently high price, particularly in fixed fee treatments where the CA need not be concerned with referral quality. For now, suppose simply that each CA i receives a number of draws of friends, who may be male or female. Each friend j is characterized by two characteristics: a social payment α_j (net of costs of recruiting that person) which (s)he will give to i if (s)he is offered the referral, and some probability of qualifying, λ_j, as well as their gender. The social payment is meant to

---

[14]In that context, men are the disadvantaged group, who are similarly less likely to receive referrals.

include ideas like altruism or expected future reciprocity as well as the costs of notifying friend j about the opportunity. Thus, when CA i is offered a contract with fixed component F_i and performance component P_i, if i refers j, then i receives in expectation

$$F_i + \alpha_j + \lambda_j P_i \tag{1}$$

Assuming that CAs do not make referrals if they cannot receive positive payments in expectation suggests a straightforward definition of connectedness.

Definition 1 *CA i is **connected** to gender g at contract (F_i, P_i) if $max_{j \in g} F_i + \alpha_j + \lambda_j P_i > 0$*

Under this definition of connectedness, CAs are unconnected under fixed fees if the largest possible social payment is less than $-F_i$, and they are unconnected under performance pay if referrals share a low α_j and a low probability of qualifying. Clearly, if male CAs are less connected to women at our contracting terms, it could generate the disadvantage that women face in referral systems.

We can analyze this in a straightforward way: define an indicator $R_i = 1$ if the CA makes a referral, and $R_i = 0$ if the CA does not. Because we randomly restricted some CAs to referring only women, and other CAs to referring only men, we can test whether CAs are more or less likely to be connected to women or men at our contracting terms. Moreover, because some CAs were allowed to refer either women or men, we will additionally be able to test whether CAs who are unconnected to men at a particular contract are likely to be connected to women at that contract: if so, it would suggest that CAs are receiving a number of draws of both men and women, so that even if all the draws of an CA's own gender fail to make the participation threshold, there is a strong chance that the other gender exceeds it. As a test,

then, we simply regress

$$R_i = \sum_k \alpha_k T_{ik} + \delta_t + u_i$$

Where T_{ik} is the exogenously assigned treatment in terms of referral gender and contract payment and δ_t are time trends. Table 2 presents this analysis, where restricted male treatments (or male fixed fee treatments in specifications which disaggregate by contract terms) are the excluded group. Overall, neither men nor women are significantly less likely to make a reference when assigned to refer women than when assigned to refer men, and point estimates on any gender differences are small in magnitude. When we disaggregate by treatment, we observe that men are less likely to make a reference when they are given performance pay than when they are given fixed fees, if the gender of their referral is restricted. The mean referral rate under fixed fees for men in restricted treatments is 90%; point estimates suggest that if these men are instead given the performance contract return rates fall to 75%. However, if men are given the choice of referring either men or women, the return rate rises back to 90% - this suggests that there are 15% of men who know only a man who is worth referring under performance pay, but also 15% who know only a woman who is worth referring. For female CAs, there is a similar trend, though the point estimate is smaller and not statistically significant.

We find these results striking in several ways. First, they reject the hypothesis that the trend of men referring men noted in section 3 occurs because of men being unconnected to women. Most male applicants are connected to suitable women, and they are as likely to be connected to women as they are to be connected to men under either contract structure. There are also a sizeable number of men who are only connected to women, when the performance of the referral matters. Second, they suggest that mean returns under performance pay are lower than under fixed fees. CA return rates under fixed fees remain at 90% for both genders of CAs assigned to both levels of fixed fees; this suggests that the expected return to performance

15

pay is lower than 1000 MWK, the lower fixed fee level. Given that the performance pay contract featured a guaranteed fee of 500 MWK and a performance premium of 1300 MWK, this suggests that the person they choose to refer under fixed fees (who they could have also referred under performance pay) has an expected qualification likelihood below 5/13, which we return to below.

Finally, there are no differences in referral rates for CAs of either gender when we restrict them to men versus when we restrict them to women. There are, however, differences in return rates between contract terms for gender-restricted referrals. Given that attrition rates are in any event low relative to most panel studies and uncorrelated with the gender treatments which are the focus of this study, we abstract from them in the main analysis, though when we analyze performance premia we will discuss the potential role of attrition in biasing our estimates.

Given that CAs are connected to both men and women, it seems likely that some other factor leads to the disadvantage women face from referral systems. In the next section, we further develop the model in the interest of identifying key differences between an individual CA's networks of men and women.

5 Model and Mechanisms

Building on equation 1, retain the notation α_j as a social incentive supplied by friend j and now suppose that CA i expects j to score Q_j points on the skills assessment. If j belongs to gender g, suppose that his (her) actual performance is $Y_j = Q_j + \varepsilon_j$, where ε_j is distributed $N\left(0, (\sigma_\varepsilon^g)^2\right)$, and the referral qualifies if $Y_j > 60$. Note that σ_ε^g may be different between men and women. Using the notation from the previous section, this suggests that $\lambda_j = \left(1 - \Phi\left(\frac{60-Q_j}{\sigma_\varepsilon^g}\right)\right)$ where $\Phi(\cdot)$ is the cdf of the standard normal distribution. As before, CA i is given a contract (F_i, P_i) and is restricted to make a referral out of individuals who belong to set $\mathcal{G}$, where $\mathcal{G}$ could be

$\{male\}, \{female\}$ or $\{everyone\}$.

While i knows a number of people in each gender specific network, we focus on the subset of those draws who could be optimal referrals under various contracting conditions. In particular, individual j will only get chosen under some contract (F_i, P_i) if $Q_j \in \arg\max_{k \in \mathcal{G}} Q_k | \alpha_k \geq \alpha_j$. That is, j will only get chosen if his or her observed quality is the best among eligible referrals who offer at least as much in social payments. For each gender g, define $h^g(\alpha_j) = Q_j$ to be the mapping between α_j and Q_j in this set, where $h^g(\alpha_j)$ is decreasing in α_j by the selection rule. Denote $\alpha_1^g = \max_{j \in g} \alpha_j$, where $j \in g$ if j is of gender g. To make analysis tractable, approximate $h^g(\alpha_j) = Q_1^g + \gamma^g(\alpha_1^g - \alpha_j)$. CA i therefore solves

$$\pi_i(\alpha_j, P_i, F_i) = \max_{j \in \mathcal{G}} P_i \left(1 - \Phi\left(\frac{60 - Q_1^g - \gamma^g(\alpha_1^g - \alpha_j)}{\sigma_\varepsilon^g}\right)\right) + F_i + \alpha_j \qquad (2)$$

Gender-specific networks can be heterogeneous, therefore, in 4 different ways: they may differ in $\alpha_1^g, Q_1^g, \gamma^g$, and σ_ε^g. The following set of definitions characterize these differences.

Definition 2 *CA i is **closer** to gender g than to gender g' if $\alpha_1^g > \alpha_1^{g'}$*

Definition 3 *CA i's network of gender g is **higher quality** than his network of gender g' if $Q_1^g > Q_1^{g'}$* [15]

Definition 4 *CA i faces a **shallower** network of gender g' than of gender g if $\gamma^g > \gamma^{g'}$* [16]

Definition 5 *CA i has **better information** about gender g than about gender g' if $\sigma_\varepsilon^g < \sigma_\varepsilon^{g'}$*

---

[15]This is a rather special definition of higher quality and does not indicate that all members of gender g are higher quality than gender g' or even that members of gender g are on average higher quality than members of gender g'.

[16]This definition of shallowness corresponds to higher tradeoffs between social incentives and expected performance. It is motivated by the idea that if a CA knows many people of one gender, the odds that he or she knows someone who would both be a good enumerator and gives reasonably high social payments is relatively high, reducing tradeoffs.

These four types of heterogeneity allow networks of men and women to be different in the degree of social payments possible, in quality of key individuals, in the tradeoff between social payments and quality, and in the usefulness of referral networks for screening. Our interest is to test whether gender differences in these four characteristics can contribute to the observed differences in referral choices. We consider separately optimal behavior under fixed fee contracts of the form $(F_i, 0)$ and performance pay contracts of the form (F_i', P_i) where $P_i > 0$.

5.1 Fixed Fees and Social Incentives

In interpreting our data in the context of the model, we start with a description of what happens when we provide contracts of the form $(F_i, 0)$.

Proposition 1 *Under fixed fee contracts of the form $(F_i, 0)$, CAs always refer the closest person of the eligible gender, friend 1. Characteristics of referrals under fixed fees will therefore be characteristics of the closest person, including gender and quality.*

Proposition 1 shows that a referral recruited under fixed fees, and restricted to a particular gender, is friend 1 of the designated gender. A first implication of this proposition is that the fraction of CAs who refer men when allowed a choice of genders can be interpreted as the fraction of CAs for whom $\alpha_1^m > \alpha_1^f$. If CAs systematically refer men under fixed fee treatments, proposition 1 suggests that this is because men are systematically closer to CAs. A similar logic applies for the trend by which women are making referrals who are unlikely to qualify: if they make these referrals when given fixed fees, then we can conclude that women are closest to potential enumerators with poor skills. Our empirical analysis in section 6 therefore begins by examining fixed fee treatments to test whether gender differences in network quality or in closeness could affect the opportunities available to women.

5.2 Performance Payments

Only one parameter of network heterogeneity, social payments, could lead to women's disadvantage under fixed fees. However, heterogeneity of any of the four parameters could lead to biases in referral behavior under performance pay contracts, or more generally if the employee making the referral perceives a benefit to referring a more highly skilled worker[17]. In particular, the returns to making a referral under performance pay are increasing in α_1^g and Q_1^g, decreasing in γ^g, and ambiguously related to σ_ε^g, depending on whether the optimal referral under performance pay is expected to qualify or not. As a result, if (for example) men's networks of women feature lower social payments, are lower quality, are more shallow, or have different information, a preference in hiring men under performance pay contracts may arise. Fortunately, we can use random referral contract variation to learn more about whether heterogeneity in key parameters other than α_1^g can contribute to women's disadvantage in the case where CAs face an incentive to make highly skilled references.

First, Proposition 1 suggests that differences in quality can be directly estimated through the quality of referrals made under fixed fees. Because quality, defined as Q_1^g, is the quality of the closest member of gender g, we can directly estimate the quality distribution of referrals from male and female CAs who are restricted to refer only men or only women and conclude whether quality differences may lead to differences in referral choices.

To test differences in other parameters, we need to return to the referral choice problem under performance payments. Recall that λ_j^g denoted the probability that referral j was hired. Suppose person p is the optimal referral when performance pay is positive ($P_i > 0$).

Definition 6 *Define the performance premium, $\Psi^g = \lambda_p^g - \lambda_1^g$, as the difference in qualification*

rates between the optimal referral under performance pay, λ_p^g, and the optimal referral under fixed fees, λ_1^g.

A consequence of proposition 1 is that the only two network characteristics related to λ_1^g are quality, Q_1^g, and information, σ_ε^g.

Proposition 2 *For $P_i > 0$: Either the optimal referral remains friend 1, or the CAs choose an optimal network member j whose qualification probability is a monotonically decreasing function of $\sigma_\varepsilon^g/\gamma^g$ and unrelated to other network characteristics.*

Proposition 2 describes the solution to the model: either there is a corner solution, and CAs maximize social incentives and ignore the performance incentives, or they choose a referral who optimizes the basket of those social and performance incentives. Interestingly, only the ratio $\sigma_\varepsilon^g/\gamma^g$ determines the optimal interior solution; this occurs as $\sigma_\varepsilon^g/\gamma^g$ describes the benefit of giving up one additional unit of social incentives, where γ^g describes how much extra Q is bought for that unit of social incentives and σ_ε^g describes the extent to which that extra Q translates into expected performance pay.

While Proposition 2 characterizes the solution to the referral problem, empirically we will be interested in how differences in network characteristics between gender-specific networks contribute to women's disadvantage. The gender g network may have a larger performance premium than gender g' either because the qualification rate of the person chosen under performance pay is higher for gender g or because the qualification rate of the index partner is lower. If a CA's network of one gender has a larger performance premium than another and at least some of that difference is attributable to improvements in qualification rates under performance pay ($\lambda_p^{g'} < \lambda_p^g$), we can make sharp predictions as to the underlying network characteristics which lead to that result. Proposition 3 describes these relationships. Note that the

condition $\lambda_p^{g'} < \lambda_p^g$ can be empirically tested in our data by looking at the average qualification rate of CAs in performance pay treatments assigned to either the male-only or female-only referral treatments.

Proposition 3 *Suppose that the performance premium when referring gender g is larger than the performance premium when referring gender g' $\left(\Psi^g > \Psi^{g'}\right)$. If $\lambda_p^{g'} < \lambda_p^g$, then at least one of the following must hold[18]:*

i): The network of gender g' is more shallow than the network of gender g $\left(\gamma^g > \gamma^{g'}\right)$

ii): The network of gender g has better information than the network of gender g' $\left(\sigma_\varepsilon^g < \sigma_\varepsilon^{g'}\right)$

iii): The network of gender g is higher quality than the network of gender g' $\left(Q_1^g > Q_1^{g'}\right)$

Proposition 3 makes clear that there are only three candidate explanations for a greater performance premium when referring gender g compared to when referring g' if at least some of the difference is attributable to improvements in λ_p^g relative to $\lambda_p^{g'}$: higher quality (which is directly estimable from fixed fee referrals), better information, or a less shallow network for gender g. Given proposition 2, it is clear that higher quality cannot cause a difference in the quality of a non-index partner but could create a larger Ψ^g by generating an internal solution (i.e. the selection of a referral other than the highest social-payment network member). In contrast, propositions 2 and 3 highlight symmetric roles of information and shallowness in the choice of referral ability which impact referral choice similarly. Intuitively, this makes sense: when individuals have high information but shallow networks, they have networks where

---

[18]Otherwise, the larger performance premium is driven exclusively by differences in the qualification rate of the index partner rather than the response of the CA to performance pay. λ_1^g is increasing in quality (Q_1^g) and ambiguously related to information (σ_ε^g), which leads to some parts of the parameter space where a larger performance premium in gender g can be attributed to lower expected quality of the index partner or different information levels. The appendix shows that in these cases, the index partner is always suboptimal under performance pay for both genders. The appendix provides a more thorough characterization of the relationship between quality and information throughout the parameter space, but since empirically it is the case in our data that larger performance premia are associated with a larger λ_p, we focus on this case in the main text.

they receive accurate signals of potential referrals, and know that they must give up large amounts of social payments to get a referral who is more likely to qualify. When CAs have networks which are not shallow but have poor information, they can find potential referrals with higher expected performance levels at relatively low cost, but poor information levels means that even a large increase in expected performance is associated with only a small increase in the likelihood of qualification. Both conditions result in CAs making referrals without a substantial performance premium. Conditions under which we can separate shallowness from poor information are discussed in section 7.2.1.

Finally, we note that while the performance premium leads to a clear test on differences in underlying parameter distributions, it also represents a useful piece of information for firms: if there are gender differences in the performance premium, then firms who use referrals to maximize screening may encourage those differences.

6   Do social payments disadvantage women?

As we discuss above, referrals under fixed fee treatments provide consistent estimates of characteristics of the index partner of each gender, partner 1. This allows us to draw several inferences: first, if we repeat Panel B of Table 1 using only CAs in fixed fee treatments, we can infer whether male and female CAs are closer to men or women. That analysis is presented in Panel D of Table 1, which indicates that 75% of male CAs are closest to men, which contrasts to 57% of female CAs. From proposition 1, we can conclude directly that male CAs get their highest social payments from other men, which can lead to across the board disadvantages to women coming from male referrals. Thus, while we rejected the idea that networks were sufficiently homophilic so that men were unconnected to women (and vice versa), we can conclude that there is homophily in social incentives, at least for men: the closest people for men tend

22

to be men.

Panel D of Table 1 also indicates that, under fixed fees, women remain much more likely to make references who will not qualify than men, with 60% of men's fixed fee references qualifying against 37% of women's. Once again, this difference is highly significant, and these numbers are very similar to the overall performance gap between men's and women's referrals. Since only closeness directly affects referral choice under fixed fees, we can conclude that for women, social payments are maximized by referring low ability people. On net, the evidence from the fixed fee treatments leads us to conclude that differences in social payments contribute to women's disadvantage both because men experience higher social payments by referring men and because women experience higher social payments by referring the unqualified. The potential of CA concerns about competing with their referrals as an explanation for these results is discussed in section 7.1.

7   The role of quality, shallowness and information

Proposition 3 suggests that three main factors could affect women's opportunities when an employee has a stake in the referral's performance (our performance pay contracts): quality (of the person who gives the largest social payments), information, and the extent of the tradeoff between social payments and network quality (network shallowness). Of these three factors, quality can be directly measured using the fixed fee treatments. Therefore before discussing the empirical evidence on performance premia, we first ask whether there are quality differences across networks.

7.1 Are there quality differences across networks?

Figure 4 presents kernel densities of the ability of men's male and female referrals recruited under fixed fees. The two distributions overlap, and a Kolmogorov-Smirnov test does not statistically differentiate them. If anything, it appears that the quality of men's networks of women dominates that of men's networks of men. We conclude, therefore, men's preference for referring men is not driven by differences in the quality of their closest women and men.

For women's networks, in contrast, there is a sharp difference. Figure 5 also presents kernel densities for the women and men who are closest to female CAs. The ability distribution of men who are closest to women clearly stochastically dominates the distribution of women who are closest to women, with the Kolmogorov-Smirnov test rejecting the distributions being the same at the 5% level. In terms of means, women who are closest to women perform on average 0.42 of a standard deviation below the CA mean, while men who are closest to women perform 0.08 standard deviations below the CA mean. Women CAs' closest men perform a bit better than the men (or women) who are closest to men, though not statistically different[19]. Our results therefore indicate that women are closest to women who are particularly low ability.

An alternative mechanism behind women's tendency to refer low ability individuals is that women in particular may be more averse to competition than men (despite the firm's motivation of wanting to hire more women)[20]. Competition is likely more salient in the context of this experiment than in other employment contexts where existing employees make referrals, though we note that competition is certainly present there as well. Existing employees may fear the referral will perform better and make the CA look bad, or compete with the CA over promotions. Compared to our setting where the referral only marginally affects the likelihood

---

[19]Columns 1 and 2 of Table 3 show similar results as the figures, but pooled across all treatments.

[20]Niederle and Vesterlund (2007) find that women shy away from competition in particular when competing with men. In our context, this would lead women to either not make a referral or refer poorly qualified men. This is not what we observe.

of qualifying or getting called for a job (given the large number of recruits)[21], competition on the job may actually be stronger.

Nevertheless, if women CAs are concerned about the competitive threat their referrals pose, they may choose to either forgo the finder's fee (and not make a referral) or refer someone who is unlikely to qualify. We do not observe the former, as the referral rate is almost identical among women CAs and male CAs. However, the the latter is consistent with the results presented in Table 1: in unrestricted treatments, women refer poor quality men and women. Nonetheless, a closer examination of the data suggest several patterns which are not consistent with a strong role of competition-aversion. Figure 5 shows that women in the restricted-men treatment refer high quality men. This is not consistent with competition. If competition were driving behavior, women in those treatments who did not know a low quality man would not make a referral. Recall from Table 2 that we do not observe any differential attrition rate among women in different gender-restricted treatments. Instead, it suggests that the women whose closest men are high quality - those leading to better quality male referrals in Figure 5 - would get higher social payments from women who are low quality, and hence refer women who are less likely to qualify in the unrestricted treatment.

Figure 3 is also inconsistent with the competition hypothesis: women who are on the margin of qualification (near a score of 60) are if anything more likely to refer someone who is qualified. Finally, we also discuss a direct test of the role of competition in the Appendix: in an additional cross-randomized treatment, we experimentally varied whether the CA was directly competing with their referral. We find no differences in the quality of the referrals. Taken together, these pieces of evidence all point towards strong social payments within women's networks, and not women having a pronounced fear of competition driving women's choices.

---

[21]On the median CA recruitment date, there were 61 CAs all of whom applied at the same time; given that all CAs were asked to make a referral this renders one's own referral just one competitor out of over 100 even ignoring CA beliefs about other recruitment dates.

Figures 4 and 5 are about the quality of closest people, rather than the full distribution of network quality. It is also possible that quality is different across the range of the network: this could make the tradeoff between social payments and quality more steep, which is captured by the model's concept of shallowness and is investigated in the next section.

7.2 Are there gender differences in information or shallowness?

Proposition 3 suggested that the performance premium, Ψ^g, is a useful statistic for identifying differences in information, shallowness, or quality. In the previous section, we found that men's networks of men and women are similar quality, while women are closer to higher quality men than women. Applying proposition 3, if we find a larger performance premium for one gender than the other among men's referrals or among women referring women[22], and if part of that difference is attributable to improved performance under performance pay, then we will have evidence of differences in either information or shallowness.

The ability to identify good potential workers from among individuals in their social network is also directly of interest to employers. To test whether men and women are willing and able to identify high quality men and women and to test Proposition 3, we regress

$$Y_i = \sum_k \alpha_k T_k + \delta_t + v_i$$

as before, where Y_i is an indicator for referring a qualified referral, T_k are the treatment categories in terms of gender and contract structure, and δ_t are time trends. Once again, CAs in restricted male, fixed fee treatments are used as the excluded group.

Table 3 presents the results of this analysis. Women do not demonstrate an overall

---

[22]Strictly, the point estimates suggest that men's networks of women may be higher ability than their networks of men, though the difference is far from being statistically significant. As a result, evidence of poor information or greater shallowness would be less conclusive if men experienced a higher performance premium when referring women; as we document below this is not the case.

performance premium under any treatments as seen in column (4) of Table 3. Whether women CAs are responding at all to performance pay is discussed further in section 7.2.2.

Male CAs, by contrast, experience a substantial performance premium when referring men: column (3) shows that when restricted to refer men, male CAs refer someone who is about 27 percentage points more likely to qualify when assigned to the performance pay treatment. Given that the qualification rate is about 50%, this is a very large premium. However, they do not experience any performance premium when restricted to refer women. We therefore conclude that male CAs have useful information for employers about men, and the price of eliciting the information is not prohibitively high. Proposition 3 applies if the qualification rates of men referred in performance premium treatments is higher than for women: Simple descriptive statistics reveal that among male CAs offered performance pay, 65% of referrals qualify in the male-only treatment vs 47% in the female-only treatment. As a result, given our results on quality, we conclude that either men have worse information about women, or their networks of women are more shallow.

As we discuss above, information and shallowness are actually quite similar in principle, as both act to make the slope between likelihood of qualification and social payments more steep. The next subsection discusses the conditions under which we can disentangle them.

7.2.1 Distinguishing shallowness from poor information

Poor information makes the expected likelihood of any referral qualifying close to one-half, as individuals with low quality signals still face a strong chance of qualifying, and individuals with high quality signals are still not tremendously likely to qualify. In other words, a network with poor information looks much like a very shallow network with $\lambda_1^g \cong 1/2$. In contrast, a network can be highly shallow with any value of λ_1^g. This distinction can lead to a test of the relative

contribution of shallowness versus information to the lower premium gender.

More specifically, note that CAs who are unrestricted in their choice of gender may pull from their network of gender g or g'. We note that even under performance contracts, CA incentives are not perfectly aligned with the employers, as they must still weigh their social incentives against expected performance pay. However, if we see a larger performance premium in the same gender which gives higher social payments, then allowing the option to refer either gender may have different effects if a network is more shallow or if it has worse information.

Proposition 4 *The effect of allowing either gender on referral performance is ambiguous if the performance premium is larger for gender g than for gender g'. Suppose that gender g is closer than gender g' $\left(\alpha_1^g \geq \alpha_1^{g'} \right)$ and the quality of gender g' networks is weakly lower than the quality of gender g $\left(Q_1^g \geq Q_1^{g'} \right)$. If the network of gender g' is more shallow than the network of gender g' but has similar information, then the unrestricted performance premium (Ψ^u) is equal to Ψ^g. If in contrast, information is better about gender g than about gender g' $\left(\sigma_\varepsilon^g < \sigma_\varepsilon^{g'} \right)$, then $\Psi^u \leq \Psi^g$. However, if quality of gender g' is greater than quality of gender g, $\left(Q_1^g < Q_1^{g'} \right)$ no sharp distinction can be drawn.*

Under performance pay, individuals are maximizing a weighted sum of expected performance pay and social incentives. Table 3 showed that the performance premium was larger for male than female referrals which could be attributed to either better information or less shallowness in networks of men. When CAs are allowed the choice between genders, the optimal choice may be the optimal male referral under performance pay, who is higher ability, or it may be the woman chosen under the female only-performance pay treatment if she gives enough social payments to compensate for the loss in expected performance pay. One difference between information and shallowness is that shallowness unambiguously lowers the returns to referring that gender under performance pay, while worse information may increase or decrease

the return of that gender. Proposition 4 formalizes this distinction: if men contribute higher social payments, are similar or higher quality, and tradeoffs in the male network are also less shallow, then the male network dominates. Allowing the CA to refer anyone does not change referral choices from the restricted male treatment and the unrestricted performance premium would resemble the restricted male performance premium. However, there is an advantage to referring the female performance pay partner in two scenarios. The first scenario is when the CA has worse information about women: in this case high social payment, but poor signal, individuals have a greater expected chance of qualifying and may induce the CA to select a woman, reducing the performance premium. The second is when the network of women is both higher quality and more shallow, in which case the unrestricted performance premium may be smaller than the restricted male premium.

Therefore, a striking result from Table 3 is that performance premia are in fact lower for men under unrestricted treatments than under male restricted treatments. Proposition 4, then, tells us that this result is indicative of poor information about women's capabilities contributing to the lower performance premium if we interpret the estimates in section 7.1 as stating that men's female referrals are similar in quality to their male referrals. Indeed there is no significant difference in their performance, and point estimates of the difference are fairly small. However, if we take the point estimates at face value, then there is the possibility that women network members are higher quality than male network members. If so, we cannot disentangle shallowness from information in the model. We note also that this result is particularly problematic for women if networks are utilized for screening: not only do employers maximize screening by having men refer men, they do better by discouraging them from even considering female referrals.

An additional test we can run compares the quality distributions of mens' referrals of

women under fixed fees with the distribution of mens' referrals of women under performance pay. The motivation for such a test is that we know that some men are choosing not to make a referral when restricted to refer one gender under performance pay. If men have good information, but shallow networks of women, we may anticipate that the men who choose not to make a referral are those who are closest to particularly low quality women. In contrast, if men have poor information about women, then return rates should be similar for men who are closest to women who are actually low quality. Figure 6 presents this result, and we see that the left tail of the distribution of female quality is at least as heavy in the performance pay treatment as in the fixed fee treatment, suggesting that there is not a correlation between women's likelihood of qualification and whether the male CA finds it worthwhile to refer them. Together, we interpret these two pieces of evidence as indicating that poor information about women is likely to be at least part of the answer.

7.2.2 Referral Characteristics

Tables 4 and 5 explore further the differences in how performance pay affects the way men choose to refer men and women by asking how men's referrals perform on various components of the test.

Table 4 finds that men referred by men under performance pay do statistically significantly better on the computer knowledge part of the exam and better (though not significantly) on most of the other components, whereas the women they refer under performance pay behave quite similarly on all components as the women they refer under fixed fees. Following the model, this is very consistent with the hypothesis that men are referring their index female partner under both contracts.

Women CAs did not demonstrate a performance premium on average when referring

men or women. Table 5 again disaggregates referral performance by component for women CAs, however, and finds that women are changing their optimal referral choices of both men and women. When we provide performance pay, women refer women with better English skills and who solve more ravens matrices correctly (though the latter is insignificant), and they refer men who are more likely to have worked for a survey firm in the past and who perform better on the practical exam. However, neither of these improvements translate to higher qualification rates because they are also associated with worse scores on other components. The more experienced men also have worse math skills, while the women with better language skills perform weakly worse on a number of characteristics[23]. These suggest that women are responding to performance pay and do have some useful information for employers, particularly about other women (as cognitive ability is likely harder to observe in a resume than past experience) but that they do not have enough information or have networks which are too shallow to choose women or men who are likely to qualify (at the level of incentives offered in the experiment).[24]

7.2.3 Attrition

In section 4, we made note of the fact that there was strong evidence that male CAs were more likely to make a referral in the presence of fixed fees than performance pay, and weaker evidence that female CAs responded similarly. In principle, these differential return rates could influence our estimates of the performance premium, though the fact that we rely on differences between restricted-gender treatments (where return rates were identical) does ameliorate this concern. Still, for example, one interpretation which would be qualitatively consistent with presented

---

[23]The total effect on women referred under performance pay is the sum of the performance pay component and the interaction between performance pay and female restricted treatment; this is never significantly different from zero though often negative in point value.

[24]Note also that this makes the competition hypothesis also less likely: if women have a hard time anticipating who will qualify, referring low quality people instead of just not making a referral is a very risky strategy.

results is that all CAs will only refer person 1, but CAs will just attrit rather than refer person 1 under performance pay if they are in a restricted male treatment and person 1 is low quality. Figure 7 suggests that among men, there is assortative matching in ability between the CA and their referral under fixed fees. However, figure 7 also shows that the performance premium exists throughout the entire distribution of CA test scores, which make attrition bias less likely to be driving the results on information / shallowness[25]. Figure 8 shows analogously that there is little evidence of female CAs responding to the performance pay incentive at any point in the CA performance distribution.

Thus, the composition of CAs making referrals does not seem consistent with attrition being the only mechanism. Moreover, even if attrition plays an important role, Table 3 is still evidence of male CAs having more information about men than about women. Male CAs were less likely to make a referral under performance pay, at the same rate, in both restricted gender treatments. However, only the male referrals in the performance pay treatment performed better. Poor information about women would be consistent with this: while male CAs attrit when they anticipate not having a high quality referral, the female referrals in the performance pay treatment are no different than those in the fixed fee treatments since the CAs' quality signals are not strongly correlated with actual performance.

8   Quality expectations and social payments

Could expectations of quality be a contributing factor to women's disadvantage even in fixed fee treatments? In particular, CAs may receive higher social payments when their referrals qualify than when they do not. This could occur because referrals who actually get the job give

---

[25]This also demonstrates that in this setting, the problem of social payments generating incentives for employees not to refer the best person to the firm takes place across the ability distribution. We therefore have more confidence that the results extrapolate to other contexts where only existing employees make referrals.

higher transfers to the references who helped them than referrals who only get the opportunity to apply, creating a relationship between the CA's expectations of quality and social payments. Alternatively, CAs could believe that they will receive a reputational premium from IPA for making a highly skilled referral. While this is conceptually very different from the social payments that we have modelled, the empirical arguments ruling out reputational concerns as the only consideration affecting the referral choice behavior we observe in the experiment are the same as social payments being tied to referral qualification, and therefore we do not present a separate sub-section on this point.

In the model, expectations of referral quality are unbiased as ε_j^g is mean zero for both men and women. Recall that the quality of men and women are similar in men's networks (Figure 4), and that women's (low quality) fixed fee referrals are also not the highest quality people they know, since some know high quality men (documented in Figure 5). If social payments were a function of qualification, with unbiased expectations we would see - in contrast to our findings - that women systematically refer men, and men refer similar numbers of men and women.

However, CA expectations of referral quality may be biased. In that case, ε_j^g is not mean zero, and differences in CA i's expectations over quality $(E_i\left[Q_1^g + \varepsilon_1^g\right])$ would have similar implications in the model to actual differences in quality (Q_1^g). Therefore, biased expectations of referral quality may be able to explain our results if social incentives are positively related to referral quality in the overall network.

We can point to several pieces of evidence that biased expectations alone are not driving the behavior that leads to few qualified women getting referred. First, Table 2 showed that attrition was higher among male CAs who were in the performance pay - male only referral treatment cell. If social incentives or reputational concerns cause men to value referral quality,

then they would refer strong candidates even in the fixed treatment, making the attrition result very surprising (particularly as performance pay has a higher expected value if you expect your referral to have at least a 5/13 chance of being above the median). Second, recall that table 3 showed that men did not refer the highest ability men in their network under fixed fees, since the performance pay led to more qualified male referrals. Therefore it can not be the case that social payments are completely indexed to qualification: if they were, there would be no additional impact of the performance pay. This intuition - that social incentives indexed to qualification in essence acts as performance pay - also leads to two additional insights.

Table 1 indicates that men systematically refer men under fixed fees. In principle, this could be explained by expectations about women's relative quality if men expect their closest male connections to outperform their closest female connections. In that case, however, men have even greater incentives to refer men under performance pay than under fixed fees, and so we should expect men to refer men even more frequently under performance pay than under fixed fees[26]. However, as Table 1 Panel E indicates, men refer men at the same rate under performance pay as under fixed fees, and those men qualify more frequently than the women they (endogenously) choose to refer (though not significantly so, owing in part to the small fraction of women referred), which is inconsistent with men referring other men due to a systematic underestimation of womens' abilities. A similar argument suggests that women are not referring low quality people due to biased expectations of performance (in particular, the people they refer under performance pay are at least as good as those referred under fixed fees).

Finally, with this explanation we would expect men, who incorrectly anticipate women being less likely to qualify and receive higher social payments when referrals qualify, to differ-

---

[26]Unless they refer women who are much better (by more than their bias) than the men they know

34

entially attrit when they are required to refer a woman. In table 2 we see that is neither the case in fixed or in performance pay treatments. Taken together, quality expectations appear to play a small, if any, role in explaining why so few qualified women get recruited through referrals.

9   Conclusion

There is a large literature in economics and sociology which has used observational data to suggest that women benefit less from job networks than men do (Ioannides and Loury, 2004). Using an experiment designed around a recruitment drive for real-world jobs, we provide direct evidence that the use of referral systems puts women at a disadvantage. We find that qualified women tend not to be referred by networks for two reasons: first, men exhibit a preference for referring men, and second, women tend to refer unsuitable candidates. This result suggests that the ubiquity of job networks as a hiring system could contribute to persistent gender gaps in wages. As with any experiment, our results are only internally valid for our sample, enumerator applicants in Malawi. However, given that they closely mirror stylized facts about gender and networks which are based on a wealth of observational studies primarily in the US and Europe[27], there is reason to believe that our findings may generalize to many other contexts.

Our experiment allows us to say several additional things about the structure of networks which could be driving the observed disadvantage for women. Social incentives within both men's and women's networks lead to fewer qualified women being referred for a job. Employers are faced with a choice of using male employees to make referrals, in which case social incentives are maximized by permitting women's disadvantage; or using female employees to make a referral, in which case relatively few of the people referred are qualified. This suggests that

---

[27]Most recently, Lalanne and Seabright (2011) find that male executives in the US and Europe have salaries which increase in numbers of executive contacts, while female executives do not receive this benefit.

unless employers design referral contracts to contradict these incentives (at additional cost), women are disadvantaged by social network-based recruitment.

We also find that men have a greater potential to screen men than to screen women, and some weaker evidence that women may be able to screen both genders in different ways. If employers use employee referrals to improve screening, then our results suggest that employers have the incentive to encourage only men to refer other men. This result suggests that in order to prevent discrimination against women, careful hiring procedures may need to be followed. One procedure which may be effective is a quota-based hiring system, requiring the firm to consider some female applicants. The firm can find female applicants through conventional hiring means and mens' (high quality) networks of women, which may allow the firm to identify the good candidates of both genders.

References

Altonji, J. and R. Blank (1999). Race and gender in the labor market. *Handbook of Labor Economics 3*, 3143–4259.

Arbache, J. S., A. Kolev, and E. Filipiak (Eds.) (2010). *Gender Disparities in Africa's Labor Market*. World Bank Publications.

Beaman, L. and J. Magruder (2012). Who gets the job referral? Evidence from a social networks experiment. *American Economic Review 102*, 3574–3593.

Bell, D. and M. Reich (Eds.) (1988). *Health, Nutrition, and Economic Crises: Approaches to Policy in the Third World*. Auburn House Publishing Company.

Bertrand, M. (2011). New perspectives on gender. *Handbook of Labor Economics 4*, 1543–1590.

Bewley, T. F. (1999). *Why Wages Don't Fall During a Recession*. Harvard University Press.

Black, S. E. and E. Brainerd (2004). Importing equality? The impact of globalization on gender discrimination. *Industrial and Labor Relations Review 57*, 540–559.

Black, S. E. and P. E. Strahan (2001). The division of spoils: Rent-sharing and discrimination in a regulated industry. *American Economic Review 91*, 814–831.

Blau, F. D. and L. M. Kahn (2000). Gender differences in pay. *Journal of Economic Perspectives 14*, 75–99.

Calvo-Armengol, A. and M. O. Jackson (2004). The effects of social networks on employment and inequality. *American Economic Review 94*, 426–454.

Doeringer, P. B. and M. J. Piore (1971). *Internal Labor Markets and Manpower Analysis.* M.E. Sharpe, Inc.

Fernandez, R. M. and M. L. Sosa (2005). Gendering the job: Networks and recruitment at a call center. *American Journal of Sociology 111*, 859–904.

Goldin, C. and L. F. Katz (2002). The power of the pill: Oral contraceptives and women's career and marriage decisions. *Journal of Political Economy 110*, 730–770.

Ioannides, Y. M. and L. D. Loury (2004). Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature 42*, 1056–1093.

Lalanne, M. and P. Seabright (2011). The old boy network: Gender differences in the impact of social networks on remuneration in top executive jobs. *IDEI working paper No. 689*.

McPherson, M., L. Smith-Lovin, and J. M. Cook (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology 27*, 415–444.

Montgomery, J. D. (1991). Social networks and labor market outcomes: Toward an economic analysis. *American Economic Review 81*, 1408–1418.

Mortensen, D. and T. Vishwanath (1994). Personal contacts and earnings: It is who you know! *Labour Economics 1*, 187–201.

Mouw, T. (2006). Estimating the causal effect of social capital: A review of recent research. *Annual Review of Sociology 32*, 79–102.

National Statistics Office (NSO) and ICF Macro (2011). *Malawi Demographic and Health Survey 2010.* NSO and ICF Macro.

Niederle, M. and L. Vesterlund (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics 122*, 1067–1101.

Rubineau, B. and R. M. Fernandez (2010). Missing links: Referrer behavior and job segregation. *MIT Sloan School Working Paper No. 4784-10*.

Tassier, T. and F. Menczer (2008). Social network structure, equality, and segregation in a labor market with referral hiring. *Journal of Economic Behavior and Organization 66*, 514–528.

World Bank Group and others (2010). *The Education System in Malawi.* World Bank Publications.

World Bank Group and others (2011). *World Development Report 2012: Gender Equality and Development.* World Bank Publications.

Table 1: Gender Distributions of CAs and Referrals

	(1)	(2)	(3)	(4)
	All CAs	Male CAs	Female CAs	Diff: p value
A. CA Characteristics				
Fraction of CAs	100%	62%	38%	
CA is qualified	53%	56%	48%	0.047
N	767	480	287	
B. CA Characteristics: Made Referral, Either Gender Treatments				
Fraction of CAs	100%	61%	39%	
CA is qualified	57%	62%	49%	0.061
N	217	130	87	
C. Referral Characteristics: Made Referral, Either Gender Treatments				
Referral is Female	30%	23%	43%	0.002
Referral is Qualified	49%	56%	38%	0.019
Referral is Qualified Male	34%	43%	22%	0.002
Referral is Qualified Female	14%	13%	17%	0.456
N	195	117	78	
D. Referral Characteristics: Made Referral, Either Gender, Fixed Fee Treatments				
Referral is Female	32%	25%	43%	0.042
Referral is Qualified	50%	60%	37%	0.012
Referral is Qualified Male	34%	44%	20%	0.007
Referral is Qualified Female	16%	16%	16%	0.983
N	117	68	49	
E. Referral Characteristics: Made Referral, Either Gender, Perf Treatments				
Referral is Female	31%	22%	45%	0.039
Referral is Qualified	46%	49%	41%	0.521
Referral is Qualified Male	35%	41%	24%	0.138
Referral is Qualified Female	12%	8%	17%	0.231
N	78	49	29	

Table 2: Probability of Making a Referral

	(1)	(2)	(3)		(4)
Female Treatment	-0.004	-0.055	-0.004		-0.042
	(0.038)	(0.054)	(0.050)		(0.074)
Either Gender Treatment	0.014	0.017	-0.052		-0.024
	(0.040)	(0.055)	(0.052)		(0.071)
Performance Pay			-0.148	***	-0.113
			(0.056)		(0.080)
Perf Pay * Female Treatment			0.004		-0.013
			(0.076)		(0.111)
Perf Pay * Either Treatment			0.152	*	0.086
			(0.079)		(0.110)
Observations	506	310	506		310
CA Gender	Men	Women	Men		Women

Notes

1 The dependent variable is an indicator for whether the CA makes a referral.

2 All specifications include CA visit day dummies.

Table 3: Referral Performance

	Referral Qualifies					
	(1)	(2)		(3)		(4)	
Female Referral Treatment	-0.030	-0.190	**	0.068		-0.181	
	(0.062)	(0.083)		(0.081)		(0.113)	
Either Gender Treatment	0.071	-0.231	***	0.227	***	-0.242	**
	(0.066)	(0.082)		(0.084)		(0.107)	
Performance Pay				0.267	***	0.021	
				(0.093)		(0.122)	
Perf Pay * Female Treatment				-0.248	*	-0.022	
				(0.127)		(0.171)	
Perf Pay * Either Treatment				-0.383	***	0.032	
				(0.132)		(0.169)	
Observations	390	227		390		227	
CA Gender	Men	Women		Men		Women	

Notes

1 The dependent variable is an indicator for the referral qualifying.

2 All specifications include CA visit day dummies.

Table 4: Screening of Male CAs on Different Characteristics

	Survey Experience	Tertiary Education	Math Score	Language Score	Ravens score	Computer Score	Practical Exam Score	Feedback points
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female Referral Treatment	-0.033	0.045	-0.017	-0.115	-0.092	0.062	1.033	3.003***
	(0.069)	(0.074)	(0.142)	(0.207)	(0.194)	(0.371)	(0.661)	(1.044)
Either Gender Treatment	0.040	0.072	0.009	0.087	0.089	0.623	1.378**	1.856*
	(0.072)	(0.077)	(0.148)	(0.215)	(0.203)	(0.387)	(0.689)	(1.089)
Performance Pay	0.080	0.067	0.134	-0.005	0.230	0.943**	0.496	1.883
	(0.080)	(0.085)	(0.164)	(0.238)	(0.224)	(0.428)	(0.757)	(1.197)
Perf Pay * Female Treatment	-0.075	0.025	-0.259	-0.027	-0.293	-0.915	-0.950	-2.443
	(0.108)	(0.116)	(0.223)	(0.325)	(0.305)	(0.583)	(1.026)	(1.622)
Perf Pay * Either Treatment	-0.165	-0.083	-0.065	-0.169	-0.367	-0.856	-1.768*	-3.371**
	(0.113)	(0.121)	(0.232)	(0.338)	(0.318)	(0.607)	(1.069)	(1.696)
Observations	386	390	390	390	390	390	383	382

Notes

1 The dependent variable is an indicator for the referral qualifying.

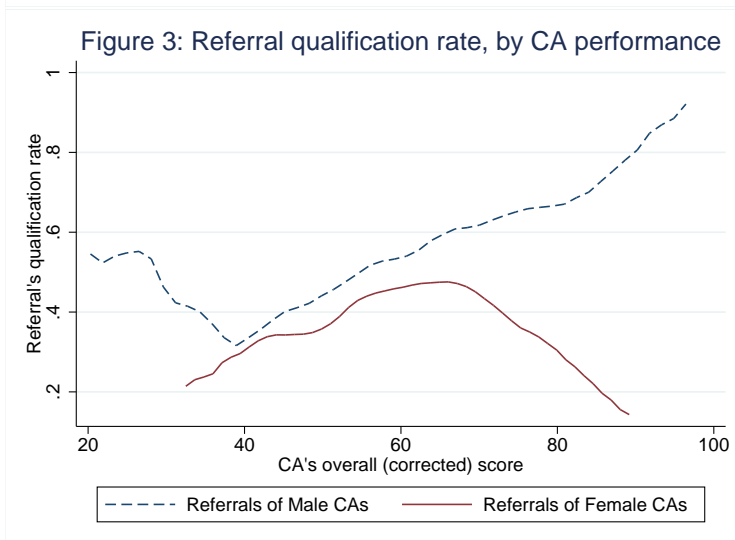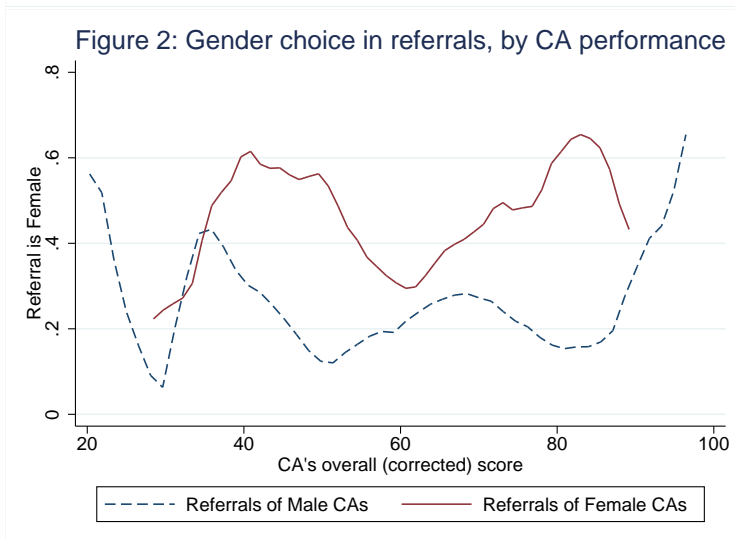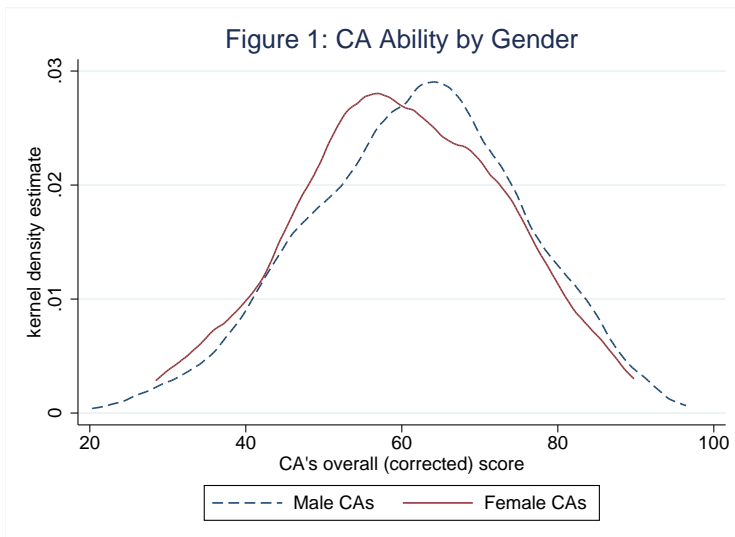2 All specifications include CA visit day dummies.

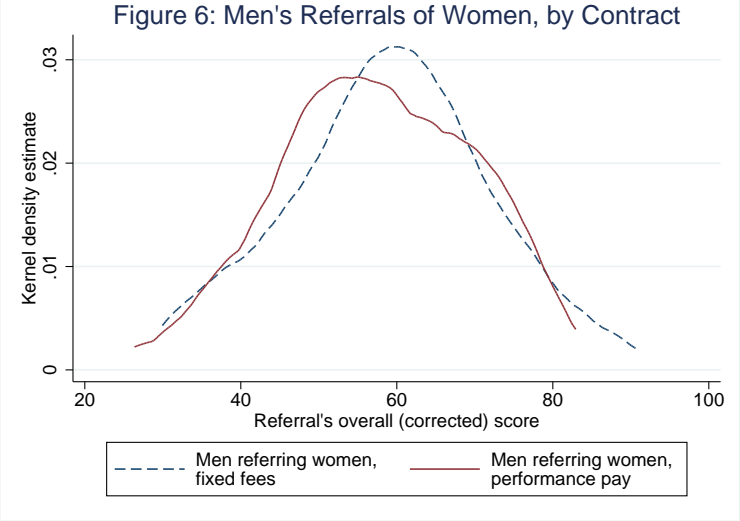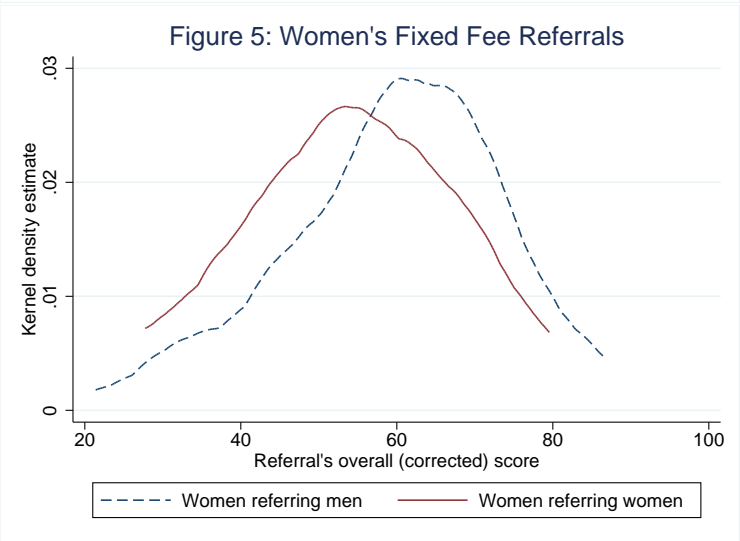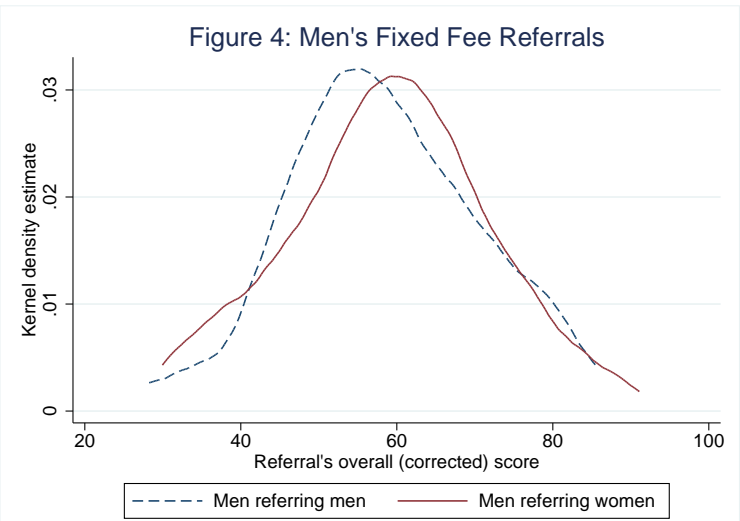Table 5: Screening of Female CAs on Different Characteristics

	Survey experience	Tertiary Education	Math Score	Language Score	Ravens score	Computer Score	Practical Exam Score	Feedback points
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female Referral Treatment	0.032	0.151	-0.332	-1.140 ***	-0.435	-0.627	0.972	2.152
	(0.091)	(0.110)	(0.216)	(0.342)	(0.270)	(0.538)	(0.963)	(1.349)
Either Gender Treatment	0.040	0.017	-0.189	-0.246	-0.172	-0.139	0.015	0.879
	(0.086)	(0.104)	(0.205)	(0.324)	(0.256)	(0.509)	(0.910)	(1.274)
Performance Pay	0.264 ***	0.143	-0.400 *	-0.465	-0.175	0.419	1.832 *	1.604
	(0.098)	(0.119)	(0.234)	(0.370)	(0.293)	(0.582)	(1.056)	(1.479)
Perf Pay * Female Treatment	-0.320 **	-0.292 *	0.402	1.330 **	0.551	0.232	-2.164	-2.134
	(0.138)	(0.166)	(0.326)	(0.515)	(0.408)	(0.811)	(1.468)	(2.055)
Perf Pay * Either Treatment	-0.270 **	-0.052	0.368	0.500	-0.260	-0.372	-1.625	-4.511 **
	(0.136)	(0.164)	(0.323)	(0.510)	(0.403)	(0.802)	(1.448)	(2.027)
Observations	226	227	227	227	227	227	222	222

Notes

1   The dependent variable is indicated in the column heading.

2   All specifications include CA visit day dummies.

41

Figure 1: CA Ability by Gender

Figure 2: Gender choice in referrals, by CA performance

Figure 3: Referral qualification rate, by CA performance

Figure 4: Men's Fixed Fee Referrals

Figure 5: Women's Fixed Fee Referrals

Figure 6: Men's Referrals of Women, by Contract

Figure 7: Referral Qualifies , by Male CA performance

– – – Men referring women, fixed fees	—— Men referring men, fixed fees
▪▪▪ Men referring women, performance pay	▬▬ Men referring men, performance pay

Figure 8: Referral Qualifies , by Female CA performance

– – – Women referring women, fixed fees	—— Women referring men, fixed fees
▪▪▪ Women referring women, performance pay	▬▬ Women referring men, performance pay

A   Appendix

A.1   Proofs

Individuals choose referral j to maximize

$$\pi_i\left(\alpha_j, P_i, F_i\right) = \max_{j \in \mathcal{G}} P_i\left(1 - \Phi\left(\frac{60 - Q_1^g - \gamma^g\left(\alpha_1^g - \alpha_j^g\right)}{\sigma_\varepsilon^g}\right)\right) + F_i + \alpha_j^g$$

which yields FOC

$$\frac{\sigma_\varepsilon^g}{P_i \gamma^g} = \phi\left(\frac{60 - Q_1^g - \gamma^g\left(\alpha_1^g - \alpha_j^{g*}\right)}{\sigma_\varepsilon^g}\right) \tag{3}$$

and SOC

$$\left(\frac{60 - Q_1^g - \gamma^g\left(\alpha_1^g - \alpha_j^{g*}\right)}{\sigma_\varepsilon^g}\right) < 0 \tag{4}$$

if there is an internal solution (that is, if a network member other than the index partner is chosen).

44

A.1.1 Proposition 1

When $P_i = 0$, the objective function reduces to $\alpha_j + F_i$. Proposition 1 immediately follows as neither quality (Q_1^g), shallowness (γ^g) or information (σ_ε^g) affect the CA's referral choice.

A.1.2 Proposition 2

Define Q_p^{g*} and λ_p^{g*} as the quality and probability of qualifying, respectively, of the network member who satisfies equations 3 and 4. That is, the network member selected when there is an internal solution. The second order condition (equation 4) tells us that $Q_P^{g*} > 60$. Using that, the first order condition tells us that $\sigma_\varepsilon^g / P_i \gamma^g$ uniquely determines the solution j: since the density is monotonically decreasing for values of $Q_j^g > 60$, we can be assured that a unique value of the density is associated with a unique value of λ_p^{g*} and Q_p^{g*} :

$$\lambda_P^{g*} = 1 - \Phi \left(\frac{60 - Q_P^{g*}}{\sigma_\varepsilon^g} \right)$$

and

$$Q_P^{g*} = 60 + \sigma_\varepsilon^g \phi^{-1} \left(\frac{\sigma_\varepsilon^g}{P_i \gamma^g} \right)$$

which gives proposition 2. This implies, among other things, that the probability that the optimal referral qualifies is independent of α_1^g and Q_1^g when there is an internal solution.

Now, we define what we call the network constraint: network member j who is of quality Q_p^{g*} gets referred if

$$P_i \left(\lambda_p^{g*} - \lambda_1^g \right) > \alpha_1^g - \alpha_j^{g*} \tag{5}$$

Define the potential performance premium, Ψ^{g*}, as $\lambda_p^{g*} - \lambda_1^g$, that is, what the performance premium would be if an internal solution was selected. Equation 5 can then be re-written as

$$\Psi^{g*} > \frac{\alpha_1^g - \alpha_p^{g*}}{P_i} \tag{6}$$

A.1.3 Proposition 3

Proposition 3 considers differences across 2 networks, where $\Psi^g = \lambda_p^g - \lambda_1^g > \Psi^{g'}$. If $\Psi^g > \Psi^{g'}$, it follows that at least one of the underlying network characteristics must be different in a way which increases the performance premium.

For a given network, performance premia can be increased in two ways: either through an increase in λ_p^{g*} (the qualification rate of the optimal performance premium referral when a partner other than the index partner is chosen under performance pay) and through inducing a switch from partner 1 to partner g in response to performance pay. Either of these can be related to network characteristics.

Proposition 3 applies when $\lambda_p^g > \lambda_p^{g'}$, that is, when at least part of the performance premium is attributable to an increase in qualification rates under performance pay (as opposed to $\lambda_1^g < \lambda_1^{g'}$). We start with a lemma which formalizes parts of proposition 2 and will be used in the subsequent proof of proposition 3.

Lemma 1 *For a given contract (F_i, P_i) if $\lambda_p^g > \lambda_p^{g'}$ then at least one of the following must be true:*

45

i): $\sigma_\varepsilon^g < \sigma_\varepsilon^{g'}$ (there is better information about gender g)
ii): $\gamma^g > \gamma^{g'}$ (the network of gender g is less shallow than the network of gender g')
iii): $\lambda_p^{g'} = \lambda_1^{g'}$ (there is no performance premium in network g')
iv): $\lambda_p^g = \lambda_1^g$ (there is no performance premium in network g)

Proposition 2 guarantees that if there is an internal solution for both gender g and g', then either (i) or (ii) is true. (iii) and (iv) cover the other cases. We then consider which network characteristics can lead to $\Psi^g > \Psi^{g'}$ and, through the lemma, their implications for the needed distribution of other network characteristics if $\lambda_p^g > \lambda_p^{g'}$.

A.1.3.1 Closeness Note that an increase in closeness of one gender g to g' implies that for any given level of expected performance, the social payments are higher from gender g than from gender g'.

From proposition 2, changes in closeness do not affect λ_p^{g*}. Since λ_1^g is also unaffected, Ψ^{g*} therefore is invariant to closeness. We need only be concerned with whether they affect the network constraint. For the network constraint, differentiate equation 5 with respect to α_1^g, we get that the network constraint unchanged. To see this, consider the following: λ_p^{g*} and λ_1^g are unaffected, and $\frac{\partial \alpha_1^g}{\partial \alpha_1^g} = \frac{\partial \alpha_p^{g*}}{\partial \alpha_1^g} = 1$. The relationship between α_p^{g*} and α_1^g can be seen by differentiating the FOC with respect to α_1^g: since the LHS equals 0, it must be that the optimal social payment adjusts as well.

Thus, changing α_1^g does not affect the tightness of the network constraint at the optimum, and differences in closeness cannot explain differences in performance premia across two networks.

A.1.3.2 Information To start, we consider Ψ^{g*}. Then:

$$\frac{\partial \Psi^{g*}}{\partial \sigma_\varepsilon^g} = \frac{\partial \lambda_P^{g*}}{\partial \sigma_\varepsilon^g} - \frac{\partial \lambda_1^g}{\partial \sigma_\varepsilon^g}$$

Note that

$$\frac{\partial \lambda_1^g}{\partial \sigma_\varepsilon^g} = \frac{60 - Q_1^g}{\sigma_\varepsilon^{g2}} \phi \left(\frac{60 - Q_1^g}{\sigma_\varepsilon^g} \right)$$

and

$$\frac{\partial \lambda_p^{g*}}{\partial \sigma_\varepsilon^g} = \left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon^{g2}} + \frac{1}{\sigma_\varepsilon^g} \frac{\partial Q_p^{g*}}{\partial \sigma_\varepsilon^g} \right) \phi \left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon^g} \right)$$

differentiating the FOC with respect to σ_ε^g, we have

$$\frac{1}{P_i \gamma^g} = \left[\frac{60 - Q_j^{g*}}{\sigma_\varepsilon^{g2}} + \frac{1}{\sigma_\varepsilon^g} \frac{\partial Q_p^{g*}}{\partial \sigma_\varepsilon^g} \right] \left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon^g} \right) \phi \left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon^g} \right)$$

so that we can plug in and get

$$\frac{\partial \lambda_p^{g*}}{\partial \sigma_\varepsilon^g} = \frac{\sigma_\varepsilon^g}{60 - Q_p^{g*}} \frac{1}{P_i \gamma^g} \tag{7}$$

which is clearly less than 0 for $Q_P^{g*} > 60$ (as guaranteed by the SOC). In other words, the qualification probability of the optimal partner under performance pay, if there is an internal

solution, is unambiguously increasing in information.

Then,

$$\frac{\partial \Psi^{g*}}{\partial \sigma_\varepsilon^g} = \frac{\sigma_\varepsilon^g}{60 - Q_p^{g*}} \frac{1}{P_i \gamma^g} - \frac{60 - Q_1^g}{\sigma_\varepsilon^{g2}} \phi \left(\frac{60 - Q_1^g}{\sigma_\varepsilon^g} \right)$$

This is in turn clearly less than 0 if $Q_1^g < 60$. if $Q_1^g > 60$, then $\partial \Psi^{g*}/\partial \sigma_\varepsilon^g < 0$ if

$$\frac{\sigma_\varepsilon^g}{Q_p^{g*} - 60} \frac{1}{P_i \gamma^g} > \frac{Q_1^g - 60}{\sigma_\varepsilon^{g2}} \phi \left(\frac{60 - Q_1^g}{\sigma_\varepsilon^g} \right)$$

Substituting in from the FOC, this will hold if

$$\frac{\sigma_\varepsilon^g}{Q_p^{g*} - 60} \phi \left(\frac{60 - Q_P^{g*}}{\sigma_\varepsilon^g} \right) > \frac{Q_1^g - 60}{\sigma_\varepsilon^g} \phi \left(\frac{60 - Q_1^g}{\sigma_\varepsilon^{g2}} \right)$$

Since $Q_p^{g*} > Q_1^g$, this will be guaranteed if $Q_p^{g*} < 60 + \sigma_\varepsilon^g$ (since $-\phi'(x) = x\phi(x)$ is maximized at 1 and $\frac{1}{x} > x$ for $x < 1$). Immediately, for much of the parameter space, better information will lead to a larger potential performance premium, which is condition (i) of proposition 3. However, in part of the parameter space, better information can also lead to a smaller potential performance premium. The specific case where better information could lead to a smaller performance premium is when the index partner is in expectation going to qualify, and the performance pay is sufficiently large (relative to information and shallowness) that if the performance partner is selected, then he is very likely to qualify. In this case, better information can increase the qualification probability of the index partner by more than it increases the qualification probability of the interior solution who is optimal under performance pay. Next, we consider the network constraint to consider whether in these parts of the parameter space, the index partner or the interior solution gets selected so as to apply results from the lemma.

Returning to the network constraint (equation 5): fixing α_j, and taking the derivative with respect to σ_ε^g, we have that for any j the constraint binds tighter as information improves for person j if

$$\frac{\partial \Psi_j^g}{\partial \sigma_\varepsilon^g} < 0$$

Note that this is a similar condition to above. Fix α_j, then

$$\frac{\partial \Psi_j^g}{\partial \sigma_\varepsilon^g} = \left(\frac{60 - Q_j^g}{\sigma_\varepsilon^{g2}} \right) \phi \left(\frac{60 - Q_j^g}{\sigma_\varepsilon^g} \right) - \frac{60 - Q_1^g}{\sigma_\varepsilon^{g2}} \phi \left(\frac{60 - Q_1^g}{\sigma_\varepsilon^g} \right)$$

If we evaluate this constraint at $Q_j = Q_p^{g*}$, then we see that this constraint becomes unambiguously looser as information improves if either $Q_1^g < 60$ or $Q_p^{g*} \leq 60 + \sigma_\varepsilon^g$. In other words, in the part of the parameter space where better information increases the potential premium, better information also makes an interior solution exist at lower levels of performance pay. Thus, better information can lead to a larger performance premium.

What about the part of the parameter space where better information leads to a smaller performance premium?

Using

$$Q_p^{g*} = Q_1^g + \gamma^g \left(\alpha_1^g - \alpha_p^g \right)$$

rewrite equation 6 to read

$$\Psi^{g*} > \frac{Q_p^{g*} - Q_1^g}{\gamma^g P_i}$$

Use the FOC to rewrite this as

$$\Psi^{g*} > \frac{Q_p^{g*} - Q_1^g}{\sigma_\varepsilon^g} \phi\left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon}\right)$$

which can be written

$$\left(\lambda_p^{g*} - \lambda_1^g\right) = \int_{\frac{60-Q_p^{g*}}{\sigma_\varepsilon^g}}^{\frac{60-Q_1^g}{\sigma_\varepsilon^g}} \phi\left(x\right) dx > \phi\left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon}\right) \frac{Q_p^{g*} - Q_1^g}{\sigma_\varepsilon^g} = \int_{\frac{60-Q_1^g}{\sigma_\varepsilon^g}}^{\frac{60-Q_p^{g*}}{\sigma_\varepsilon^g}} \phi\left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon}\right) dx \quad (8)$$

Since $Q_p^{g*} > Q_1^g > 60$, we can be assured that

$$\phi\left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon}\right) \le \phi\left(x\right) \qquad \forall x \in \left[\frac{60 - Q_p^{g*}}{\sigma_\varepsilon^g}, \frac{60 - Q_1^g}{\sigma_\varepsilon^g}\right] \qquad (9)$$

and the inequality holds.

Thus, whenever we are in the part of the parameter space where the performance premium is decreasing in information, there is guaranteed to be a performance premium so that partner p is chosen. In other words, in the part of the parameter space where worse information could lead to a larger performance premium for gender g than for gender g', $\Psi^{g'} > 0$ and $\Psi^g > 0$. From the lemma, this indicates that if $\lambda_p^g > \lambda_p^{g'}$, then either $\sigma_\varepsilon^g < \sigma_\varepsilon^{g'}$ or $\gamma^g > \gamma^{g'}$, satisfying the conditions of proposition 3. Therefore, when $\lambda_p^g > \lambda_p^{g'}$ and $\Psi^g > \Psi^{g'}$, either $\sigma_\varepsilon^g < \sigma_\varepsilon^{g'}$ or differences in shallowness or quality leads to the larger performance premium in gender g.

A.1.3.3 Shallowness Since λ_1^g is unrelated to shallowness, we can state

$$\frac{\partial \Psi^{g*}}{\partial \gamma^g} = \frac{\partial \lambda_p^{g*}}{\partial \gamma^g}$$

Note from the FOC that a larger γ^g suggests that $\phi\left(\frac{60-Q_p^{g*}}{\sigma_\varepsilon}\right)$ must be diminishing. Since $Q_p^{g*} > 60$, this suggests that Q_p^{g*} must be increasing in γ^g. This in turn indicates that λ_p^{g*} and Ψ^{g*} are increasing in γ^g (and hence decreasing in shallowness).

What about the network constraint? Taking the derivative of both sides of inequality 5, we see that the an interior solution becomes more likely if:

$$\left[\frac{\left(\alpha_1^g - \alpha_p^{g*}\right)}{\sigma_\varepsilon^g} - \frac{\gamma^g}{\sigma_\varepsilon^g}\frac{\partial \alpha_p^{g*}}{\partial \sigma_\varepsilon^g}\right] \phi\left(\frac{60 - Q_1^g - \gamma^g\left(\alpha_1^g - \alpha_p^{g*}\right)}{\sigma_\varepsilon^g}\right) > -\frac{1}{P_i}\frac{\partial \alpha_j^{g*}}{\partial \gamma^g}$$

or, substituting in from the FOC, if

$$\left[\frac{\left(\alpha_1^g - \alpha_p^{g*}\right)}{\gamma^g} - \frac{\partial \alpha_p^{g*}}{\partial \sigma_\varepsilon^g}\right] > -\frac{\partial \alpha_p^{g*}}{\partial \gamma^g}$$

which holds. Thus, an increase in shallowness (a decrease in γ^g) causes the constraint to

48

become tighter and makes a corner solution exist for more other parameter values. Therefore, a decrease in shallowness can lead to a larger performance premium (and would also be associated with an increased λ_p^g).

A.1.3.4 **Quality** From proposition 2, network quality (Q_1^g) is unrelated to the performance of the optimal referral ignoring the network constraint. Therefore, $\partial \Psi^{g*}/\partial Q_1^g = -\partial \lambda_1^g/\partial Q_1^g < 0$. However, it may also affect the network constraint. Differentiating 5 with respect to Q_1^g and evaluating at the optimum, we have that the network constraint becomes tighter (looser) if

$$P_i \frac{\partial \lambda_1^g}{\partial Q_1^g} > (<) \frac{\partial \alpha_p^{g*}}{\partial Q_1^g} \tag{10}$$

Recall the FOC:

$$\frac{\sigma_\varepsilon^g}{P_i \gamma^g} = \phi \left(\frac{60 - Q_1^g - \gamma^g \left(\alpha_1^g - \alpha_j^* \right)}{\sigma_\varepsilon^g} \right)$$

Differentiating with respect to Q_1^g, we have

$$0 = \frac{60 - Q_p^{g*}}{\sigma_\varepsilon^g} \phi \left(\frac{60 - Q_1^g - \gamma^g \left(\alpha_1^g - \alpha_j^* \right)}{\sigma_\varepsilon^g} \right) \left[\frac{1}{\sigma_\varepsilon^g} - \frac{\gamma^g}{\sigma_\varepsilon^g} \frac{\partial \alpha_p^*}{\partial Q_1^g} \right]$$

so that

$$\frac{\partial \alpha_p^*}{\partial Q_1^g} = \frac{1}{\gamma^g}$$

substituting into 10 we have

$$\frac{\partial \lambda_1^g}{\partial Q_1^g} > (<) \frac{1}{\gamma^g P_i}$$

and substituting in from the FOC, that becomes

$$\phi \left(\frac{60 - Q_1^g}{\sigma_\varepsilon^g} \right) > (<) \frac{\sigma_\varepsilon^g}{\gamma^g P_i} = \phi \left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon^g} \right)$$

This is in principle ambiguous, and depends on whether Q_1^g or Q_p^{g*} is closer to 60. As a result, the constraint may tighten or loosen when quality increases. However, in the specific case that the constraint is tightening in quality, it turns out that there will always be an interior solution. To see this, recall from above that there is an interior solution if

$$\left(\lambda_p^{g*} - \lambda_1^g \right) = \int_{\frac{60 - Q_p^{g*}}{\sigma_\varepsilon^g}}^{\frac{60 - Q_1^g}{\sigma_\varepsilon^g}} \phi(x) \, dx > \phi \left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon} \right) \frac{Q_p^{g*} - Q_1^g}{\sigma_\varepsilon^g} = \int_{\frac{60 - Q_1^g}{\sigma_\varepsilon^g}}^{\frac{60 - Q_p^{g*}}{\sigma_\varepsilon^g}} \phi \left(\frac{60 - Q_p^{g*}}{\sigma_\varepsilon} \right) dx$$

Suppose we are in the situation where the constraint tightens. Then

$$\left(\lambda_p^{g*} - \lambda_1^g\right) = \int_{\frac{60-Q_p^{g*}}{\sigma_\varepsilon^g}}^{\frac{60-Q_1^g}{\sigma_\varepsilon^g}} \phi\left(x\right) dx > \int_{\frac{60-Q_1^g}{\sigma_\varepsilon^g}}^{\frac{60-Q_p^{g*}}{\sigma_\varepsilon^g}} \phi\left(\frac{60-Q_p^{g*}}{\sigma_\varepsilon}\right) dx$$

and we are guaranteed an interior solution. So there will definitely be an interior solution when parameters are such that an increase in quality reduces the relative payoff of an interior solution.

Practically, this means that there is a range of levels of network quality where an increase in quality can generate a shift from referring the index partner at a particular level of performance pay to referring the optimal performance premium partner. As such, it is possible that $Q_1^g > Q_1^{g'}$ to imply that $\Psi^g > \Psi^{g'}$ and $\lambda_p^g > \lambda_p^{g'}$ if $Q_1^{g'}$ is sufficiently low. Once that threshold is crossed, we have just shown that partner $p > 1$ will be referred under performance pay for the remainder of the quality distribution. Recall that λ_p^{g*} is unrelated to Q_1^g. If the only difference between two networks is in the quality of the index partner, and both networks refer partner $p > 1$ at the levels of performance pay we provide, then the magnitude of the performance premium is decreasing in quality. In principle, then, if both networks exhibit a positive performance premium, a larger performance premium in gender g could be generated by gender g being lower quality. In proposition 3, we rule this possibility out through the assumption $\lambda_p^g > \lambda_p^{g'}$. In this example, $\lambda_p^{g*} = \lambda_p^g$ and $\lambda_p^{g'*} = \lambda_p^{g'}$. In that case, following the lemma, it must be the case either that $\sigma_\varepsilon^g < \sigma_\varepsilon^{g'}$ or that $\gamma^g < \gamma^{g'}$, so that again one of the three conditions in Proposition 3 holds in practice.

A.1.4 Proposition 4

Let network characteristics superscripted with u represent the unrestricted optimum (across both genders). Note that if $\alpha_1^g > \alpha_1^{g'}$ then $\lambda_1^u = \lambda_1^g$. Thus, $\Psi^u = \lambda_p^u - \lambda_1^g$. Note that the optimal referral under performance pay will be solve $\max_{k\in\{g,g'\}} \alpha_p^k + P_i\lambda_p^k$. Then define $\alpha_{pg'}^g = \alpha^g\left(\lambda_p^{g'}\right)$, the social payment associated with the gender g person who is as likely to qualify as the optimal performance pay referral of gender g'. If $Q_1^g \geq Q_1^{g'}$, $\gamma^g > \gamma^{g'}$, and $\sigma_\varepsilon^g = \sigma_\varepsilon^{g'}$, then $\alpha_{pg'}^g\left(\lambda_p^{g'}\right) > \alpha_p^{g'}$ and gender g person $p^{g'}$ is prefered to the optimal choice under performance pay for gender g'. Since by definition person p^g is preferred to all other gender g persons under performance pay, he is also preferred to person $p^{g'}$ and $\lambda_p^u = \lambda_p^g \Rightarrow \Psi^u = \Psi^g$. However, if $Q_1^g \geq Q_1^{g'}$, $\gamma^g = \gamma^{g'}$, and $\sigma_\varepsilon^g < \sigma_\varepsilon^{g'}$ then λ_p^u could $= \lambda_p^g$ or $\lambda_p^{g'}$. To see this, suppose $\lambda_p^{g'} < 0.5$ (this case happens when $\Psi^{g'} = 0$ and $Q_1^{g'} < 60$) and $Q_1^g = Q_1^{g'}$. In that case $\alpha_{pg'}^g\left(\lambda_p^{g'}\right) < \alpha_p^{g'}$ (since, if we define $\Lambda^k\left(Q\right) = P\left(Y > 60|Q, gender = k\right), \Lambda^{g'}\left(Q\right) > \Lambda^g\left(Q\right) \forall Q < 60$) and, for different levels of P_i, it could be that $\alpha_p^{g'} + P_i\lambda_p^{g'} >$ or $< \alpha_p^g + P_i\lambda_p^g$ which means that either person $p^{g'}$, who is unlikely to qualify but gives high social payments, or person p^g who is likely to qualify but gives relatively low social payments could be chosen under perfomrance pay. In fact, the two outcomes of this specific case are the only two potential outcomes; as a result, when information and quality are worse about g' relative to g we cannot rank person $p^{g'}$ relative to p^g, and so $\Psi^u = \Psi^g$ or $\Psi^u = \lambda_p^{g'} - \lambda_1^g < \Psi^g$. A similar argument reveals that when $Q_1^{g'} > Q_1^g$, we cannot rank persons p^g against $p^{g'}$ in either the case of greater shallowness or worse information about gender g'.

A.2   Competition

In order to directly look at the role of competition in referral decisions, we experimentally varied how salient competition was to CAs. CAs were told the qualification threshold was either (i) determined using an absolute standard (receiving a score greater than 60) or (ii) in relative terms (scoring in the top half of applicants). Table A2 shows that referrals, both men and women, are not statistically less likely to qualify when CAs are directly competing with their referrals to become qualified. While this treatment should not alter perceptions of competition in the post-qualification phase, and is therefore a fairly weak test, it provides suggestive evidence that, on average, competition is unlikely to be driving our main results.

A.3 Appendix Tables

Appendix Table 1: Summary Statistics and Randomization Check

Dependent Variable	Mean and SD: Male	p value of joint test of treatments	N	Mean and SD: Female	p value of joint test of treatments	N
	(1)	(2)	(3)	(4)	(5)	(6)
CA Age	25.52	0.441	445	24.61	0.787	271
	[3.88]			[4.62]		
CA qualified	0.56	0.188	480	0.48	0.390	287
	[0.50]			[0.50]		
CA Overall Test Score (corrected)	61.66	0.373	480	59.98	0.085	287
	[13.59]			[13.22]		
CA Has Previous Survey Experience	0.31	0.410	480	0.26	0.189	288
	[0.46]			[0.44]		
CA Has Tertiary Education	0.69	0.367	480	0.78	0.186	287
	[0.46]			[0.42]		
CA MSCE Math Score	5.65	0.867	419	6.84	0.061	242
	[2.30]			[1.80]		
CA MSCE English Score	5.68	0.651	435	5.75	0.594	256
	[1.49]			[1.41]		
CA Job Comprehension Score	0.80	0.894	480	0.81	0.573	288
	[0.40]			[0.39]		
CA Math Score	0.21	0.245	480	0.18	0.351	288
	[0.10]			[0.09]		
CA Ravens Score	0.61	0.146	480	0.56	0.460	288
	[0.40]			[0.39]		
CA Language Score	0.15	0.302	480	0.14	0.602	288
	[0.03]			[0.03]		
CA Practical Component Z-score	-0.10	0.102	476	0.17	0.101	284
	[1.03]			[0.90]		
CA Computer Score	0.44	0.533	480	0.43	0.523	288
	[0.21]			[0.20]		
CA Feedback Points	25.90	0.037	474	27.92	0.252	284
	[7.28]			[6.31]		

Notes

1 The displayed p value is from the joint test of all the treatment variables and their interactions from a regression of the dependent variable listed at left on indicators for each treatment and CA visit day controls. The regressions are done separately for men and women.

2 All specifications include CA visit day dummies.

Appendix Table 2: Competition incentives in the fixed fee treatments

	CA Qualifies	Referral Qualifies	Referral Qualifies			
	(1)	(2)	(3)	(4)	(5)	(6)
Competitive Treatment	-0.055	0.072	0.052	0.014	0.090	0.227
	(0.062)	(0.069)	(0.121)	(0.086)	(0.095)	(0.165)
Female Treatment			0.094			-0.024
			(0.116)			(0.177)
Either Treatment			0.175			-0.160
			(0.123)			(0.169)
Competitive * Female Treatment			0.007			-0.263
			(0.166)			(0.236)
Competitive * Either Treatment			0.103			-0.142
			(0.176)			(0.236)
Observations	276	232	232	166	133	133
CA Gender	Men	Men	Men	Women	Women	Women

Notes

1 The dependent variable is indicated in the column heading.

2 All specifications include CA visit day dummies.