

Do Job Networks Disadvantage Women?

Evidence from a Recruitment Experiment in Malawi *

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

October 2012

Abstract

Using a field experiment in Malawi where men and women apply for future surveyor positions with a local firm, we find that highly skilled women are systematically disadvantaged through the use of referrals. This happens both because most men recommend other men, and because women refer fewer candidates who qualify for the position. We document that segregated networks do not cause this behavior. We develop a theoretical model of referral choice and exploit random variation in referral contract terms to find that both men's and women's biases result from social incentives rather than expectations of performance. We also document that the screening potential of networks is maximized when men refer men. This paper suggests that the use of social networks in hiring is an additional channel through which women are disadvantaged in the labor market.

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,

*We thank IPA-Malawi field staff and Sam Arenberg for research assistance. We also thank participants at Brown University's Mini-conference on Networks, Northwestern University, SITE, the North America Summer Meetings of the Econometrica Society, and seminar audiences at the Universities of New South Wales, Toronto, Illinois Urbana-Champaign, Houston, and California, Berkeley for helpful comments. All errors are our own.

[†]Northwestern University. Email: l-beaman@northwestern.edu

[‡]Innovations for Poverty Action. Email: nkeleher@poverty-action.org

[§]University of California, Berkeley. Email: jmagruder@berkeley.edu

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 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 disadvantage 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 referral systems to create inequality between groups has been described theoretically (Calvo-Armengol and Jackson, 2004).² Mortensen and Vishwanath (1994) also show theoretically that network-based job information dissemination can disadvantage women, even if men and women are equally productive but men have a higher contact probability. Observational data seems to support the hypothesis that women benefit less from job networks than men. Ioannides and Loury (2004) document stylized facts that women are less likely to be hired through a referral and that unemployed women are less likely than unemployed men to search using family and friends³.

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

¹Additional 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).

²For 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.

³Moreover, 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).

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. Labor market disadvantages may result in female applicants with weaker easily observable characteristics, like job experience, but network screening 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, it remains an open question whether women are made worse off by the use of employee referrals.

We used a competitive recruitment drive conducted by a research organization in Malawi, Innovations for Poverty Action (IPA-Malawi), as an opportunity to test how job referrals affect the recruitment of men and women in an experimental setting. IPA-Malawi historically had struggled to hire female enumerators, and was interested in exploring whether referrals could reveal an otherwise untapped pool of qualified female applicants specifically, and qualified applicants in general⁴. 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 interview, 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

⁴Often times 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.

a man, or that they may refer someone of either gender, and their finder's fee was randomly selected to be a fixed fee of either 1000 or 1500 Malawi Kwacha (MWK) or a performance incentive (a guaranteed 500 MWK with the potential to earn an additional 1300 MWK if the referral attained a certain threshold).

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. 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 14% of men and 17% of women refer qualified women, which compares to 42% of men and 21% of women who refer qualified men. This is consistent with the finding from observational data from a call-center in Fernandez and Sosa (2005)⁵ and is the first experimental evidence that we know of that 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). 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.

⁵In that context, men are the disadvantaged group, who are similarly less likely to receive referrals.

Since 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. They also receive a noisy signal of each network members' 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: (i) for any type of CA, networks of men and women may differ in the magnitude of top social payments ("closest gender"); (ii) the observable quality of friends who provide top social payments ("quality"); (iii) the extent of the tradeoffs between social payments and observable quality ("network shallowness"); and (iv) and the accuracy of the signal of observed quality relative to actual quality ("information").

Because of random variation in the structure of the finder's fee, we are able to observe the characteristics of both potential referrals when CAs are motivated in their choice only by network incentives and when firms provide an additional incentive to find a person who is high ability. This facet allows several tests of the various sources of heterogeneity. First, the model suggests that referrals selected under fixed fee payments should reflect the people closest to CAs, so that both the quality of closest people and any gender difference in top social payments will be reflected in referral choices under fixed fees. 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 give highest social payments to men 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 CA⁶. Overall, social incentives among both men and women's networks make it harder for qualified women to get job opportunities.

Second, we derive that if tradeoffs between social payments and observed quality are higher in one gender, or if information is relatively worse for that gender, we should observe a smaller performance premium. In other words, incentivizing CAs to find the best possible referral will be less effective in the presence of either bad information or high tradeoffs, which may have implications for firm motivation to address any gender disparities which exist for reasons of social preference. Taking this prediction to the data, 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. Further examining the distribution of test scores when men refer women under performance pay suggests that worse information about women is likely to be at least part of the explanation. By contrast, women 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.

Taken together, 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

⁶This 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; however, the quality of women's referrals in the unrestricted treatments are closer to the quality of the female-restricted treatments than the male-restricted treatments.

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.

The paper is organized as follows. The experiment 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 section 5. The empirical evidence for each of the three hypotheses are presented in section 6, followed by a discussion of alternative interpretations in 7 and the conclusion.

2 Experimental Design

2.1 Setting and Overview

Gender-based difference in employment is common in many developing countries (Bell and Reich, 1988). Women in Africa are more likely to be in the informal sector and the proportion of women in the formal employment is less than half that of men. A survey of 14 African countries found that women are more likely to be employed (formally or informally) in agricultural jobs and less likely to hold jobs in the manufacturing and services (Arbache, Kolev, and Filipiak, 2010).

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.

The standard method for recruiting enumerators is to post announcements at community centers, technical schools, and government offices. An initial screening session is open to all applicants with minimum qualifications. 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.

The standard IPA-Malawi screening session consists of submitting a CV to IPA and sitting for a written test. The written test assesses reading comprehension, hand writing, math ability, and computer literacy (via self-assessment). Following the screening session, applicants deemed to be qualified may be invited for a survey-specific training of enumerators.⁷ At the end of the training, job offers are made to a group of individuals deemed to be adequate for work on the survey.

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

⁷These trainings consist of a multiple-day workshop on proper technique and procedures for conducting paper-assisted or computer-assisted personal interviews. Each training is tailored to a specific survey; however, interview techniques for facilitating and documenting interviews is rather standardized. Also, during a training workshop, practical skills are assessed through a field pilot of the given survey.

IPA enumerators, the dates and times of the recruitment sessions, and a solicitation to bring a CV and certificate of secondary school completion (MSCE). 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. Participants were given one of two distinct written tests. The two tests were distributed at random to limit cheating. Each test consisted of several math problems, ravens matrices, English skills assessment, job comprehension component, and a computer skills assessment. For richer survey data, our screening session integrated the practical test.

For the practical test, the participant played the role of the enumerator for a computer assisted personal interview.⁸ 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.

⁸All 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. (The CAPI software used was Blaise. To our knowledge, no participant had used this particular software prior to the recruitment session.) Once finished with the self-administered CAPI training, participants moved to 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. All original participant letters described a specific set of instructions about the referral process. We randomly varied the content of the letters.

Each letter included an instruction about the gender requirement 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 either gender. The referral needed to be someone who had not worked for or been tested by IPA in the past. 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 would need to complete the same written and practical assessments as done by the CA.

Conventional applicants were 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 all evaluators. All conventional applicants 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.

Conventional applicants were asked to complete an anonymous questionnaire as an assessment of their referral's quality and whether or not they shared any of their payment with the referral. In addition, CAs were contacted by an IPA staff member to ask how the conventional applicant identified the referral and how the payment was used.

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. 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.

Appendix Table A1 displays summary statistics for the sample of CAs, for men and women separately. It also shows that the randomization lead to balance along most characteristics. The p value for the joint test of all the treatment variables, and their interactions, 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 a 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), 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: women refer women at approximately the rate by which women apply themselves (43% of the time), while men refer women only 23% of the time when given the choice. 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 women⁹. 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. As it turns out, 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 this difference is held up by each gender of CA, regardless of which gender they choose to refer: male CAs very rarely refer women, but the women who they do refer are fairly likely to qualify, while female CAs often refer women but rarely refer anyone who IPA would actually hire. 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 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. Overall, we conclude that the use of referral systems strongly disadvantages qualified women in this

⁹In both cases, the sample is restricted to CAs who have the choice of which gender to refer.

context.

4 Are Men Connected to Women (and Women Connected to Men)?

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. 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. Since 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 statistically significantly less likely to make a reference when they are given performance pay than when they are given fixed fees if they are required to refer either men or

women. 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 only know a man who is worth referring under performance pay, but also 15% who only know 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 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 language 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

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

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 $Q_1^g > Q_1^{g'}$ ¹⁰

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

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

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 Behavior under Fixed Fees

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 $(F_i, 0)$. Examining the optimization condition leads to

¹⁰This 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' .

the following proposition.

Proposition 1 *Under fixed fee contracts of the form $(F_i, 0)$, CAs always refer the closest person of the eligible gender, friend 1. Differences in α_1^g can lead to different return rates between genders, but differences in Q_1^g , γ^g , and σ_ε^g will not result in different return rates between genders.*

Proposition 1 reveals that if we examine referrals who were recruited under fixed fees when restricted to a particular gender, then we are directly examining friend 1 of the appropriate gender. As a result, the quality of referrals recruited under fixed fee treatments restricted to referring a particular gender can be taken as an estimate of Q_1^g , and the fraction of CAs who refer men when allowed a choice of genders can be interpreted as the fraction of CAs for whom $\alpha_1^1 > \alpha_1^2$. If CAs systematically refer men under fixed fee treatments, proposition 1 suggests that the only interpretation of this fact is that men are systematically closer to CAs. Our empirical analysis in section 6 therefore begins by examining fixed fee treatments to estimate both Q_1^1, Q_1^2 and the relationship between α_1^1 and α_1^2 for male and female CAs and determine whether gender differences in network quality or in closeness could affect the opportunities available to women.

5.2 Behavior under Performance Payments

Employers should be interested in what the potential of networks are, given proper (financial) motivation. To describe how network heterogeneity affects incentivized referrals, we commence by defining a performance premium in this context.

Definition 6 *There is a performance premium if CA i selects referral p with social incentive α_p and observed quality q_p , $p > 1$ when $P_i > 0$ ¹¹.*

¹¹This definition focuses on the “extensive margin” of the performance premium as, by definition, whenever

This definition of a performance premium is essentially whether the CA selects an internal solution or a corner solution - the individual who gives the highest social payment - to the referral choice equation.

Proposition 2 *CAs will select referrals which feature a performance premium in gender g for a larger set of potential performance incentives if*

A): the gender g network is less shallow

B): CAs have more information about gender g and the index partner has sufficiently low expected performance¹²

C): CAs have a higher (lower) quality network of gender g than gender g' and $|60 - Q_1^g| < (>) |60 - Q_1^g - \gamma^g (a_1^g - \alpha_p^g)|$ (for person p who solves equation 2)

Corollary 1 *The set of performance incentives which will lead to selecting a referral with a performance premium is unaffected by differences in α_1^g .*

In other words, as networks become more shallow, the return of someone other than the closest friend declines relative to the closest friend. If index partners have sufficiently low expected performance, then the opposite is true for better information: in that case, the index partner is not expected to qualify and so better information increases the chance that the CA is correct about the index partner's (expected) failure. In terms of network quality, the non-linearity of the normal distribution causes the distinction in effects: as quality increases, the returns to recruiting the index partner increases relative to other potential partners if the index partner is at a point of higher density in the normal distribution, or if (s)he is closer to

the person chooses a referral other than the closest (and lowest ability) friend, the performance premium will rise. Marginal effects on the magnitude of the performance premium are theoretically ambiguous with respect to most key parameters due to the non-linearity of the normal distribution.

¹²See Lemma 1 in the appendix for the specific condition on expected performance. Though strictly, the condition is weaker than in Lemma 1: it must be that $\frac{60-Q_1^g}{\sigma_\varepsilon^2} \phi \frac{60-Q^g}{\sigma_\varepsilon} > \frac{60-Q^g-\gamma^g \alpha_1-\alpha_j}{\sigma_\varepsilon^2} \phi \frac{60-Q^g-\gamma^g \alpha_1-\alpha_j}{\sigma_\varepsilon}$.

Table 3: Estimated directions of network differences

Parameter	$\pi_i(F_i, 0)$	$\pi_i(F'_i, P_i)$		$E[Y_i P_i > 0]$
		$E[Y_i P_i > 0] = Y_i$	$E[Y_i P_i > 0] > Y_i$	
α_1^g	↑	↑	↑	0
Q_1^g	0	↑	↑	Ambiguous
γ^g	0	0	↓	↓
σ_ε^g	0	↑	↓	↓

the mean. In this case, the added quality of the index partner increases the probability that the index partner qualifies, making referring someone else less desirable.

Proposition 3 *If there is no Performance Premium for gender g , then the overall return under performance pay is: increasing in α_1^g ; increasing in Q_1^g ; increasing in σ_ε^g ; and unrelated to γ^g .*

Proposition 4 *If there is a Performance Premium for gender g , then the overall return under performance pay is increasing in α_1^g ; increasing in Q_1^g ; decreasing in σ_ε^g ; and decreasing in γ^g ¹³.*

Table 3 summarizes Propositions 1, 2, 3, and 4 and documents that heterogeneity in any of the underlying parameters could put men at an advantage under performance pay. In particular, if one gender (say, men) is closer than the other, then men should be referred more under performance pay just as they would be in fixed fees, all else equal. Unlike in fixed fee treatments, if men (or women) are higher quality, then they also will be referred preferentially under performance pay. This result holds whether CAs intend to return with the index partner or intend to choose a partner with a performance premium, as either will be higher quality (with higher expected return) when the network's quality shifts upwards. Therefore, if we conclude (from our tests on fixed fee CAs) that gender-specific networks are different in closeness or quality, we should expect those results to translate to gender preference under performance pay. Shallowness and bad information affect gender preference in ambiguous ways

¹³Propositions 3 and 4 make use of lemma 1, contained in the proofs appendix, which states that under the functional forms contained here, the expected performance of partner p is greater than 60.

under performance pay. If one gender network is shallower than the other, then it affects gender preference under performance pay only if a referral other than the index partner was going to be selected. Bad information can go either way: if the optimal referral under performance pay was the (low quality) index partner¹⁴, then bad information should make that gender more likely to be selected in unrestricted treatments as the chance that the index partner actually qualifies is increased (conditional on the signal). On the other hand, if the CA intended to select someone who he expects to qualify, worse information about that gender can only hurt his payoff to selecting that gender. Shallowness and bad information are similar, however, in how they affect the CA's choice of optimal partner: if one gender-specific network is more shallow, or if information is worse, then the CA will unambiguously be more likely to select the index partner rather than a partner who he expects to perform well on the recruitment exercise. This leads to a test: if there is more of a performance premium in one gender than the other, and if there is no difference in quality of the index partner, then we can conclude that the gender with a larger performance premium has a network which is either less shallow or better informed. Given that employers are no doubt interested in using referral systems to screen, this may suggest that employers would want to encourage referrals of the less shallow or better informed network.

5.3 Empirical Tests

The model above outlines how differences in underlying heterogeneity may lead to different referral choices. This section summarizes how these differences may disadvantage one gender vis-a-vis the other when referrals are given the option to refer either gender. As our focus in this paper is on understanding potential mechanisms which could disadvantage women, discussion

¹⁴Recall that at least for some individuals, the expected probability of qualification for the index partner was less than 5/13

in this section will presume that women are the disadvantaged gender under any hypothesized mechanism though of course algebraically either gender could be the disadvantaged g' . Under “either” treatments, CAs will exhibit a preference to refer men over women if the overall returns to referring men are higher. Table 3 highlights four potential reasons that returns may be higher under different contracts: men could be closer than women to CAs; men could be higher quality than women; men could be in less shallow networks than women; and CAs could have different information about men than about women (which in principle could have ambiguous sign as to the advantage). Ultimately, we present 4 tests: first, using the choice of gender under fixed fee, unrestricted treatments, we infer which gender (male or female) is closer to CAs. If one gender is systematically closer, we should also observe that gender being chosen more under performance pay. Second, we estimate the quality of each gender-specific network for CAs by examining the quality of references under fixed fees. Third, we test for the presence of a performance premium. If there is no performance premium, it will be consistent with bad information or shallowness. To separate these, we compare distributions of gender-specific referral scores under fixed fees and performance pay to evaluate whether very low quality referrals are less likely to take place under performance pay, which would be consistent with well-informed shallowness, or not, which would be consistent with bad information being at least part of the story.

6 Results

6.1 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 the social incentives which live over the network, at least for men: the closest people for men tend 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. Given that the two sources of women's disadvantage that we highlighted earlier are both present when we introduce no direct financial incentives to make a qualified referral, we 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.

6.2 Are there quality differences across networks?

A second reason why women may be disadvantaged by social networks would be that men (or women) did not have high quality women in their network,¹⁵ or that men or women are very distant from high quality women. The model provides a clear formulation of network quality, being defined as the expected quality of the person who an CA is closest to. Figure 4 presents kernel densities of the ability of men's male and female networks 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, that differences of the quality of women in men's networks as opposed to men in men's networks does not contribute to men's preference for referring 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 0.42 of a standard deviation below the CA mean, on average, while men who are closest to women perform 0.08 standard deviations below the CA mean, which is higher than the men (or women) who are closest to men though not statistically different. Our results indicate that women are closest to women who are particularly low ability.

These observations are about the quality of closest people, rather than the full distri-

¹⁵The purest test for the highest potential quality man or woman that an CA could bring in would put extremely large performance premium attached to a high threshold. In the interests both of maintaining a somewhat normal employment scenario, we were unable to include extremely high stakes treatments. Of course, those treatments also seem unlikely to be reflective of the referral incentives which exist in other employment settings.

bution of network quality. It is also possible that quality is different across the range of the network, which could contribute to women’s disadvantage. We therefore now look at referrals recruited under performance pay contracts.

6.3 Are there gender differences in information or shallowness?

Results from fixed fees are informative about the characteristics of closest people within networks, but cannot speak to other characteristics of gender-specific networks. What is potentially most interesting to employers is not the ability of the closest individuals to men and women, but rather the ability of people that men and women can identify, when properly motivated. In the theoretical model, we identified that CAs would not change their referral choices if either information was too poor, or if networks of a particular gender were too shallow. To test whether men and women are willing to identify high quality men and women, 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 4 presents the results of this analysis. Male CAs experience a substantial performance premium when referring men: 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. However, they do not experience any performance premium when restricted to refer women. Female CAs show positive coefficients on performance pay for all gender treatments, but they are never significant and always small in magnitude. We therefore conclude that male CAs have useful information for employers about men, and that tradeoffs are not too high to

prohibit using it.

However, men do not seem to choose female referrals with a performance premium, which could be consistent with bad information about women or with shallow networks of women. These two explanations are fairly similar: in one, men would observe a similar signal for all women that was quite accurate; in the other, men would observe different quality signals about women but know to place little weight on these noisy signals. One way to differentiate these two explanations is to examine whether men are less likely to make really poor female referrals under performance pay: with high information but shallow networks, we would expect men who know their female networks are low quality to be less likely to make a referral under performance pay, while they would not feel the same pressure under fixed fees. In contrast, if men have low information about the ability of women in their networks, then we would expect them to be similarly likely to make a low ability female referral under either choice of contract. For male CAs who were asked to refer women, their referrals are similarly likely to score more than a standard deviation below the CA mean under performance pay and fixed fee treatments (20% versus 17% respectively). These male CAs are also at least as likely to refer women who score more than 1/2 of a standard deviation below the CA mean under performance pay than under fixed fee treatments (38% versus 32%, though neither difference between contract groups is statistically significant). This suggests that, at least, male CAs are not declining to make a referral under performance pay as one might expect if they had highly accurate but homogeneous signals of low ability. This does not rule out, of course, that there are also more stark tradeoffs among women network members than among men.

Table 5 explores further the differences 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 5 finds that men referred be 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, in contrast, do not demonstrate an overall performance premium under any treatments. However, they are nonetheless changing their referral choices. Tables 6 again disaggregates referral performance by component, for women CAs, 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 than the men being referred under fixed fees, while the women with better language skills perform weakly worse on a number of characteristics¹⁶. 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 face networks which are too shallow overall to find women or men who are able to qualify.

Finally, a striking result from table 4 is that performance premia are in fact lower for men under unrestricted treatments than under male restricted treatments. Returning to the model, we observe that there is no guarantee that allowing the choice of either genders will lead

¹⁶The 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.

to a larger performance premium, which leads to the following proposition.

Proposition 5 *Allowing the choice of either gender under performance pay leads to an unambiguously higher return rate. However, the effect of allowing either gender on referral performance is ambiguous if there are different performance premia for genders g and g' .*

Based on our definition of performance premia, if CAs refer someone for both genders who is not the index partner when they receive performance pay, then there will be a performance premium for all CAs under both restricted treatments and under the unrestricted either gender treatment. If CAs always refer the index partner when they receive performance pay under either restricted treatment, then there will be no performance premia under either either of the restricted treatments or the unrestricted either gender treatment. However, when some CAs optimize for one gender with a performance premium, and optimize for the other gender without a performance premium, the same CAs may end up referring the performance premium candidate under the “either” option while others refer the index candidate. To better understand when we expect a performance premium, consider an CA who would opt to refer a person with a performance premium in gender g but not in gender g' . In this event, the CA will prefer to refer his premium candidate, candidate p (of gender g) over the close candidate of gender g' if the differential probability of qualification is large enough relative to the difference in social incentives, or if

$$P_i \left(\Phi \left(\frac{60 - Q_1^{g'}}{\sigma_\varepsilon^{g'}} \right) - \Phi \left(\frac{60 - Q_1^g - \gamma^g (\alpha_1^g - \alpha_p^g)}{\sigma_\varepsilon^g} \right) \right) > \alpha_1^{g'} - \alpha_p^g$$

Obviously, all of the dimensions of heterogeneity we have discussed before enter here. What is interesting for prediction is to consider how different characteristics which may be related to performance premia interact with which candidate gets referred. More specifically, if

we increase Q_1 , α_p or decrease γ for the gender with the performance premium, we anticipate that CAs with the option of referring either gender will be more likely to refer the performance premium candidate. The same is true if we decrease σ_ε for either gender, or if we decrease Q_1 or α_1 for the candidate with no performance premium. In contrast, if the opposite is true, so that we increase the quality (Q_1), social payments (α_1) of the low-performance premium gender, or if information gets worse for either gender, then we anticipate that allowing the opportunity to refer either gender would be associated with a lower performance premium than restricting referrals to gender g . Given the preceding evidence that men's information about women appears poor, the lack of a performance premium observed when men have unrestricted choices would be quite consistent with men having some women who are close and who might qualify (since information is poor, any women would have a strong chance of qualifying) and sometimes optimizing by choosing these women over the men who they know will perform well but who give them low social payments.

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 consistent with presented 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. Interestingly, the gender differences here still require differences in information: clearly CAs have good information if they attrit because they know that their optimal referral is low quality, and so they would have to have poor information about women's capabilities if they

do not attrit in the same way when required to refer women¹⁷. Figure 6 suggests that among men, there is assortative matching in ability between the CA and their referral under fixed fees. This presents an additional concern for interpreting the results in Table 4: if performance pay reduces the likelihood of high ability CAs making referrals, and ability is correlated within networks, then Table 4 could be biased. However, figure 6 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 / shallowness¹⁸. Figure 7 shows analogously that there is little evidence of female CAs responding to the performance pay incentive at any point in the CA performance distribution.

7 Alternative Interpretations

7.1 Quality expectations and social payments

While Table 1 established that social payments contribute to women’s disadvantage, could expectations of quality be a contributing factor? In particular, CAs may receive higher social payments when their referrals qualify than when they do not, perhaps because referrals who actually get the job give higher transfers to the references who helped them than referrals who only get the opportunity to apply or because they believe that they will receive a reputational premium from IPA for making a highly skilled referral. Either explanation would create a relationship between the CA’s expectations of quality of men versus women and social payments.

This potential is consistent with our model, as we made no assumptions about the joint distri-

¹⁷This explanation is, however, inconsistent with the fact that male CAs decline to make a referral at the same rate whether they are required to refer men or women under performance pay. Together with the evidence on poor information about women, it suggests that male CAs are indeed making different optimal male referrals under performance pay.

¹⁸This 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.

bution of α_j and Q_j in the underlying network, relying instead on characteristics of the subset of non-dominated referrals. It would, of course, change our interpretation of social payments.

In the model, expectations of referral quality are unbiased as ε_j^g is mean zero for both men and women. Figure 4 shows that men's fixed fee referrals of men are not higher ability than their referrals of women, so that unbiased expectations about referral quality should not deliver the result that men systematically refer men. Second, women's (low quality) fixed fee referrals are also not the highest quality people they know, since they know high quality men (documented in figure 5). Therefore expectations about referral quality, as long as they are unbiased, are not leading women to make low quality referrals under fixed fees. If social payments were a function of qualification, with unbiased expectations we would expect to see women systematically referring men, and men referring 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 ($E_i [Q_1^g + \varepsilon_1^g]$) would have similar implications in the model to actual differences in quality (Q_1^g). Given that we found that women's disadvantage is preserved when CAs optimize only social incentives, biased expectations of referral quality could explain our results if social incentives are positively related to referral quality in the overall network. In that case, biased expectations of quality could lead to men referring men even though their closest women are, in reality, at least as skilled as their closest men. Bias could also lead to women referring low ability women despite being close to high ability men.

The presence of performance pay suggests an easy test of whether biased expectations could contribute to our result. CA incentives to refer someone who they believe to be higher quality are unambiguously higher under performance pay than under fixed fees, as they receive social incentives (and any relationship those have with referral quality) in both types of contract

but receive additional performance pay indexed to referral performance in performance pay contracts. Therefore we can point to several pieces of evidence that biased expectations are not driving the behavior that leads to few qualified women getting referred. First, recall that Table 4 showed that men did not refer the highest ability men in their network under fixed fees, since the performance pay lead to more qualified male referrals. Therefore it can not be the case that social payments are completely indexed to qualification, or that CAs do not respond to our performance pay: 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 unless they also know women who are much better (by more than their bias) than the men they know. 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 still small fraction of women referred), which is inconsistent with women's disadvantage resulting from expectations over quality. 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-

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 minimal role in explaining why so few qualified women get recruited through referrals.

7.2 Competition

An additional explanation of our findings is that women in particular are more averse to competition than men, thereby referring less qualified applicants. 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. In contrast, if women's competition aversion explained our results, it would have to be that women perceive more of a competitive threat from other women than from men. In reality, given the firm wants to increase the number of women enumerators, competition - once they have become qualified - is actually weaker for female applicants. Moreover, the large number of applicants in the recruitment drive means that any one additional referral by and large does not impact the CA's chances of becoming qualified or getting a job conditional on qualification. Nevertheless, if women CAs perceive there to be a limited number of positions for women, they may choose to either forgo the finder's fee or refer a woman who is unlikely to qualify. We do not observe the former but the latter is consistent with the results presented thus far, if the perceived threat from other women is sufficiently high to overcome the performance premium offered. 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.

In order to shed light on this potential mechanism, 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.

Moreover, women CAs who have the biggest incentive to fear competition do not make systematically the worst referrals. Figure 3 shows that when female CAs are at the margin of qualification, around 60 points, they make referrals who are *more* likely to qualify. This is additional evidence that competition from all other applicants is not driving the results. However, an interpretation which remains is that women make references with a goal of avoiding competition from other women. Figure 2 suggests that the marginal female CAs switched from referring women to referring men, which explains some of why their referrals perform more strongly (given that women's male referrals are higher ability than their female referrals, across the distribution of women's ability, as visible in Figure 7). Figure A4 also shows that women CAs close to the qualification threshold of 60 who must refer women choose referrals who are less likely to qualify. This occurs in both fixed and performance pay treatments, so our performance incentive was not large enough to overcome this perceived competitive threat. These patterns are consistent with the hypothesis that women view fellow women as competition, making it even more difficult for qualified women to obtain opportunities through social networks.

8 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, so much so that two of the stylized facts presented in Ioannides and Loury (2004) survey emphasize gender differences in network usage. Using an experiment designed around a recruitment drive for real-world jobs, we provide the first experimental 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 exhibit a preference for referring 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 Europe¹⁹, 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 lead to a disadvantage for women. First, we confirm that women and men are present in each other's networks. Thus, while we cannot describe the relative gender homophily of networks, we can conclude that women are not being left out simply because they are fully absent from men's networks. Using variation in contract structure, instead, we find that several network characteristics work against women. First, men's networks are characterized by other men being systematically the closest members, who give the highest social incentives for a referral. However, the men and women who are closest to men tend to be high ability. This contrasts with women's networks, which are not characterized by a gender preference in

¹⁹Most 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.

social incentives but are characterized by low quality people and especially low quality women providing the highest social incentives. This suggests that unless employers design referral contracts to contradict these incentives (at additional cost), network incentives put women at a disadvantage: 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 ultimately qualified.

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. Men also exhibited a preference for performance premium when we required them to refer men, but the preference eroded when they were allowed to refer women as well. Our framework allowed us to say little about why job referral systems are used, but a frequent hypothesis is that they are used to help employers in screening (Montgomery, 1991; Beaman and Magruder, 2012). If employers use referral systems because they hope to screen, then this result suggests that employers may want to emphasize men referring men when hiring through networks. In other words, if screening is a productive use of networks, then this result suggests that a profit-maximizing employer may be encouraged to discriminate when using referrals to hire new workers. This result suggests that in order to prevent discrimination against women, careful hiring procedures may need to be followed. One such procedure which these results suggest may work well is a quota-based referral system, where people making references are required to make references of either gender - given that male CAs can both screen men and tend to be close to high quality women, these estimates suggest that a quota based system may be effective at identifying high quality workers 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. *Forthcoming, The American Economic Review*.
- 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: <i>p</i> value
<u>A. CA Characteristics</u>				
Fraction of CAs	100%	62%	38%	
CA is qualified	53%	56%	48%	0.047
N	767	480	287	
<u>B. CA Characteristics: Made Referral, Either Gender Treatments</u>				
Fraction of CAs	100%	61%	39%	
CA is qualified	57%	62%	49%	0.061
N	217	130	87	
<u>C. Referral Characteristics: Made Referral, Either Gender Treatments</u>				
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	
<u>D. Referral Characteristics: Made Referral, Either Gender, Fixed Fee Treatments</u>				
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	
<u>E. Referral Characteristics: Made Referral, Either Gender, Perf Treatments</u>				
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 (0.038)	-0.055 (0.054)	-0.004 (0.050)	-0.042 (0.074)
Either Gender Treatment	0.014 (0.040)	0.017 (0.055)	-0.052 (0.052)	-0.024 (0.071)
Performance Pay			-0.148 (0.056)	*** (0.080)
Perf Pay * Female Treatment			0.004 (0.076)	-0.013 (0.111)
Perf Pay * Either Treatment			0.152 (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 4: Referral Performance

	Referral Qualifies			
	(1)	(2)	(3)	(4)
Female Referral Treatment	-0.030 (0.062)	-0.190 (0.083)	** 0.068 (0.081)	-0.181 (0.113)
Either Gender Treatment	0.071 (0.066)	-0.231 (0.082)	*** 0.227 (0.084)	*** -0.242 (0.107)
Performance Pay			0.267 (0.093)	*** 0.021 (0.122)
Perf Pay * Female Treatment			-0.248 (0.127)	* -0.022 (0.171)
Perf Pay * Either Treatment			-0.383 (0.132)	*** 0.032 (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 5: Screening of Male CAs on Different Characteristics

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

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

Figure 1: CA Ability by Gender

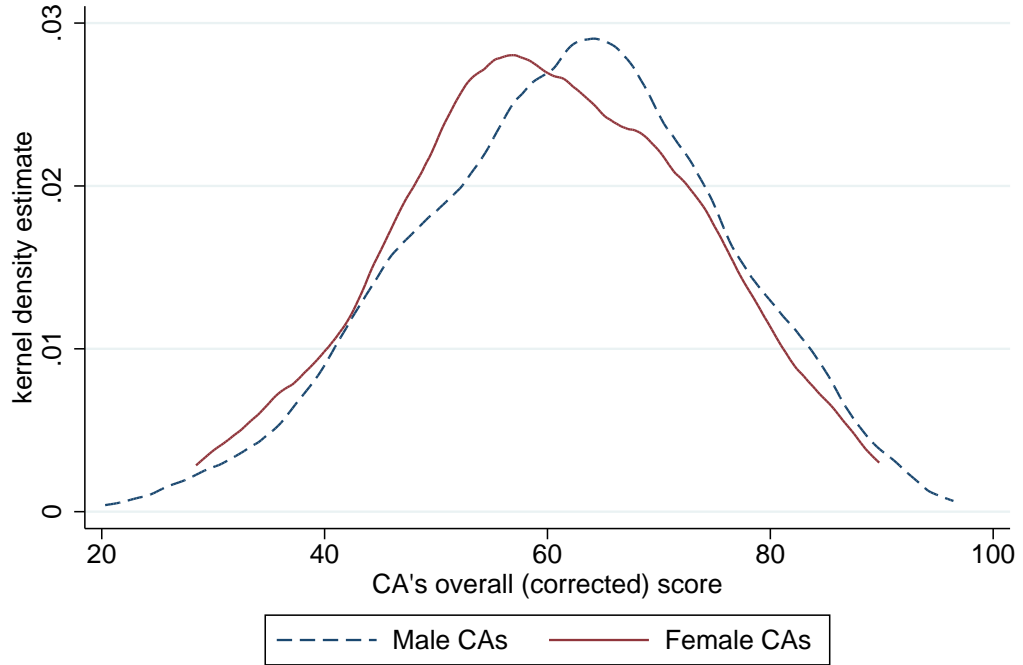


Figure 2: Gender choice in referrals, by CA performance




Figure 3: Referral qualification rate, by CA performance




Figure 4: Men's Fixed Fee Referrals




Figure 5: Women's Fixed Fee Referrals




Figure 6: Referral Qualifies , by Male CA performance

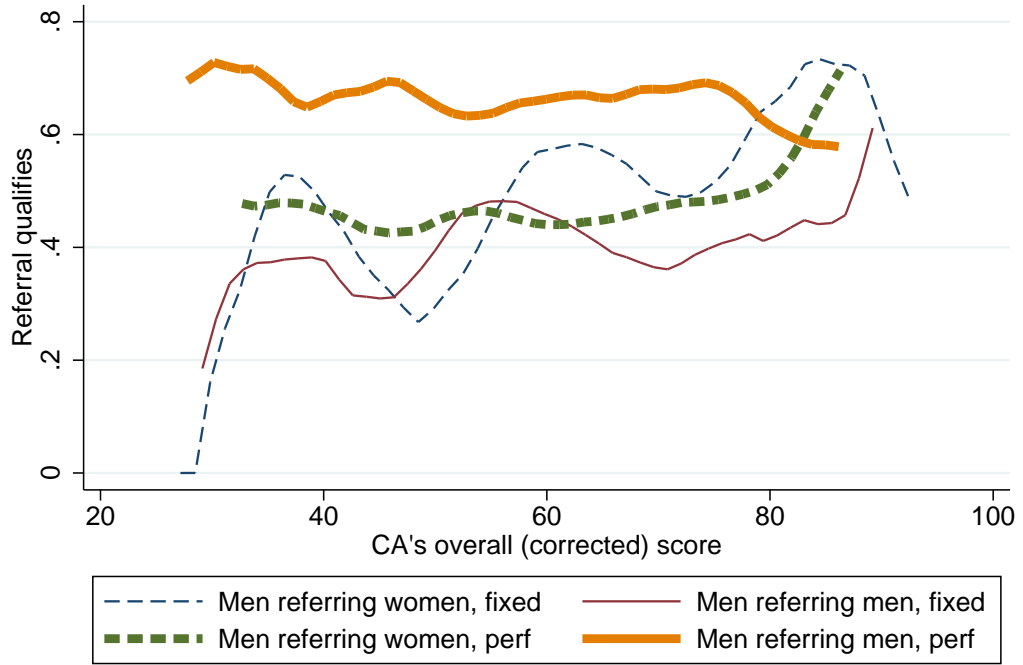
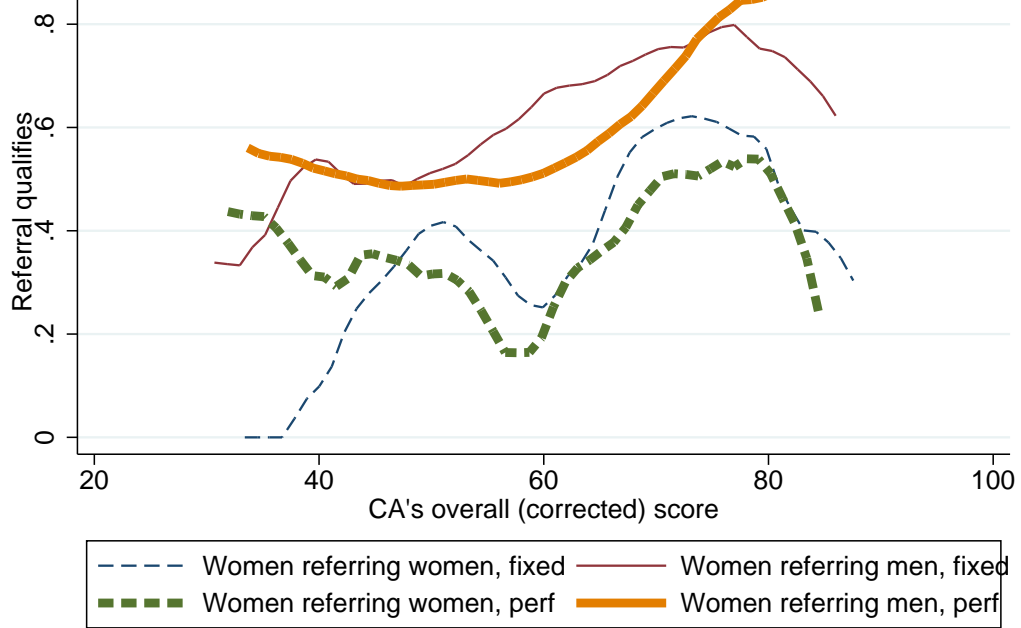


Figure 7: Referral Qualifies , by Female CA performance



A Appendix

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

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