Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi *

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

October 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. The main reason is that men systematically refer few women. We show this is not because there are too few women who are qualified for the job. Instead we show that factors which are not related to women's qualifications but are instead due to the social environment – such as men having worse information about women's abilities and receiving more social benefits from referring men – play a role. Firms cannot just rely on their female employees either since in this context, at least, women referred lower quality candidates.

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

*We thank IPA-Malawi field staff and Sam Arenberg for research assistance. We also thank participants at numerous seminars and conference audiences 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

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). The potential of social network-based hiring to create inequality between groups has been described theoretically (Calvo-Armengol and Jackson, 2004), but there is limited empirical evidence.[2] However, stylized facts in the literature suggest that there may be gender differences in referral behavior: Ioannides and Loury (2004) document 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; Burks et al., 2013; Hensvik and Skans, 2013), 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 may succeed in identifying the women who have strong hard-to-observe but productive characteristics. 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

[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).

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 main dimensions: candidates were either told that they may refer a woman, that they may refer a man, or that they may refer anyone; second, 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 performed above the median qualified for future positions and were informed that they would be called as positions open.

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, 77% of men referred other men.

The random variation in the structure of the finder's fee and the requested gender of the candidate enables us to look at three main explanations of why men overwhelmingly refer other men.[6] Our motivation is to disentangle factors which are useful for the firm - such as the candidate's true underlying quality - from those which are not useful for the firm and stem from the social environment the candidates are embedded in. First, women may be scarce in men's networks, especially because men are more likely to complete secondary school than women in

[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.

[6]We focus on male CAs both since their behavior leads to fewer women being referred and since most firms will start off male-dominated (as is this case) and therefore the behavior of male employees will have the largest impact on recruitment.

Malawi (27% vs 16% in urban Malawi (National Statistical Office of Malawi, 2009)).[7] Second, men may be responding to incentives (either explicit or implicit) provided by the firm to refer good candidates, who may be men. Finally, characteristics of the social environment, such as altruism, expected future reciprocity, or more accurate information about job-specific skills may lead men to prefer other men.

This disadvantage for women does not appear to be due to men not knowing women or because women are particularly scarce in men's networks: men make referrals at identical rates when required to refer either women or men. We use a theoretical model to show that this could not occur if the number of women in men's networks was very small compared to the number of men unless there were also other important differences between men and women in men's networks. Moreover, we show that performance pay leads to a lower referral rate when men have to refer other men but not if they have the choice of referring a man or a woman. Therefore there are some men who only know sufficiently good women to make a referral when offered performance pay.

We show that at least part of the reason women are infrequently referred by men has to do with the social environment, and not women's qualifications. Our test focuses on fixed fee treatments, where incentives to find a high ability referral are the weakest. Men overwhelmingly refer other men in fixed fee treatments. But when we ask men explicitly to refer women, the referral's probability of qualifying is the same - if anything slightly higher - as when the male CAs were asked to refer only men. Therefore, women's disadvantage in fixed fee settings does not stem from actual quality differences. Our model highlights that there are at least two social considerations which could push men to refer other men in fixed fee treatments. First, men may receive higher social benefits, as in Bandiera et al. (2009); Beaman and Magruder (2012); Prendergast and Topel (1996), from referring other men. These could include: higher future payments in the risk sharing arrangement, more utility from referring a man due to altruism or a direct kickback. They could also pay lower search costs from referring men if there are societal norms restricting men's interactions with women. While these factors are quite diverse, they are all things that do not affect firm profits and are simply imposing restrictions on the firm's ability to use its employees to get good quality candidates. Second, even though there is no direct incentive to refer someone good, the CA may internalize the firm's problem if they

[7] These figures are among those 10 years and older according to the 2008 census, with a smaller gap among the younger generations.

believe there is a reputational benefit (cost) of referring someone good (bad). If reputational incentives are concave, CAs are less likely to refer someone whose quality signal is noisy. The social environment may be such that male CAs receive a noisier signal about women: gender homophily in networks is widely documented around the world (McPherson et al., 2001) and could mean men have fewer opportunities to learn about women's abilities.

While we provide evidence that social considerations put women at a disadvantage, are there also productive reasons women may be left out of network-based job recruitment? Yes, but not all contracts will allow the firm to capitalize on them. When male CAs can refer either men or women, we see no more men being referred under performance pay than under fixed fees. We also see no difference in the quality of the candidates referred. In this context, therefore, increasing the incentives for CAs to search for a high quality candidate did not make women worse off.

However, our experiment provides some evidence cautioning that contracts which provide larger returns to candidate quality could put women further at a disadvantage.[8] Among male CAs who are asked to refer men, we find that those in the performance pay treatment refer better qualified candidates than CAs who were offered a fixed fee. This confirms that CAs know which of their male friends represent stronger candidates, and that there are important social benefits which cause them to choose men who they know are less likely to succeed under fixed fees. However, we do not find this response when male CAs must refer women. Firms therefore have an incentive to allow men to continue to refer only other men (and even encourage it) since the best candidates come from men referring other men with a performance bonus. Since men do refer higher ability men when given performance pay, it may be surprising that men do not make higher ability referrals when they can choose referrals of either gender. The model suggests an explanation for this: some CAs may prefer a referral who has a noisy quality signal (a woman) who gives higher social benefits over a high quality man who gives lower social benefits. We demonstrate that this can happen when men have worse information about women than men in their network, and that the net effect of this tradeoff depends on how much quality is rewarded from the firm. We conclude therefore that larger performance bonuses could lead male CAs to refer even fewer women.

Can a firm rely on their female employees to offset men's behavior? Our experiment indicates that using women to make referrals would not be effective in this context because

[8] Our experiment provides a moderate performance incentive, about a day's worth of income.

women systematically referred people who were less likely to qualify. While women refer other women at about the same rate as women apply through the traditional method, a female CA is 18 percentage points less likely to refer someone who qualifies than a male candidate. Since men are systematically referring men, and women are unlikely to refer someone who qualifies, the net result is that few qualified women get referred to the firm. We also find that women CAs do not respond to the performance pay incentive, indicating that either women have little information about people (men and women) in their social network or that the incentive pay was too small to induce a change in the CAs' behavior. Either way the results caution that firms at a minimum may find it more expensive to find good female candidates by eliciting referrals from their female employees than by tacitly allowing their male employees to refer high ability male candidates.

The experiment provides clean evidence that women fare worse than men when IPA-Malawi uses referrals to make hires. As with any field experiment, there is a risk that results would not generalize to other contexts. In this case, however, while there is an absence of experimental evidence on global hiring practices, there is documentation of consistent trends in observational data in contexts much different than Malawi. For example, Lalanne and Seabright (2011) find that women executives in the U.S. and Europe don't leverage their contacts into higher salaries as well as their male counterparts. Loury (2006) using the NLSY found that male workers referred by women get lower on average wages than those who applied through formal channels. Seabright (2012) even suggests that women are more likely to invest in strong ties rather than weak ties, which could hurt them in labor markets which rely on contacts as in Granovetter (1973)'s classic work. The results are also consistent with the finding from observational data from a call-center in Fernandez and Sosa (2005)[9] 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). While the internal validity of our estimates cannot directly speak to the potential internal validity of these other studies, they do allow us to conclude that the apparent disadvantage of women in referral-based hires is a causal relationship in Malawi, and may well be similarly causal in these other contexts.

The paper is organized as follows. The experimental design and data are described in section 2. The main results are discussed in section 3. Section 4 develops a model of referral

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

choice which suggests predictions for how scarcity, the quality distribution of men and women, and the social environment may manifest in our experiment, and these predictions are brought to the data on referrals of male CAs in section 5. Section 6 describes women CAs' referral choices, and finally we conclude.

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 (Arbache et al., 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% 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).[10] 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

[10] 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.

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.[11] Participants were given one of two distinct written tests[12]. 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.

For the practical test, the participant played the role of the enumerator for a computer assisted personal interview.[13] 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.

[11]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.

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

[13]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.

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. The letter also said that the referral should be highly qualified for the enumerator position and gave a suggestive guide about what this would entail. Namely, the letter 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 three to four 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.[14]

[14] To maintain a quick turn-around in notifying applicants of qualifying, real-time test-scoring and data entry

2.4 Internal Validity and CA Characteristics

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, 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 at the 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.

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.

3 Are Women Disadvantaged?

Figure 2 documents 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 (Table 1 shows this difference is 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 approximately the rate by which women apply themselves through the traditional method (43% of the time), while men refer women only 23% of the time. The difference between male and female CAs

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 significant at the 1% level, as shown in column (4) of Table 1. Moreover, these differences persist across the range of CA performance: Figure 3 presents local polynomial regressions of the gender choice of referral on CA overall test score, disaggregated by men and women.[15] CA men are less likely to refer women than CA women across the distribution, with particularly large differences at the top and bottom of the distribution of CA test scores. We will discuss in section 6 whether women CAs can be relied upon to offset male CAs' preference for referring other men. Overall, we conclude that the use of referral systems strongly disadvantages women in this context.

4 Theory

In this section, we develop a model of referral choice to investigate which characteristics of CA behavior may lead to women's disadvantage. CAs each have a network of N_M men and N_F women. These men and women each have three characteristics: an actual quality Y; a noisy signal of that quality that the CA observes Q, where $Y = Q + \varepsilon$ and ε is distributed $N\left(0, \sigma_{\varepsilon}^{g}\right)$, and an idiosyncratic social benefit α, which may be negative or positive and can be interpreted as the cost to CA i of bringing that person in or the reward that that person would give the CA for bringing him or her in. Social benefits are meant to include both the cost of alerting the potential referral to the job opportunity, and any altruistic or reciprocal transfers that the referral would make for being given this opportunity. α_j may therefore be positive or negative, and we make no assumptions about it's relationship to Q_j or Y_j. Each potential referral of gender g is independently drawn from a joint distribution $f^g\left(\alpha, Q\right)$. In addition to social payments, CAs may also consider ambient incentives to refer a high quality worker $\left(E\left[R\left(Y\right)|Q\right]\right)$, which perhaps derive from reputational effects, as well as any direct financial incentives provided by the firm $\left(E\left[P_i\left(Y\right)|Q\right]\right)$. $R\left(Y\right)$ is presumed to be increasing in Y. For simplicity, we consider contracts of the form $F_i + P_i I\left(Y_j > c\right)$, that is, contracts where the CA receives a fixed fee F_i for referring anyone, and an additional P_i if their referral qualifies by performing better than some qualification threshold.

The CA problem is to find the optimal referral. The entire network is $\mathcal{N}_i = \mathcal{M}_i \cup \mathcal{F}_i$, where $\mathcal{M}_i$ $(\mathcal{F}_i)$ is the set of potential male (female) referrals. In an unrestricted setting, when

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

CAs can choose from the entire network $\mathcal{N}$, CAs solve

$$\max_{j \in \mathcal{N}_i} E\left[R\left(Y_j\right)|Q_j\right] + \alpha_j + E\left[P_i\left(Y_j\right)|Q_j\right] + F_i$$

With these contracts, the level of fixed fees does not affect the relative returns to referring different network members. Therefore, we can summarize the solution to this referral problem in terms of the level of performance pay. Suppose person N_P^* is the optimal referral from the full network $\mathcal{N}$ under contract (F, P), and person G_P^* is the optimal referral in network of gender $\mathcal{G}$. Finally, define a contact j as referrable at contract (F_i, P_i) if the CA can expect positive profits from referring j at that contract, that is, if $E\left[R\left(Y_j\right)|Q_j\right] + \alpha_j + E\left[P_i\left(Y_j\right)|Q_j\right] + F_i > 0$.

4.1 Mechanisms

In this framework, men may be systematically chosen as referrals for three reasons: first, if $N_M > N_F$, then even if the underlying distributions of social costs and quality are similar, men will maximize that distribution more frequently just because there are additional draws to find the maximum. Second, men may be chosen systematically if workers believe there are higher quality male referrals and because they are trying to maximize the quality of the worker who is referred either because of ambient reputational incentives or because of explicit performance incentives. Finally, factors from the social environment may induce CAs to refer men. We consider two types of social considerations: the distribution of social benefits, α, and the accuracy of quality signals, which may interact with the firm incentives and social payments to refer more men or women. We consider the implications for each of these in turn.

4.1.1 Scarcity

Definition 1 *CAs choose men more frequently under contract (F_i, P_i) due to* ***scarcity of potential female references*** *if $N_M > N_F$ and $P\left(j = N_{P_i}^* | j \in \mathcal{M}_i\right) = P\left(j = N_{P_i}^* | j \in \mathcal{F}_i\right)$*

If a potential referral is equally likely to be the best referral under contract (F_i, P_i) whether that person is male or female, and the only difference is that there are more draws of men in the network than of women, then the probability that a man is referred under contract $(F_i, P_i) = N_M / (N_M + N_F)$. In practice, N_M and N_F are unobserved to the econometrician. Intuitively, however, if referrable women are much more scarce in CA networks than referrable men, then we should observe two things. First, CAs will refer other men more frequently (when they can choose from the entire network). Second, CAs will make a referral more often when they are restricted to refer men than when they are restricted to refer women. Proposition 1 formalizes this intuition to create a bound on the fraction of men we should see referred if the only difference between men and women is that women are more scarce.

Proposition 1 *Define* $N_M^0(F_i, P_i)$ $\left(N_F^0(F_i, P_i)\right)$ *to be the number of referrable men (women) in network* $\mathcal{N}_i$ *under contract* (F_i, P_i)*, that is, those for whom the expected payoffs to the CA from making them as a referral are positive. Define* $\psi(F_i, P_i) = E[N^*_{P_i} \in \mathcal{M}_i | E[R(Y_{N^*_{P_i}})|Q_{N^*_{P_i}}] + \alpha_{N^*_{P_i}}$ $+ E[P_i(Y_{N^*_{P_i}})|Q_{N^*_{P_i}}] + F_i > 0]$ *Then:*

(a): If scarcity of potential female referrals is the only cause of women's disadvantage in being referred, then $\psi(F_i, P_i) = N_M^0(F_i, P_i) / \left(N_F^0(F_i, P_i) + N_M^0(F_i, P_i)\right)$

(b): Suppose CA's choose not to make a referral A *times as often when required to refer women as when required to refer men, and that each network member of gender* G *has probability* $\Pi^G(F_i, P_i)$ *of being referrable at contract* (F_i, P_i)*. Then if scarcity fully explains male preference, then* $\exists \gamma < 0$ *s.t.*

$$\frac{(1-\psi(F_i, P_i))}{\psi(F_i, P_i)} = \frac{\Pi^F(F_i, P_i)}{\log(1-\Pi^F(F_i, P_i))} \frac{\log\left(1-\Pi^M(F_i, P_i)\right)}{\Pi^M(F_i, P_i)} + \gamma \log A$$

Proposition 1 formalizes a bound for the scarcity of women among potential referrals, which depends on the attrition rates of men relative to women when induced to make a referral and the probability that an individual person of gender G is referrable. In practice $\frac{\Pi^F(F_i,P_i)}{\log(1-\Pi^F(F_i,P_i))} \frac{\log\left(1-\Pi^M(F_i,P_i)\right)}{\Pi^M(F_i,P_i)}$ serves as a lower bound for $\frac{(1-\psi(F_i,P_i))}{\psi(F_i,P_i)}$ if $A < 1$ and an upper bound if $A > 1$. Of course, we also don't observe the probability that an individual network member of gender G is referrable; however, we can use the fact that the overall attrition rate among gender G is analytically $\left(1 - \Pi^G(F_i, P_i)\right)^{N_G}$ to back out what ψ must be at different values of N_F and N_M.

4.1.2 Search for Quality

A second possibility is that men refer men more frequently because CAs are trying to refer the highest quality worker in their network because of ambient or explicit incentives provided by the firm, and that person is more likely to be male than female. In the model, this is suggested if $E\left[R\left(Y_{M^*_{P_i}}\right) + P_i\left(Y_{M^*_{P_i}}\right)\right] > E\left[R\left(Y_{F^*_{P_i}}\right) + P_i\left(Y_{F^*_{P_i}}\right)\right]$.

Since both $R(Y_j)$ and $P_i(Y_j)$ are non-decreasing in Y_j, we can simply test for whether optimal male referrals are higher or lower quality than optimal female referrals. Moreover, if the search for a high quality worker leads to women's disadvantage, then we would expect the optimal referral in the full network to be at least as skilled as the optimal referral in either restricted network. Thus, if responses to employer incentives and scarcity are the only causes of women's disadvantage, then we would anticipate that $E\left[Y_{j_N^*}\right] \geq E\left[Y_{j_M^*}\right] > E\left[Y_{j_F^*}\right]$.[16] Comparing quality distributions of referrals made under various gender restrictions and contract types allows a direct test of this hypothesis.

4.1.3 Social Environment in Referral Choice

If scarcity and firm incentives are insufficient to explain the male preference exhibited by CAs, then we know that one of the other aspects of the referral choice problem must be an important contributor. There are two other aspects to the problem: social benefits, α_j, and information, σ_ε^g.[17] Both of these can be considered part of the social environment in the referral choice problem: social benefits contain search costs and altruistic payments, which are surely related to social interaction, and the precision of quality signals seems likely to be informed by the frequency or appropriateness of social interaction, as well. We begin by characterizing the solution to the referral choice problem when these social considerations are important.

Proposition 2 *$E\left[Y_{G_{P_i}^*}\right]$ is non-decreasing in P_i. $P\left(Y_{G_{P_i}^*} > Y_{G_0^*}\right)$ is increasing in P_i iff (i): $N_G > 1$; (ii): there is positive probability of observing someone who is both better in expectation than the person who is being referred under fixed fees and whose social payments are not much lower in gender G networks*[18]*; and (iii): $\sigma_\varepsilon^g < \infty$. If any of conditions (i),(ii), or (iii) fail than $P\left(Y_{G_{P_i}^*} > Y_{G_0^*}\right) = 0$.*

This proposition allows us to identify situations where social payments and information are important by examining how referral performance changes with performance incentives. All three of these conditions are necessary, and together they are sufficient. Condition (ii) means in practice that social incentives are not perfectly correlated with quality, and that

[16] Note that this test is incorrect if the relationship between quality signals Q_j and actual quality Y_j are different between the two genders, either because CAs signals are biased for one gender or because of informational differences. We consider this possibility below.

[17] One thing not explicitly considered here is that there may be bias in quality signals when referring one gender or the other, so that $E\left[\varepsilon_j | j \in \mathcal{G}\right] \neq 0$. We ignore this potential for now as our test for response to firm incentives under different information will be robust to this possibility as well.

[18] "Not that much lower" depends on how much higher quality the person could be. The specific condition is $\int_{Q_0}^{\infty} \int_{\alpha_0+E[R(Y_0)-R(Y)|Q_0,Q]+P_i\left(\Phi\left(\frac{c-Q}{\sigma_\varepsilon^g}\right)-\Phi\left(\frac{c-Q_0}{\sigma_\varepsilon^g}\right)\right)}^{\alpha_0+E[R(Y_0)-R(Y)|Q_0,Q]} f^g\left(\alpha, Q\right) d\alpha dQ > 0$

social incentives aren't discontinuously lower for higher quality people. Therefore, if we observe referral quality increasing with performance incentives, we will know that: CAs have networks with multiple potential referrals; there are important social benefits in those networks which are not perfectly correlated with referral quality; and that CAs have useful information about the quality of their potential referrals. The failure of any one of these assumptions, however, suggests that referral quality should be unaffected by increased performance incentives.

Social Benefits The most direct social considerations are the social benefits, α_j. If men's distribution of social benefits dominates women's, then CAs may systematically refer men in an effort to receive these social benefits. Our experimental framework does not allow a direct test of the differences in social benefits across genders and to a large extent it will be a residual explanation. However, as proposition 2 shows, we will only see the performance of referrals increase in response to a sufficiently large increase in performance pay if social benefits are important and not perfectly correlated with referral ability, providing evidence of the importance of social benefits.

Information If CAs have different information about male and female referrals, then men may be referred more often under fixed fee payments if reputational incentives are concave, and they may be referred more often under performance pay incentives both because of concave reputational incentives and because of efforts to obtain performance pay. We can provide evidence that useful information exists for each gender if referral quality improves when performance pay is increased (when CAs must refer that gender). However, if referral quality does not respond to performance pay in one gender, we will not know whether information or other characteristics of the referral pool are different. The role of information can, though, be isolated when CAs can choose from their entire network, $\mathcal{N}$.

Proposition 3 *When individuals choose referrals from the full network $\mathcal{N}_i$, the probability of referral qualification is increasing in P_i. If social incentives are not important, or if P_i is large enough, then $P\left[Y_{N_P^*} > c\right] \geq P\left[Y_{G_P^*} > c\right]$ $\forall G$. If information is finite and the same between men and women $\left(\sigma_\varepsilon^F = \sigma_\varepsilon^M < \infty\right)$ then Proposition 2 applies to unrestricted choices and performance premia will be positive unless condition (ii) fails for at least one of the genders. If CAs have worse information about women $\left(\sigma_\varepsilon^F > \sigma_\varepsilon^M\right)$, the relationship between referral quality and performance pay is ambiguous.*

When the full network can be drawn upon for a referral, CAs have the option of referring the same men and women they choose to refer under performance pay. This means that if they have useful information about men, then then they have the opportunity to use that information when their referral choices are unrestricted across genders. However, they may not: while loosening restrictions on referral choices is guaranteed to bring in referrals who generate larger payoffs for CAs, these payoffs could be larger in terms of either social payments or expected performance pay. Proposition 3 suggests that when information is the same about men and women, any CA who changes their referral choice under performance pay will do so to bring in referrals who are higher quality in expectation.[19] However, when information is worse about women, CAs may opt to choose referrals who are worse in expectation under performance pay. This happens because the low ability women face a higher probability of earning the performance bonus than similarly low ability men from the CA's perspective. In other words, when information is worse about women, CAs may choose to take a gamble on a high social payment but apparently low ability woman, rather than a low social payment but high ability man. This can reduce the performance of referrals when CAs can choose from the entire network $\mathcal{N}$ for small enough performance incentives.

5 Empirical Evidence of Mechanisms

Our approach for assessing the causes of male bias is as follows: First, we examine CAs' choice of whether or not to refer someone to determine whether we can bound the effects of scarcity of women on the referral behavior. Next, we look at the quality of men and women referred by male CAs to look for evidence of CAs responding to direct firm incentives versus other social considerations. Finally, we assess whether information can serve as a source of productive bias against women and how social payments may interact with differential information in referral choices by examining how referral performance changes as performance incentives change.

5.1 Are Women Very Scarce in Men's Networks?

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,

[19] This could either because they are identifying a woman who is higher quality than the man who would have been referred under a fixed fee, or because they are bringing in a better person of the same gender.

referrals serve to perpetuate job segregation due to the limited overlap of groups from which referrals are drawn.

The theory above suggested that the rates at which CAs make referrals can help us assess whether women are very scarce in the network. We randomly restricted some CAs to referring only women, and other CAs to referring only men: this allows us to look at how likely CAs are to know men and women who are referrable at our contracting terms. 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. 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.

Columns (1)-(2) of Table 2 presents this analysis, where treatments where the CAs were restricted to referring only men (or male fixed fee treatments in specifications which disaggregate by contract terms) are the excluded group. Overall, men are not 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 contract type, as in column (2), 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. In terms of the model, this suggests that CAs are receiving a number of draws of both men and women. Even if all the draws of a CA's own gender fail to make the participation threshold to bother making a referral, for some CAs there is a draw from the other gender who exceeds it.

Under both contract types, this suggests that A, the relative return rate of women compared to men, is about one. Intuitively, this behavior seems unlikely to be generated by a situation where suitable women referrals are tremendously scarce relative to men, as male

CAs are similarly likely to know at least one man and at least one woman who is worth referring. To generate this result if women are really scarce in networks relative to men, any woman in a network must be much more likely to individually qualify than any man in a network. Proposition 1 allows us to more directly bound the role of scarcity in women's disadvantage. Since $A = 1$, we infer that if scarcity is the only source of women's disadvantage, then $\frac{(1-\psi(F_i,P_i))}{\psi(F_i,P_i)} = \frac{\Pi^F(F_i,P_i)}{\log(1-\Pi^F(F_i,P_i))} \frac{\log\left(1-\Pi^M(F_i,P_i)\right)}{\Pi^M(F_i,P_i)}$ or that the scarcity of suitable women relative to suitable men can be summarized by the probability that either women or men are referrable. We of course do not know the probability that any individual woman or man is referrable; however, we know that $1 - \left[1 - \Pi^F\left(F_i, P_i\right)\right]^{N_F} = E\left[R_i | T_i = female, F_i, P_i\right]$ and $N_F \geq 1$. Our approach is to fix N_F, solve for the implied $\Pi^M\left(F_i, P_i\right)$ using the empirical $\psi\left(F_i, P_i\right)$, and evaluate the plausibility of the implied network characteristics to determine whether scarcity alone can achieve the gender referral ratios that we see in the data.

From Table 1, we see that men are referred under fixed fees 75% of the time, while men are referred under performance pay 78% of the time. Thus, for scarcity to explain why women are left out under fixed fee contracts, we would need $\left(\log\left(1 - \Pi^M\left(F_i, P_i\right)\right) / \Pi^M\left(F_i, P_i\right)\right) * (0.9/\log 0.1) = 0.33$, if $N_F = 1$. Appendix Table 2 works out the cases for $N_F \in \{1, 2, 3\}$ (these are a bound on the potential of scarcity if $N_F > 3$) and reveals that there is one type of network under which scarcity can come close to the gender disparities we observe: if there is exactly one woman who is referrable 90% of the time under fixed fees and very many men who are almost never referrable, then scarcity can create a situation where men are referred 72% of the time. However, if there is more than one woman in the network or if men are closer in probability of being referrable to women, then scarcity cannot explain the gender preference that men exhibit for other men under fixed fee contracts. Given that networks with exactly one woman who is very likely to be referrable and many men who are very unlikely to be referrable under fixed fees seem a priori unlikely, we feel comfortable rejecting scarcity as an explanation for our fixed fee results. Moreover, under performance pay, scarcity can never come close to explaining observed patterns in referral behavior under any assumptions about the network structure: with a 75% return rate under performance pay, it is impossible for scarcity of women to explain why men refer men more than 65% of the time and yet refer someone similarly frequently when required to refer a man or when required to refer a woman, when in fact they refer men 78% of the time under performance pay.[20]

[20]We will discuss attrition and its impact on our subsequent analysis in section 5.3.1. Note, however, that in

5.2 How Do the Women in Men's Networks Perform?

Since scarcity of women cannot fully explain men's preference for men, we examine whether gender differences in the quality distributions of men's network members combine with CA efforts to identify a high quality worker to generate a sufficient explanation for the observed bias against women. Since real world hires feature a range of contracts with different degrees of firm incentives, we focus on contracts without explicit firm incentives in this investigation (our fixed fee contracts). Our model highlights that in this treatment, CAs may have an incentive to refer a high quality person due to concerns over their reputation. Under our fixed fee contracts, we continue to find that men are referred significantly more often than women. If firm incentives can generate the bias under fixed fee contracts, they remain a plausible explanation across the range of contract types. The random variation in the gender of the referral allows us to directly investigate whether the women that men would have referred are on average worse quality than the men they prefer to refer.

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 entirely driven by differences in men's and women's qualifications in the network.

Therefore there must be other factors affecting men's referral patterns. The model highlights two which stem from the CAs' social environment: social benefits and information. If referring male network members generates more social benefits (or are less costly to recruit), then this could explain why men are referring other men in fixed fee treatments - even when the women that they could have referred appear to be just as good. A second explanation is that men have less information about women. If the reputational incentives are concave, CAs may prefer to refer a man simply because their quality signal is more accurate. Both of these explanations are consistent with the experiment and are factors which generate a bias against women for reasons that do not affect firm profits. For example, if men have limited information about women, then that limited information impedes CAs from referring the best person in the network for the job. While we can not definitely say which of these two factors are driving

section 5.2 we compare the performance of men and women referred in performance pay treatments, where there is no different attrition.

men's preference for men in fixed fee treatments, the results show that CAs' social environment is affecting their choice of referral rather than gender disparities being driven exclusively by underlying differences in ability among members of their social network.

5.3 Financial Incentives & the Social Environment

The fixed fee treatment suggested that factors stemming from the CAs' social environment led to women benefiting less from referrals than men. When the firm intensifies productive incentives, are women placed at a further disadvantage? We find no evidence of the performance incentives favoring men in our experiment. Comparing panels C and D of Table 1 shows that male CAs refer only marginally fewer women (22% vs 25%) in performance pay than under fixed, and this difference is not statistically significant. The intensification of firm incentives in this case did not further disadvantage women. However, some results from the experiment suggest that CAs' search for high quality candidates could further disadvantage women if firm incentives were higher stakes than ours.

How CAs change their referral choices in response to performance pay, relative to the fixed fee treatments, also provides insights into CAs' underlying social environment. Propositions 2 and 3 suggest that social benefits and information - our key characteristics of the social environment - determine how CAs will respond to performance pay. We will be able to provide evidence that the social environment for men referring other men differs from that governing men's referrals of women.

We examine differences in referral behavior comparing the different gender treatments across fixed and performance pay treatments using the following specification:

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

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. Columns (3)-(4) of Table 2 presents the results of this analysis for male CAs.

We focus first on male CAs' behavior when referring other men. Column (4) shows that they refer significantly better candidates when given a performance pay incentive: candidates are approximately 27 percentage points more likely to qualify if the CA was in a performance pay treatment than in fixed. Given that the qualification rate is about 50%, this is a very large premium. This demonstrates two points. First, CAs were not referring the best person in the network for the job in the fixed fee treatments. As in Beaman and Magruder (2012), we see evidence of social benefits which skew the CA's behavior away from what would benefit the firm most. Even in this setting where CAs may internalize some of the firm's objective function because of reputation, the firm still needs to offset the incentives created within the CA's social network.[21] Second, CAs must have useful information about the male members of their network. Otherwise, even if they attempted to recruit a better person, we would not see any increase in the actual qualification rate. We therefore conclude that male CAs have useful information for employers about men, and the price of eliciting the information is not prohibitively high.

However, column (4) also shows that male CAs do not create a performance premium when restricted to refer women (the sum of the interaction term with Female Treatment and Performance pay is essentially zero). Simple descriptive statistics demonstrate clearly among CAs in performance pay treatment, the referred men outperform the referred women: 65% of referrals qualify in the male-only treatment vs 47% in the female-only treatment. As a result, we conclude that either men have worse information about women, or that it is more costly to elicit this information about women (i.e. the performance premium would have to be more steep in order to induce them to forego higher social payments - or incur higher recruitment costs - and refer a higher equality female candidate). In either case, the firm gets the highest quality candidates by asking male CAs to refer other men and providing a performance incentive.

Table 2 further shows that there is no performance premium in the either-gender treatments, as the sum of the Performance pay coefficient and the Either*Perf coefficient is approximately zero. While men CAs respond to the performance incentive when they must refer other men by referring better quality people, they don't have this response when they can refer whomever they wish. The model provides an explanation for this: a noisy signal of women's ability combined with social benefits can lead CAs to prefer referring a potentially lower quality woman who gives high social benefits over a higher quality man with certain low social

[21]Reputational concerns may be higher or lower in this setting, where candidates are making referrals and not existing employees. On the one hand, the probability of getting the job is less than one - reducing the CA's worry about his reputation. On the other hand, the firm has very little information about the candidate - compared to existing employees - and therefore a bad referral may be much more damaging to the firm's opinion of the candidate.

benefits.[22] In this case, poor information about women and social benefits combine to limit the firm's ability to use men's superior information about other men. Taken together, the differences in response to performance incentives when men have to pick only one gender or can refer anyone suggest that there are important differences in the social environment between men's networks of men and men's networks of women, supporting the conclusions of section 5.2. Moreover, these differences suggest that there is the potential for firm incentives to increase male bias: after all, the highest quality referrals overall are men's referrals of men and presumably a strong enough incentive would induce greater male bias. Thus, even though our performance pay contract does not elicit this behavior, we conclude that other contracts may: in particular, contracts which put more emphasis on the quality of the referred candidate may induce CAs to ignore the forgone social benefits and refer even more men than we observed in the experiment.

Table 3 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. This is consistent with men having little information about women, or the financial incentives not being large enough to induce the CAs to choose a better candidate and give up associated social benefits.

We have thus far discussed men's signals of women as simply being noisier. CA expectations of referral quality may also be biased. In that case, ε_j^g is not mean zero. Biased beliefs are also likely to stem from factors in the CA's social environment. However, men's referral patterns are largely inconsistent with biased expectations alone driving the behavior that leads to few qualified women getting referred. If men (incorrectly) underestimate women's ability, we would anticipate that this bias is increased under performance pay. As a result, if bias in men's beliefs about women is driving male preference, we would expect men to refer even more men under performance pay than under fixed fees. However, as Table 1 indicates, men refer men at the same rate under performance pay as under fixed fees. Similarly, men would also

[22] In fact, there are two explanations within the model which can yield better male referrals under performance pay and similar unrestricted and female referrals (who are statistically worse than the male performance pay referrals). One is worse information about women combined with social incentives, and the other is that the underlying joint distribution of women's quality and social incentives is discontinuous, with there being no probability of observing a woman who is both higher ability than the women referred under fixed fees and who gives social payments which are only slightly less. Both of these explanations suggest the social environment as the key contributing factor; however we focus on the information explanation as discontinuities in the distribution of potential women network members seems ex ante less plausible.

make a referral less often when they have to refer a woman in a performance pay treatment if they have biased beliefs. Table 2 showed that this is not the case either. Together with the evidence from section 5.2, we conclude that (biased or unbiased) quality expectations appear to play a small, if any, role in explaining why men refer so few women.

5.3.1 Attrition

In section 5.1, 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.[23] 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 results is that all CAs will only refer one particular person, but CAs will just attrit rather than refer that person under performance pay if they are in a restricted male treatment and that person is low quality. Figure 5 suggests that among men, there is assortative matching in ability between the CA and their referral under fixed fees. However, Figure 5 also shows that the performance premium exists throughout the entire distribution of CA test scores, which makes attrition bias less likely to be driving the results in table 2.[24]

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 2 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.

[23] In Section 6, we will also note that female CAs responded similarly.

[24] Since the referral patterns are similar across the entire distribution of CA ability, we have more confidence that the results extrapolate to other contexts where only existing employees make referrals.

6 Women

Figure 2 showed that women refer other women about 40% of the time, which is statistically the same rate that women apply themselves through the traditional method. Given that women CAs exhibit less of a gender preference in selecting referrals than men CAs, a natural hypothesis is that firms could use women to make references and avoid gender bias while recruiting highly skilled employees. A closer look at our experiment, however, rejects this hypothesis, as the prospective employees (and particularly the women) referred by women CAs are significantly less likely to qualify for the position than either the pool of traditional applicants or men's referrals. Figure 6 and Table 4 reveal that, on average, women refer people who are eighteen percentage points less likely to qualify for a position than men's referrals, a difference which is highly statistically significant. Figure 7 demonstrates that the people referred by women are systematically less likely to qualify than the people referred by men across the range of CA abilities, and again with particularly large differences for highly able CAs.

Figure 8 presents kernel densities of female CAs' referrals' scores in the fixed fee treatments to test whether there may be differences in the quality of referrals in women's networks of men and women. The ability distribution of referred men clearly stochastically dominates the distribution of referred women, with the Kolmogorov-Smirnov test rejecting the distributions being the same at the 5% level. In terms of means, the referred women perform on average 0.42 of a standard deviation below the CA mean, while men referred by women CAs perform 0.08 standard deviations below the CA mean. Moreover, the introduction of moderate performance incentives does not lead to higher quality referrals by women CAs, as Column 4 of Table 5 shows. Our results therefore indicate that women's referrals of other women are too unlikely to qualify to be hired to offset men's referral behavior and create balance in the workforce.

Table 6 shows referral performance disaggregated by component for women CAs. This is suggestive evidence 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. 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 this information does not translate into a choice of women or men who are likely to qualify (at the level of incentives offered in the experiment).[25]

In the model, one mechanism that could explain this is social benefits: women CAs may receive very high social benefits from referring women whose skills translate poorly to survey enumeration. Our experiment says nothing directly about the relationship between social benefits and members' abilities in women's networks. There is, however, very suggestive evidence from other literatures that women tend to invest more in close ties and less in weak ties that - according to Granovetter (1973) - are most useful for a job search (Seabright, 2012). Social psychology also suggests that women do more helping in long-term, close relationships while men display helping behaviors with a wider range of people(Eagley and Crowley, 1986). It is possible that a larger performance reward could induce women to refer better quality candidates. However, it would still be cheaper for firms to get good quality candidates from their male employees.

Competition

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).[26] 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 of qualifying or getting called for a job (given the large number of recruits)[27], competition on the job may actually be stronger.

Nevertheless, if women CAs are concerned about the competitive threat their referrals

[25] Appendix figure A1 shows that there is little evidence of female CAs responding to the performance pay incentive at any point in the CA performance distribution.

[26] 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.

[27] On the median CA recruitment date, there were 61 CAs who 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.

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 latter is consistent with the results presented in Table 5: in unrestricted treatments, women refer poor quality men and women. However, several additional pieces of evidence seem inconsistent with the competition aversion hypothesis. Figure 7 shows that women who are on the margin of qualification (near a score of 60) are if anything more likely to refer someone who is qualified. This is inconsistent with women making referral choices based on a fear of losing out of the job due to their referral. Second, Tables 5 and 6 suggest that women have a hard time anticipating who will qualify. In that case, referring low quality people instead of just not making a referral is a very risky strategy. 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 as shown in Appendix Table 3.

While there are overall a few patterns in the data that suggest competition-aversion is not the only factor driving women to refer low quality candidates, we do not have conclusive evidence that rules out competition as a contributing factor. Given that in our experiment, women refer more able men than women, future research should explore this possibility as it could suggest that women need not always shy away from competing with men as in Niederle and Vesterlund (2007).

7 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. Much of this difference occurs as men exhibit a preference for referring men. We document that men's preference to refer other men cannot be fully explained by scarcity of women, or by efforts to find the highest ability worker: instead, aspects of the social environment contribute to women's disadvantage. We also document that using women to make referrals is similarly unsuccessful at identifying high ability female workers. While women CAs in our experiment do not exhibit the gender preference that men do, they do refer people (and particularly women) who are unlikely to qualify for positions. This result suggests that the ubiquity of job networks as a hiring system could contribute to persistent gender gaps in labor market outcomes. 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[28], there is reason to believe that our findings may generalize to many other contexts.

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.

Bandiera, O., I. Barankay, and I. Rasul (2009). Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica 77*, 1047–1094.

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.

Burks, S., B. Cowgill, M. Hoffman, and M. Housman (2013). "You′d be perfect for this": Understanding the value of hiring through referrals. *IZA Discussion Paper No. 7382*.

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

[28]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.

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

Eagley, A. and M. Crowley (1986). Gender and helping behavior: A meta-analytic review of the social psychological literature. *Psychological Bulletin 100*, 283–308.

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.

Granovetter, M. (1973). The strength of weak ties. *American Journal of Sociology 78*, 1360–1380.

Hensvik, L. and O. N. Skans (2013). Social networks, employee selection and labor market outcomes. *IFAU - Institute for Labour Market Policy Evaluation Working Paper No. 2013:15*.

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*.

Loury, L. D. (2006). Some contacts are more equal than others: Informal networks, job tenure, and wages. *Journal of Labor Economics 24*, 299–318.

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 Statistical Office of Malawi (2009). 2008 population and housing census: Education and literacy report. Technical report, Government of Malawi.

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.

Prendergast, C. and R. H. Topel (1996). Favoritism in organizations. *Journal of Political Economy 104*, 958–978.

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

Seabright, P. (2012). *War of the Sexes*. Princeton University Press.

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) All CAs	(2) Male CAs	(3) Female CAs	(4) 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. Referral Characteristics: Made Referral, Either Gender Treatments				
Referral is Female	30%	23%	43%	0.002
N	195	117	78	
C. Referral Characteristics: Made Referral, Either Gender, Fixed Fee Treatments				
Referral is Female	32%	25%	43%	0.042
N	117	68	49	
D. Referral Characteristics: Made Referral, Either Gender, Perf Treatments				
Referral is Female	31%	22%	45%	0.039
N	78	49	29	

Table 2: Male CA's Referral Choices

	Made a Referral				Referral Qualifies			
	(1)		(2)		(3)		(4)	
Female Treatment	-0.004		-0.004		-0.030		0.068	
	(0.038)		(0.050)		(0.062)		(0.081)	
Either Gender Treatment	0.014		-0.052		0.071		0.227	***
	(0.040)		(0.052)		(0.066)		(0.084)	
Performance Pay			-0.148	***			0.267	***
			(0.056)				(0.093)	
Perf Pay * Female Treatment			0.004				-0.248	*
			(0.076)				(0.127)	
Perf Pay * Either Treatment			0.152	*			-0.383	***
			(0.079)				(0.132)	
Observations	506		506		390		390	

Notes

1 The dependent variable in columns (1)-(2) is an indicator for whether the CA makes a referral and in (3)-(4) an indicator for whether the referral qualifies.

2 All specifications include CA visit day dummies.

Table 3: 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)
Mean of Dep Variable	0.244	0.708	2.031	7.746	1.544	4.985	15.775	27.144
SD	(0.430)	(0.455)	(0.869)	(1.279)	(1.220)	(2.336)	(4.029)	(6.304)
Observations	386	390	390	390	390	390	383	382

Notes

1 The dependent variable is listed in the column heading.

2 All specifications include CA visit day dummies.

Table 4: Quality of CAs and Referrals

	(1)	(2)	(3)	(4)
	All CAs	Male CAs	Female CAs	Diff: *p* value
A. Referral Characteristics: Made Referral, Either Gender Treatments				
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	
B. Referral Characteristics: Made Referral, Either Gender, Fixed Fee Treatments				
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	
C. Referral Characteristics: Made Referral, Either Gender, Perf Treatments				
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 5: Female CA's Referral Choices

	Made a Referral		Referral Qualifies	
	(1)	(2)	(3)	(4)
Female Referral Treatment	-0.055	-0.042	-0.190 **	-0.181
	(0.054)	(0.074)	(0.083)	(0.113)
Either Gender Treatment	0.017	-0.024	-0.231 ***	-0.242 **
	(0.055)	(0.071)	(0.082)	(0.107)
Performance Pay		-0.113		0.021
		(0.080)		(0.122)
Perf Pay * Female Treatment		-0.013		-0.022
		(0.111)		(0.171)
Perf Pay * Either Treatment		0.086		0.032
		(0.110)		(0.169)
Observations	310		227	227

Notes

1 The dependent variable in columns (1)-(2) is an indicator for whether the CA makes a referral and in (3)-(4) an indicator for whether the referral qualifies.

2 All specifications include CA visit day dummies.

Table 6: 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)
Mean of Dep Variable	0.195	0.652	1.828	7.498	1.617	4.780	16.194	27.802
SD	(0.397)	(0.477)	(0.932)	(1.486)	(1.200)	(2.407)	(4.023)	(5.800)
Observations	226	227	227	227	227	227	222	222

Notes

1 The dependent variable is listed in the column heading.

2 All specifications include CA visit day dummies.

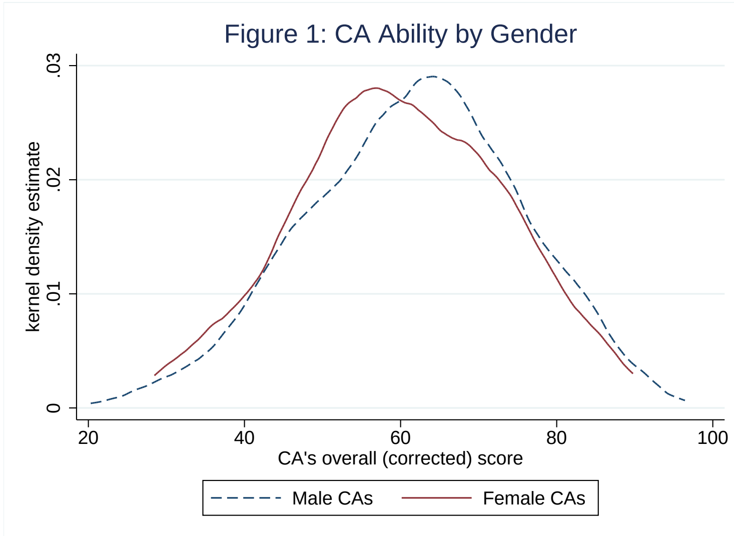
Figure 1: CA Ability by Gender
kernel density estimate
0
.01
.02
.03
20
40
60
80
100
CA's overall (corrected) score
Male CAs
Female CAs

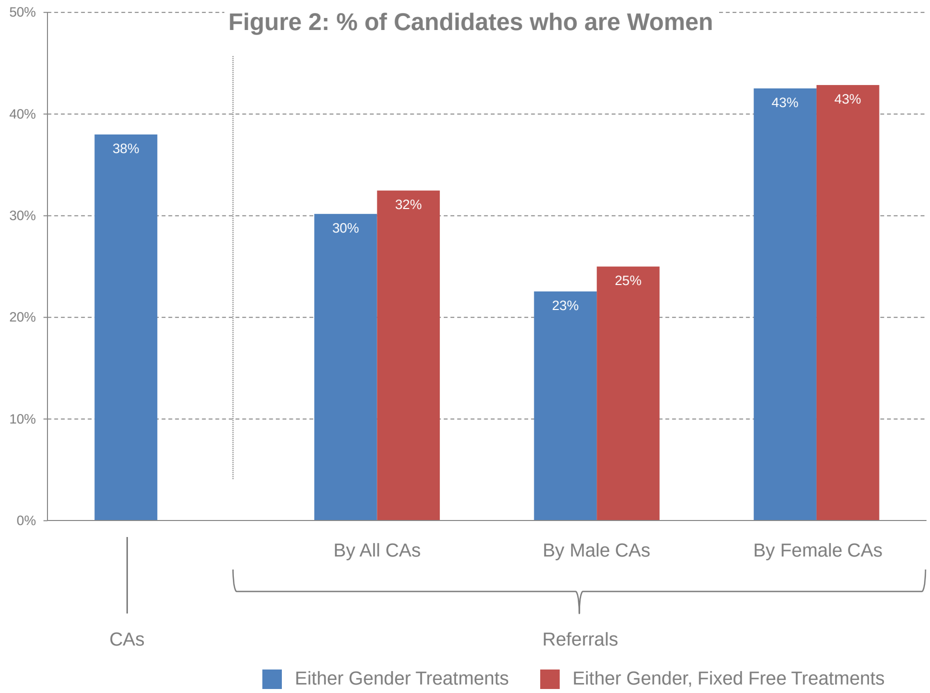
Figure 2: % of Candidates who are Women
0%
10%
20%
30%
40%
50%
38%
30%
32%
23%
25%
43%
43%
By All CAs
By Male CAs
By Female CAs
CAs
Referrals
Either Gender Treatments
Either Gender, Fixed Free Treatments

Figure 3: Gender choice in referrals, by CA performance

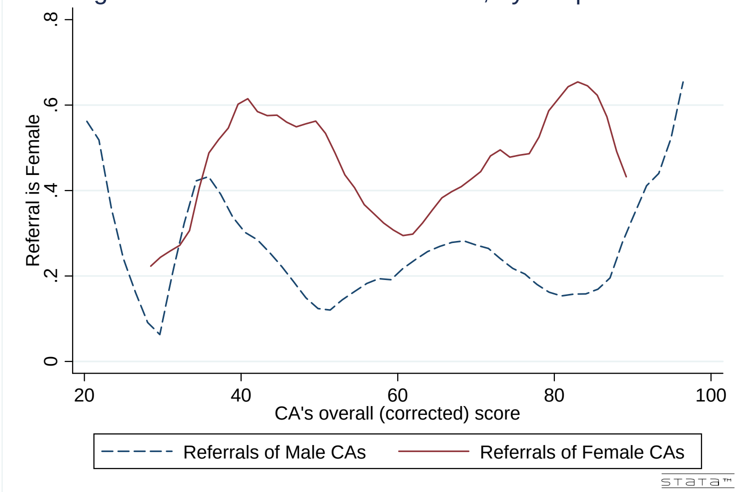

Figure 4: Men's Fixed Fee Referrals

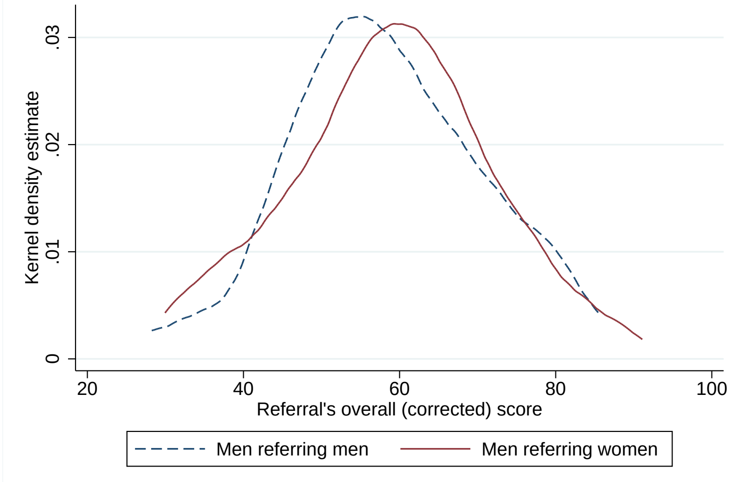

Figure 5: Referral Qualifies , by Male CA performance

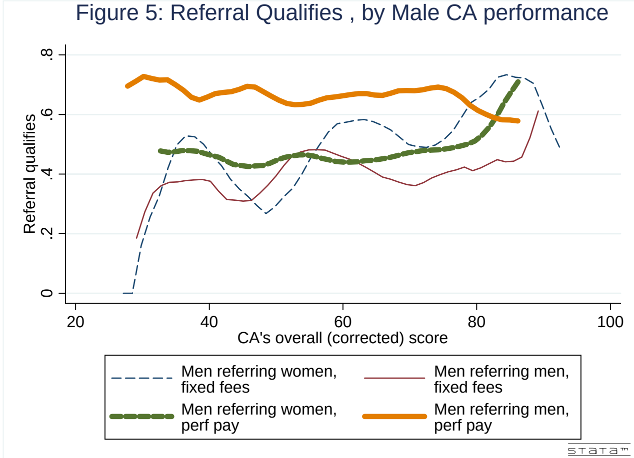

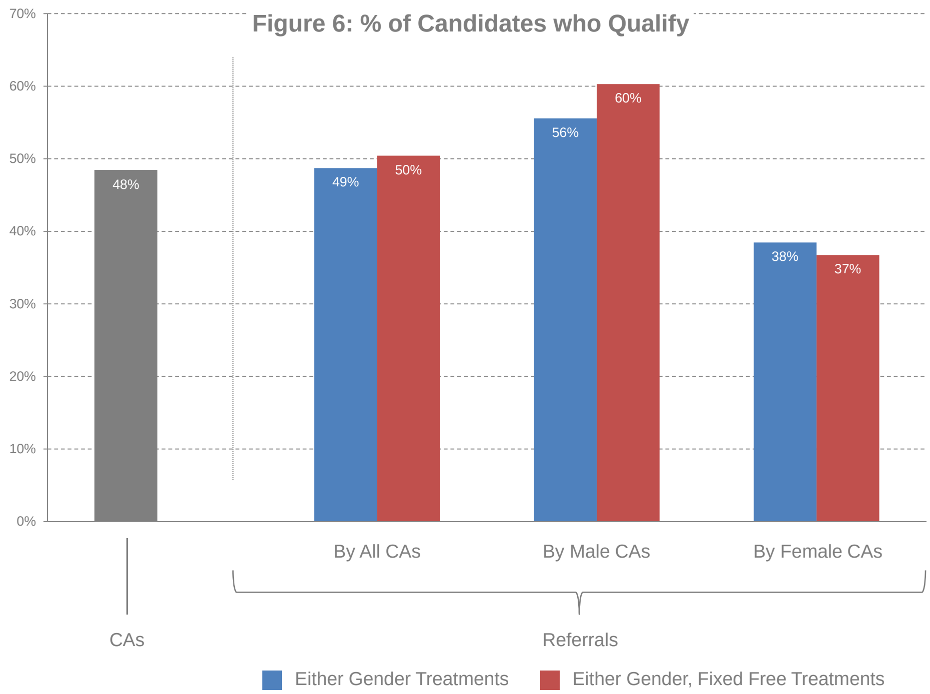
Figure 6: % of Candidates who Qualify
70%
60%
50%
40%
30%
20%
10%
0%
48%
49%
50%
56%
60%
38%
37%
By All CAs
By Male CAs
By Female CAs
CAs
Referrals
Either Gender Treatments
Either Gender, Fixed Free Treatments

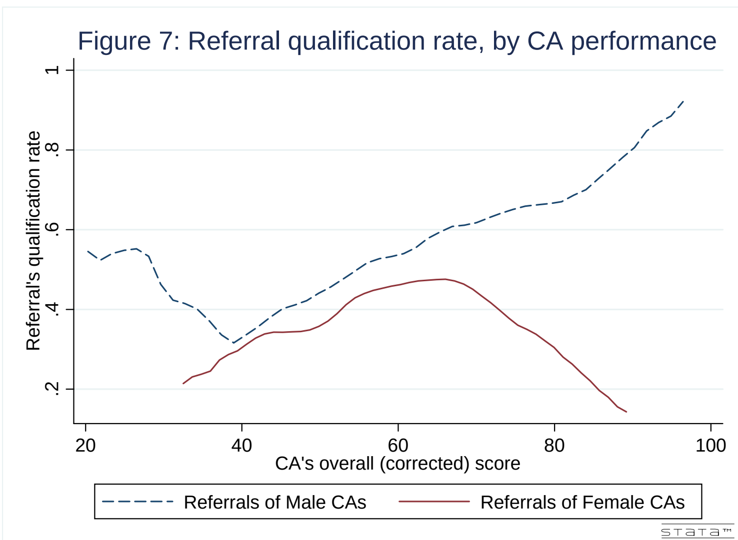
Figure 7: Referral qualification rate, by CA performance
Referral's qualification rate
.2
.4
.6
.8
1
20
40
60
80
100
CA's overall (corrected) score
Referrals of Male CAs
Referrals of Female CAs

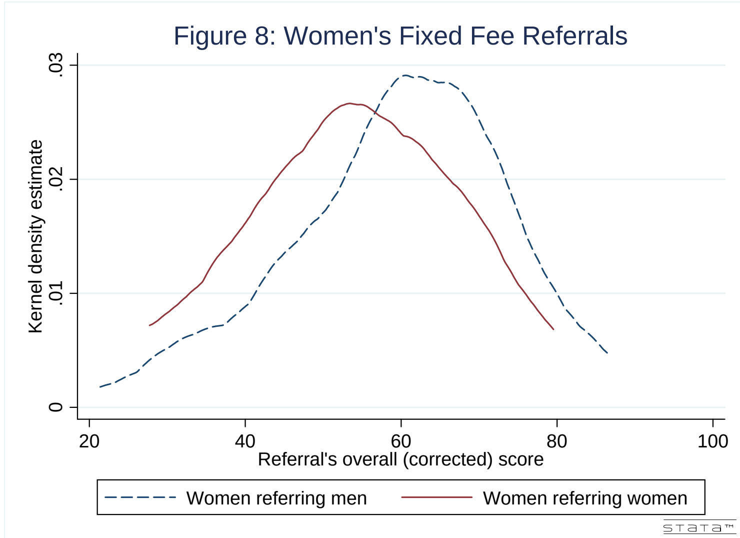

A Appendix

A.1 Proofs

A.1.1 Proposition 1

First, note that under scarcity, $\psi(F_i, P_i) = N_M^0(F_i, P_i) / \left(N_F^0(F_i, P_i) + N_M^0(F_i, P_i)\right)$, so that $N_F^0(F_i, P_i) / N_M^0(F_i, P_i) = (1 - \psi(F_i, P_i)) / \psi(F_i, P_i)$. Second, note that $E\left[N_G^0\right] = N_G \Pi^G(F_i, P_i)$ so that $N_F / N_M = \left(\Pi^M(F_i, P_i) / \Pi^F(F_i, P_i)\right) * E\left[\left(N_F^0(F_i, P_i) / N_M^0(F_i, P_i)\right)\right]$. Taking expectations and plugging back in for $N_F^0(F_i, P_i) / N_M^0(F_i, P_i)$, we have that

$$\frac{N_F}{N_M} = \frac{\Pi^M(F_i, P_i)}{\Pi^F(F_i, P_i)} \frac{(1 - \psi(F_i, P_i))}{\psi(F_i, P_i)} \tag{1}$$

Finally, note that $E[1 - R_i | \text{gender } G, F_i, P_i] = \left(1 - \Pi^G(F_i, P_i)\right)^{N_G}$, so that if male CA's decline to make a referral A times as frequently when referring women, we have

$$\left(1 - \Pi^F(F_i, P_i)\right)^{N_F} = A\left(1 - \Pi^M(F_i, P_i)\right)^{N_M} \tag{2}$$

or

$$N_F \log\left(1 - \Pi^F(F_i, P_i)\right) = \log A + N_M \log\left(1 - \Pi^M(F_i, P_i)\right) \tag{3}$$

Solving for N_F/N_M, we have

$$\frac{N_F}{N_M} = \frac{\frac{1}{N_M}\log A + \log\left(1 - \Pi^M(F_i, P_i)\right)}{\log(1 - \Pi^F(F_i, P_i))} \tag{4}$$

Substituting in from 1,we find

$$\frac{(1 - \psi(F_i, P_i))}{\psi(F_i, P_i)} = \frac{\Pi^F(F_i, P_i)}{\log(1 - \Pi^F(F_i, P_i))} \frac{\log\left(1 - \Pi^M(F_i, P_i)\right)}{\Pi^M(F_i, P_i)} + \frac{\frac{1}{N_M}\left(\Pi^F(F_i, P_i) / \Pi^M(F_i, P_i)\right)}{\log(1 - \Pi^F(F_i, P_i))} \log A \tag{5}$$

Define $\gamma = \frac{1}{N_M}\left(\Pi^F(F_i, P_i) / \Pi^M(F_i, P_i)\right) / \log\left(1 - \Pi^F(F_i, P_i)\right)$, which is less than 0 since $1-\Pi^F(F_i,P_i)<1$ and other expressions in γ are positive.

A.1.2 Proposition 2

Some notation: $k >_P j$ suggests that person k is preferred under contract (F_i, P) to person j. Thus, person $k >_0 j$ implies that k is preferred under fixed fee contracts to person j. First, note

$$j >_P k \;\Rightarrow \alpha_j > E\left[R\left(Y_k\right)|Q_k\right] - E\left[R\left(Y_j\right)|Q_j\right] + \alpha_k + P_i\left(\Phi\left(\frac{c-Q_j}{\sigma_\varepsilon^g}\right) - \Phi\left(\frac{c-Q_k}{\sigma_\varepsilon^g}\right)\right) \tag{6}$$

If person k is preferred under fixed, then

$$\alpha_j < E\left[R\left(Y_k\right)|Q_k\right] - E\left[R\left(Y_j\right)|Q_j\right] + \alpha_k \tag{7}$$

Thus, if $j >_P k$ and $k >_0 j$, it must be the case that $\Phi\left(\frac{c-Q_j}{\sigma_\varepsilon^g}\right) < \Phi\left(\frac{c-Q_k}{\sigma_\varepsilon^g}\right)$, or that $Q_j > Q_k$. Thus, Q is non-decreasing in P_i, and since $Y = Q + \varepsilon$ (and ε is independent of Q), so is Y.

Next, note that since draws are independent, $P\left(j <_0 k | l <_0 k\right) = P\left(j <_0 k\right)$. Thus

$$\begin{aligned} P\left(Y_p > Y_F\right) &= 1 - P(\text{no one is better under } P \text{ than } Y_k | k \text{ is best under fixed}) = \\ & \quad 1 - P\left(1 <_P k | k \text{ is best under fixed}\right) P\left(2 <_P k | k \text{ is best under fixed}\right) \cdots \end{aligned} \tag{8}$$

so that

$$P\left(k >_P j\right) = \left[\begin{array}{c} \frac{\int_{-\infty}^{Q_k}\int_{-\infty}^{E[R(Y_k)|Q_k]-E\left[R\left(Y_j\right)|Q_j\right]+\alpha_k} f^g(\alpha_j,Q_j)d\alpha_j dQ_j}{\int\int_{-\infty}^{E[R(Y_k)|Q_k]-E\left[R\left(Y_j\right)|Q_j\right]+\alpha_k} f^g(\alpha_j,Q_j)d\alpha_j dQ_j} + \\ \frac{\int_{Q_k}^{\infty}\int_{-\infty}^{E[R(Y_k)|Q_k]-E\left[R\left(Y_j\right)|Q_j\right]+\alpha_k+P_i\left(\Phi\left(\frac{c-Q_j}{\sigma_\varepsilon^g}\right)-\Phi\left(\frac{c-Q_k}{\sigma_\varepsilon^g}\right)\right)} f^g(\alpha_j,Q_j)d\alpha_j dQ_j}{\int\int_{-\infty}^{E[R(Y_k)|Q_k]-E\left[R\left(Y_j\right)|Q_j\right]+\alpha_k} f^g(\alpha_j,Q_j)d\alpha_j dQ_j} \end{array}\right] \tag{9}$$

and

$$P\left(k >_p j\right)\forall j = \left[\begin{array}{c} \frac{\int_{-\infty}^{Q_k}\int_{-\infty}^{E[R(Y_k)|Q_k]-E\left[R\left(Y_j\right)|Q_j\right]+\alpha_k} f^g(\alpha_j,Q_j)d\alpha_j dQ_j}{\int\int_{-\infty}^{E[R(Y_k)|Q_k]-E\left[R\left(Y_j\right)|Q_j\right]+\alpha_k} f^g(\alpha_j,Q_j)d\alpha_j dQ_j} + \\ \frac{\int_{Q_k}^{\infty}\int_{-\infty}^{E[R(Y_k)|Q_k]-E\left[R\left(Y_j\right)|Q_j\right]+\alpha_k+P_i\left(\Phi\left(\frac{c-Q_j}{\sigma_\varepsilon^g}\right)-\Phi\left(\frac{c-Q_k}{\sigma_\varepsilon^g}\right)\right)} f^g(\alpha_j,Q_j)d\alpha_j dQ_j}{\int\int_{-\infty}^{E[R(Y_k)|Q_k]-E\left[R\left(Y_j\right)|Q_j\right]+\alpha_k} f^g(\alpha_j,Q_j)d\alpha_j dQ_j} \end{array}\right]^{N_G-1} \tag{10}$$

and

$$P\left(Y_{G_P^*} > Y_{G_0^*}\right) = 1 - \left[\frac{\int_{-\infty}^{Q_k} \int_{-\infty}^{E[R(Y_k)|Q_k] - E\left[R\left(Y_j\right)|Q_j\right] + \alpha_k} f^g(\alpha_j, Q_j) d\alpha_j dQ_j}{\int \int_{-\infty}^{E[R(Y_k)|Q_k] - E\left[R\left(Y_j\right)|Q_j\right] + \alpha_k} f^g(\alpha_j, Q_j) d\alpha_j dQ_j} + \frac{\int_{Q_k}^{\infty} \int_{-\infty}^{E[R(Y_k)|Q_k] - E\left[R\left(Y_j\right)|Q_j\right] + \alpha_k + P_i\left(\Phi\left(\frac{c - Q_j}{\sigma_\varepsilon^g}\right) - \Phi\left(\frac{c - Q_k}{\sigma_\varepsilon^g}\right)\right)} f^g(\alpha_j, Q_j) d\alpha_j dQ_j}{\int \int_{-\infty}^{E[R(Y_k)|Q_k] - E\left[R\left(Y_j\right)|Q_j\right] + \alpha_k} f^g(\alpha_j, Q_j) d\alpha_j dQ_j} \right]^{N_G - 1} \tag{11}$$

For Proposition 2, notice that equation 11 $= 0$ iff $N_G = 1$ or

$$\begin{aligned} &\int_{Q_k}^{\infty} \int_{-\infty}^{E[R(Y_k)|Q_k] - E[R(Y_j)|Q_j] + \alpha_k + P_i\left(\Phi\left(\frac{c - Q_j}{\sigma_\varepsilon^g}\right) - \Phi\left(\frac{c - Q_k}{\sigma_\varepsilon^g}\right)\right)} f^g\left(\alpha_j, Q_j\right) d\alpha_j dQ_j \\ &= \int_{Q_k}^{\infty} \int_{-\infty}^{E[R(Y_k)|Q_k] - E[R(Y_j)|Q_j] + \alpha_k} f^g\left(\alpha_j, Q_j\right) d\alpha_j dQ_j \end{aligned} \tag{12}$$

For this not to hold, we would need

$$\int_{Q_k}^{\infty} \int_{E[R(Y_k)|Q_k] - E[R(Y_j)|Q_j] + \alpha_k + P_i\left(\Phi\left(\frac{c - Q_j}{\sigma_\varepsilon^g}\right) - \Phi\left(\frac{c - Q_k}{\sigma_\varepsilon^g}\right)\right)}^{E[R(Y_k)|Q_k] - E[R(Y_j)|Q_j] + \alpha_k} f^g\left(\alpha_j, Q_j\right) d\alpha_j dQ_j > 0 \tag{13}$$

As $\sigma_\varepsilon^g \to \infty$, expression 13 converges to 0. For fixed $\sigma_\varepsilon^g < \infty$, it is the density expression which motivates condition (ii) in proposition 2. As this holds for any possible (α_k, Q_k) s.t. $k >_0 j \; \forall j \in \mathcal{G}$, proposition 2 is satisfied.

A.1.3 Proposition 3

If $j >_p k$ and $k >_0 j$, then as before we are guaranteed that $\Phi\left(\frac{c - Q_j}{\sigma_\varepsilon^{g_j}}\right) < \Phi\left(\frac{c - Q_k}{\sigma_\varepsilon^{g_k}}\right)$. If $\sigma_\varepsilon^W = \sigma_\varepsilon^M$ then we are guaranteed by a similar proof as proposition 2 that referral quality is non-decreasing in performance pay and that an analogous proposition holds. If there is a performance premium in at least one of the genders G, then by proposition 2 we are guaranteed that $N_G > 1$ and $\sigma_\varepsilon^G < \infty$ which, again under proposition 2, suggests that the choice of referrals in the unrestricted network should be improving in quality so long as density condition (ii) is met in both genders. However, if $j \in \mathcal{M}_i$ and $k \in \mathcal{F}_i$ and $\sigma_\varepsilon^M < \sigma_\varepsilon^F$, then we are no longer guaranteed that $Q_j > Q_k$ since $\exists Q_k < Q_j$ s.t. $\Phi\left(\frac{c - Q_j}{\sigma_\varepsilon^M}\right) < \Phi\left(\frac{c - Q_k}{\sigma_\varepsilon^F}\right)$. As a result, there remain monotonicity predictions on qualification probability, but not on referral quality.

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 and Figure

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 the 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: Scarcity Bounds

	Number of Women (N_F)		
	1	2	3
A. Fixed Fee			
Implied Π^F	0.9	0.684	0.536
Maximum Implied ψ if $\Pi^M = 0.001$	0.719	0.627	0.589
Implied ψ if $\Pi^M = 0.1 * \Pi^F$	0.709	0.619	0.582
Implied ψ if $\Pi^M = 0.25 * \Pi^F$	0.693	0.605	0.572
Implied ψ if $\Pi^M = 0.5 * \Pi^F$	0.658	0.579	0.552
Implied ψ if $\Pi^M = 0.75 * \Pi^F$	0.6269	0.546	0.528
Implied ψ if $\Pi^M = 0.1 * \Pi^F$	0.5	0.5	0.5
Actual ψ: 0.75			
B. Performance Pay			
Implied Π^F	0.75	0.5	0.415
Maximum Implied ψ if $\Pi^M \geq 0.001$	0.649	0.581	0.564
Implied ψ if $\Pi^M = 0.1 * \Pi^F$	0.64	0.575	0.559
Implied ψ if $\Pi^M = 0.25 * \Pi^F$	0.6253	0.565	0.55
Implied ψ if $\Pi^M = 0.5 * \Pi^F$	0.596	0.546	0.535
Implied ψ if $\Pi^M = 0.75 * \Pi^F$	0.563	0.525	0.519
Implied ψ if $\Pi^M = 0.1 * \Pi^F$	0.5	0.5	0.5
Actual ψ: 0.78			

Appendix Table 3: 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.

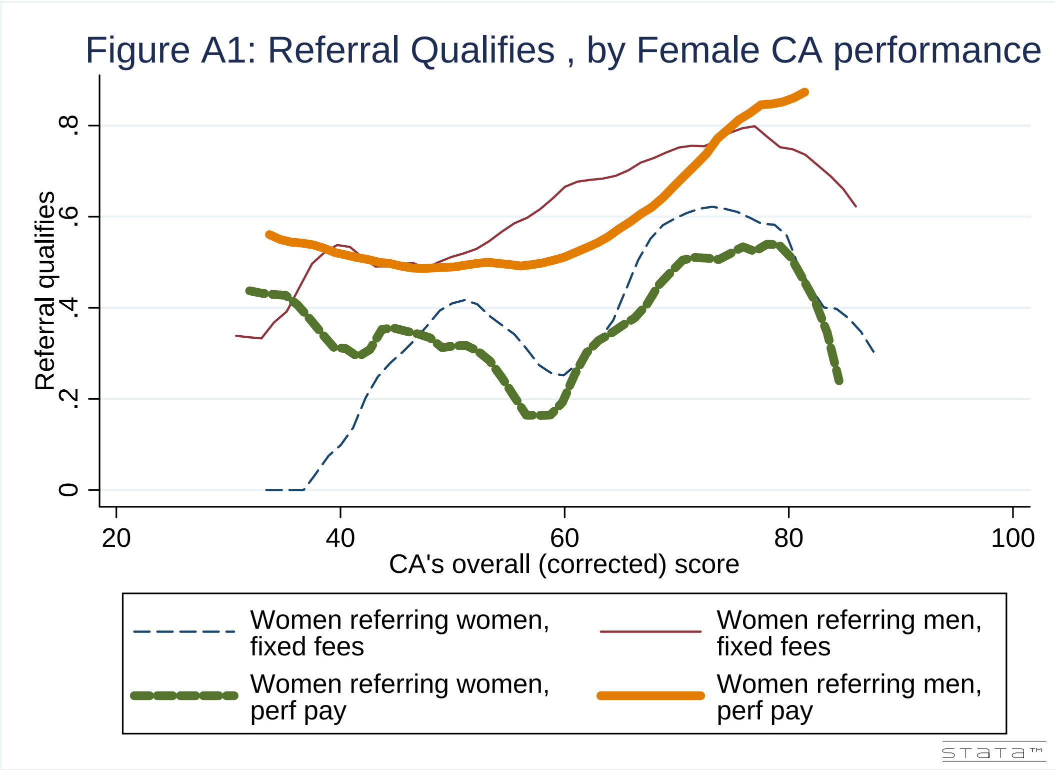
Figure A1: Referral Qualifies , by Female CA performance
Referral qualifies
0
.2
.4
.6
.8
20
40
60
80
100
CA's overall (corrected) score
Women referring women, fixed fees
Women referring men, fixed fees
Women referring women, perf pay
Women referring men, perf pay