

Do Job Networks Disadvantage Women?

Evidence from a Recruitment Experiment in Malawi *

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

August 2015

Abstract

We use a field experiment to show referral-based hiring has the potential to disadvantage qualified women, highlighting another potential channel behind gender disparities in the labor market. Through a recruitment drive for a firm in Malawi, we look at men's and women's referral choices under different incentives and constraints. We find that men systematically refer few women, despite being able to refer qualified women when explicitly asked for female candidates. Performance pay also did not alter men's tendencies to refer men. In this context, women refer lower quality candidates than men - limiting their ability to offset men's behavior.

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

*We thank IPA-Malawi field staff for dedicated and careful implementation, and Sam Arenberg for research assistance. We also thank participants at numerous seminars and conference audiences for helpful comments. We are grateful to Pascaline Dupas, Matt Jackson, Seema Jayachandran, Cynthia Kinnan, Ethan Ligon, Kaivan Munshi, Imran Rasul, Mark Rosenzweig, Aloysius Siow, among others, for helpful comments and discussions. All errors are our own.

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

[‡]University of Washington and Innovations for Poverty Action. Email: nkeleher@uw.edu

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

Group, 2011). A large portion of the literature in economics has focused on labor market discrimination (taste-based or statistical) or differences in human capital accumulation as reasons for the gender gap in earnings (Altonji and Blank, 1999).¹ Another possibility is that hiring processes themselves disadvantage women. We conduct a field experiment generating a list of qualified candidates for a job in which men and women regularly compete in order to ask whether the use of referrals could disadvantage women in the labor market, and examine some of the conditions which could encourage this potential source of bias.

A large fraction of jobs - up to 50% - are attained through informal channels, including employee referrals (Bewley, 1999; Ioannides and Loury, 2004). Firms use referrals to overcome a number of frictions in the labor market, including incomplete information while screening applicants² and to induce effort (Heath, 2015; Kugler, 2003). This paper highlights a potential cost of firms using informal networks to address labor market imperfections: groups who are socially distant from the employed may be left out. If employees make referrals, then unemployed people in networks with few employed contacts receive fewer job referrals and fewer jobs, thus exacerbating inequality. This potential is described theoretically in Calvo-Armengol and Jackson (2004), though there is limited empirical evidence for this channel at present.³ An extensive literature in sociology

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

²See Beaman and Magruder (2012); Brown et al. (2015); Burk et al. (2015); Dustmann et al. (2015); Hensvik and Skans (2013); Pallais and Sands (2014); Montgomery (1991).

³Mortensen and Vishwanath (1994) also show theoretically that network-based job information dissemination can disadvantage women, even if men and women are equally productive but men have a higher contact probability.

(reviewed in McPherson, Smith-Lovin, and Cook, 2001) suggests that networks, particularly workforce networks, are quite gender homophilous. This observation suggests that the use of referrals could put women at a disadvantage. On the other hand, if female applicants are, on average, weaker along easy-to-observe dimensions - like job experience - network screening may succeed in identifying the women who have strong hard-to-observe but productive characteristics. These two channels taken together suggest that whether the use of referrals results in better or worse employment outcomes for women is theoretically ambiguous, and likely depends on whether employees use referrals as an opportunity to benefit their network, or as an opportunity to benefit their employer.

In this study, we used a competitive recruitment drive conducted by a research organization in Malawi, Innovations for Poverty Action (IPA-Malawi), as an opportunity to document whether women may be advantaged or disadvantaged through the use of referrals. Moreover, we examine how such an advantage or disadvantage is affected by experimental variation in the incentives in the referral choice process. IPA-Malawi⁴ advertised positions for survey enumerators using a traditional method of posting flyers. 38% of the applicants who respond to these fliers were female, suggesting that both men and women are interested in survey enumerator jobs in Malawi. At the conclusion of a half-day application process, candidates were asked to refer a friend or relative to apply for the position. Conventional applicants were also offered a finder's fee.

The referral process was cross-randomized along three dimensions. First, candidates were either told that they may refer a woman, that they may refer a man, or that

⁴IPA-Malawi was interested in exploring whether referrals could increase the pool of qualified female applicants specifically and qualified applicants in general since the firm needs female enumerators when surveying women about sensitive questions, such as family planning practices.

they may refer a person of either gender. Second, their finder's fee was randomly selected to be a fixed fee of either 1000 or 1500 Malawi Kwacha (MWK; \$1=153 MWK, 2011) 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). The full performance incentive is approximately a day's wage for an enumerator.⁵ Third, applicants 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).

In our setting, qualified female candidates are disadvantaged by the use of referrals in this hiring drive. Among the conventional applicants (CAs) who were allowed to choose either gender for a referral, only 30% of referrals are women. This is statistically significantly lower than the fraction of women who apply through traditional recruitment channels. The low number of women referred is driven largely by male candidates: when given the choice, 77% of men referred other men. While men systematically refer other men for these positions, they are, in fact, able to refer women when their choices are constrained. Men make referrals at identical rates when required to refer either women or men. Moreover, women referred by men whose options are limited to women referrals are also just as likely to qualify for the short-list as the men referred by men who can only refer other men. These two facts suggest that men's systematic choice of male referrals is not the result of men simply not knowing any women they could refer. Men are capable of referring women, but usually choose not to.

⁵The daily wage for an enumerator at the time was typically MWK 1875, though enumerators working outside the cities would also earn a per diem worth approximately the same as the salary.

Given that women seem less biased in the gender of their referrals than men are, we ask whether a firm can rely on their female employees to offset men's behavior. Even though the women who apply themselves are almost as likely to qualify for the short-list as male applicants (48% compared to 54%), our experiment indicates that using women to make referrals would not be effective in this context because women systematically refer people who were less likely to qualify. While women refer other women at about the same rate as women apply through the traditional method, when compared with male CAs, a female CA is 18 percentage points less likely to refer someone who qualifies. Since men systematically refer men, and women are not very likely to refer someone who qualifies, the net result is that few qualified women get referred to the firm.

Our experimental design allows us to learn about the types of referral contracts which may result in better outcomes for firms who want to hire more women as well as firms who are indifferent to the gender of their workers but want to find the most able workers possible. For example, if the bias towards referring men were driven by taste-based discrimination, then economic incentives should diminish discriminatory behavior, resulting simultaneously in more women and better workers being hired. On the other hand, if this referral bias were driven by a difference in (actual or perceived) ability of women, we may expect it to be exacerbated in the presence of referral performance incentives. An unbiased firm who prioritizes high quality workers over diversity may then have a good reason to prefer hiring male referrals, depending on the accuracy of workers' information. In practice, we find that men who could choose to refer anyone referred a similar fraction of women in both fixed fee and performance pay treatments (23% vs 21%). This suggests that among a range of contracts similar to those considered

here, increasing explicit or implicit incentives to identify high ability workers may not improve outcomes for women. We also find that men who could refer anyone do not on average bring in higher ability workers when given explicit financial incentives, suggesting that (again, in the range of contracts we consider) financial incentives may not always improve hiring outcomes for a firm. Women CAs are also not responsive to the financial incentives. We do, however, find that men who were asked to refer other men recruit better candidates when offered the performance incentive. The best quality candidates therefore come from male CAs who are asked to recruit other men and are properly incentivized. We offer a model in the appendix which can explain these results. While this is only one experiment with one firm, this result suggests that even an unbiased firm may have an incentive to allow bias in their employees' referral patterns in order to maximize the quality of their candidate pool.

One weakness of our experimental design is that we ask job applicants, and not existing employees, to make referrals.⁶ Candidates may be leery to refer high quality candidates because they do not want to compete with them for the (numerous) available positions. This is a key difference with our experiment compared to the way job recruitment is usually conducted. However, in many cases in the real world, employees will compete with their own referrals. Some employees work directly with the people they refer as in Heath (2015), and many employees refer individuals who will work at the same level in the company as they do (Brown et al., 2015). Individuals who work at the same level or on the same production team as their referrals will compete with them

⁶Since most IPA employees are contractors, hired for individual surveys, there are not many full-time, permanent staff to use for such an experiment.

for promotions. Our set up of using candidates to make referrals, while generating some threats to external validity, allows us to experimentally vary the extent that candidates are competing with their own referrals. To do this, we use the treatment where some individuals were told that they would qualify for the short-list if they were in the top 50% of applicants while the others were informed that qualification was based on a fixed threshold. The former would then deduce that they were competing with their referrals, while those in the fixed threshold group would only be competing with referrals after the short-list stage. This should have raised the salience of competition. We find no evidence that increasing the salience of competition decreases the quality of the referrals made by either men or women - though the estimates are quite noisy.

The experiment provides cautionary evidence that women could fare worse than men when firms use social networks to make hires. Future research should explore the robustness of these results in other contexts.

2 Experimental Design

2.1 Setting and Overview

Women in Africa are more likely to work 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. Enumerator jobs are relatively well paid but offer only short-term contract work, typically for a few months at a time.⁷ 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 asking sensitive questions).⁸ For this experiment, we introduced incentives for conventional job applicants (CAs) to make referrals during IPA's recruitment sessions in the two main Malawian cities, Blantyre and Lilongwe. There were a total of 55 sessions (including CAs and referrals) in the two cities, over 31 days from late June 2011 through August 2011. We had two interview sites within Lilongwe and one in Blantyre. After the initial conventional applicant session at each site, CAs and referral sessions were interspersed with one another overtime. In some recruitment sessions, we interviewed both CA and referral applicants. However, CAs

⁷See Godlonton (2014) for a fuller description of the data collection industry in Malawi. According to the 2010/11 Integrated Household Survey, Godlonton (2014) states that the typical urban man aged 18-49 who completed secondary school earned \$4.75 per day. IPA pays \$6.50 plus \$12 in per diem per day.

⁸Informal 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.

were never at the session at the same time as the person they referred.

To recruit conventional applicants, IPA posted fliers indicating a hiring drive at a number of visible places in urban areas. The posters included information on the minimum requirements for IPA enumerators, the dates and times of the recruitment sessions, and a solicitation to bring a CV and certificate of secondary school completion (MSCE). Minimum requirements to be hired for an enumerator position are: a secondary certificate, fluency in the local language (Chichewa), and English reading and oral comprehension. Candidates with data collection experience, good math skills, and basic computer skills are given preferential review. Participants then attended an interview session, where they submitted their CV and were registered with a unique applicant number. Participants were limited to those individuals who had never worked for IPA. 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 standard test that IPA had previously used and a practical test which served as a condensed version of a skills assessment that IPA had previously used to evaluate enumerators.⁹ Participants were given one of two distinct written tests. Each test consisted of several math problems, Raven's matrices, English skills assessment, job comprehension component, and a computer skills

⁹The 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.

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.¹⁰ 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.¹¹

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

2.3 Referral Instructions and Experimental Treatment Arms

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. The experimental treatment arms were motivated by the simple model in Beaman and Magruder (2012). In the model, a CA chooses who to refer by maximizing (i) a social benefit they get from the network member they refer and (ii) a benefit they get from the

¹⁰All 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.

¹¹The two sets of written tests and the two versions of the practical exam were randomly distributed to applicants to limit cheating. We wanted to minimize the the ability of CAs - particularly those in performance pay treatments - to simply tell referrals the correct answers. We do not observe any significant differences between CAs and referrals treated with the same or different versions of the test.

firm, which may depend on the ability of the person referred (which may not be perfectly observed by the CA). As long as the distributions of social benefits and the ability of network members are not perfectly correlated, CAs face a tradeoff between maximizing social benefits and maximizing the ability of the person referred. A financial incentive which depends on the ability of the referral can induce CAs to choose a network member who is higher ability than is referred in a fixed fee treatment, and in that case the CA on average foregoes social benefits to capture the higher payment from the firm. However, CAs will only bring in better referrals in the performance pay incentive if they have sufficiently accurate information about people in their social network. In the appendix we further develop this framework to allow for heterogeneity in the men versus women in a given CA's network. There are three key types of heterogeneity between male and female network members we explore: first, the precision of the signal about ability; second, the distribution of social benefits; and third, the number of network members who are male vs female. We designed the experiment with variation in the financial terms offered to CAs for a referral (fixed vs performance pay) and cross-randomized whether we asked CAs to refer a man, a woman or had a choice. We return to the model in section 3.3.

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. Along with the letter, the applicant received a card to give to his referral, and the referral used the card to gain admission to the interview site. The card is also how we track referrals to particular CAs, as we did not solicit names directly from the CAs; instead we wanted CAs to be able to talk with potential referrals

before making their referral choice. 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 conventional applicant who qualified for a payment was informed and given payment in a sealed envelope.¹² Most CAs did not know their score or whether they qualified before making their referral.

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 and the Practical Component Z-score is significant at the 10% level for both men and women CAs. For women CAs, there is also a baseline difference in MSCE math scores at the 5% level.

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, at the 10% level. 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.

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

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 Empirical Results

3.1 Number of Women Recruited

Figure 2 documents the primary result of this paper. While 38% of applicants themselves were women, only 30% of referrals are women when we allow CAs to choose which gender to refer. This difference is significant at the 5% level.¹³ This difference in application rates is driven by men systematically referring other men 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 is significant at the 1% level, as shown in column (4) of panel C in 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.¹⁴ 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 (excluding the tails where there are very few observations). Table 1 also shows that women’s referrals are less likely to qualify than men’s, which means that women’s referrals are also unlikely to be

¹³Table 1 Panel B shows the equivalent figures for the specific subset of CAs randomized into the either-gender treatments: in this subsample the pattern is even more striking as 40% of CAs are women.

¹⁴In both cases, the sample is restricted to CAs who have the choice of which gender to refer.

qualified women. We discuss this result in greater detail in section 3.4. Here we examine implications of referral systems for the pool of qualified candidates. 35% of the pool of qualified CAs are women. Of the pool of qualified referrals, only 28% are women.¹⁵ Therefore the same trend in getting fewer women through referrals than through the traditional recruitment method continues if you look at only qualified applicants. Overall, we conclude that the use of referral systems disadvantages women in this context.

One possible concern with these findings is that at each of the three interview sites, we started interviewing conventional applicants before the referrals (in order to have candidates to make referrals). We do not want to conflate a possible reduction in the number of women applicants over time with the difference in the number of women recruited through different hiring channels. Therefore we designed the experiment to have oscillating rounds in which we interviewed CAs and referrals so as to minimize this problem. On many days we interviewed both referrals and CAs. Perhaps as a result of this design, this concern (while *ex ante* quite serious) appears to have little empirical support. We can document trends in the characteristics of people who remain interested in the job by looking at how CA characteristics change with the number of recruitment sessions held at each site. Appendix Figure A1 documents that, if anything, the fraction of women among conventional applicants increased over time at each site. Appendix Figure A2 also shows that the quality of women applying as conventional applicants is variable but largely increasing over time. By contrast, the qualification rate among men is largely flat. There is little evidence then that qualified women overall were unavailable

¹⁵This difference is not statistically significant at conventional levels ($p=.19$), but this is likely because the sample size is halved since the effect size is quite similar.

after the initial interview session.

3.2 Are Qualified Women Not in Men’s Networks?

3.2.1 Rates of Referring Women

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

The experiment 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 referable 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. Making a referral means that a referral actually showed up to an interview session. 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 dummy variables for each CA recruitment day.

Columns (1)-(2) of Table 2 presents this analysis, where treatments CAs who 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 89%; point estimates suggest that if these men are instead given the performance contract, return rates fall to 74%.

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.

3.2.2 Performance of Female Candidates Referred by Male CAs

Perhaps men know other women but choose not to refer women because they are not well qualified for the position.

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.

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 dummy variables for each CA recruitment day. 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. Consistent with Figure 4, Column (3) shows that the probability of qualifying for the short-list is the same whether the referral had to be a man, woman or the CA had the choice.

3.3 Financial Incentives

Men appear to be capable of referring women but typically choose not to. In this section we explore what changes in the contract terms does to referral patterns. We observe a variety of incentives offered employees in labor markets around the world, including direct financial incentives like we offer in our experiment. The firm in Brown et al. (2015) provides a small cash bonus if a referred worker stays for at least 6 months; in the Burk et al. (2015) data, the trucking firm provided a bonus to referred workers who stay a for at least 3 or 4 months, and the cell center firms provide no bonus at some locations but at other locations offer a bonus of about \$50 if the referred worker stays for a minimum amount of time (between 30 and 90 days). This is in addition to likely non-monetary benefits that a worker would receive if they bring in a good worker, ranging from their reputation with their boss to getting to work with a friend.

Financial incentives which are contingent on referral quality may affect both the quality of applicants brought to the firm but also the gender mix. For example, if men don't refer women because of taste-based discrimination, then economic incentives should diminish discriminatory behavior, resulting simultaneously in more women and better workers being hired. On the other hand, if men have beliefs (either founded or unfounded) that women are of lower ability, we may expect even fewer women referred in the presence of referral performance incentives.

We find no evidence of the performance incentives favoring men or women in our experiment. Comparing panels D and E of Table 1 shows that male CAs refer only marginally fewer women (21% vs 23%) 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.

Column (4) of Table 2 shows that male CAs in the male-gender treatment refer significantly better candidates when given a performance pay incentive: candidates are approximately 25 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. This is consistent with Beaman and Magruder (2012), who argue that social benefits can 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.¹⁶ 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. However, column (2) of Table 2 showed that fewer CAs made referrals in the performance pay-male referral treatment. There may be selection of CAs who make referrals affecting the estimate in column (4), and this complicates the interpretation of the finding. We discuss this possibility in section 3.3.1. Ignoring the potential selection issue for now, we conclude that male CAs have useful information for employers about men, and the price of eliciting the information is not prohibitively high.¹⁷

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 that among CAs in performance pay treatment, the referred men outperform the referred women: 62% of referrals qualify in the male-only treatment vs 44% in the female-only treatment. It may be that men cannot identify which women are well qualified. Or it may be costlier for men to get high quality women to apply for the job (so a larger incentive is needed). In either case, the firm gets the highest quality candidates by asking male CAs to refer other men and providing a performance incentive.

¹⁶Reputational 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.

¹⁷A component of the overall score includes the subjective assessment of the candidate (feedback score) by the IPA employee. This represents about 10% of the overall score. Since this component is subject to potential biases related to gender, we have also re-done the analysis excluding that component. To do so, we re-calculate who the median candidate is to determine who would have qualified under this alternative scoring mechanism. We find very similar results. They are omitted for brevity but available from the authors upon request.

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 Perf Pay*Either 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. This is a surprising result. We offer a model in the appendix which provides a potential explanation. In the model, a CA is choosing a referral to maximize both a social benefit they get from the network member they refer and a benefit they get from the firm, which may depend on the ability of the person referred. We assume that CAs get a noisy signal of the ability and a precise signal of social benefits of each network member. Under pretty weak assumptions, the CA will face a tradeoff between choosing a network member who offers a high social benefit and a network member who is high ability. The model can generate the pattern in response to performance pay across genders we observe in the data if men may get a noisier signal of the ability of women in their network compared to the men in their network. CAs may then prefer to refer a woman with high social benefits and uncertain ability to a (known) high ability man who gives low social benefits, when given the choice. This could lead to no change in the average ability of referred candidates between the fixed and performance pay treatments.

The results suggest that there is the potential for firm incentives to increase bias in favor of men: after all, the highest quality referrals overall are men's incentivized referrals of men. Higher stakes incentives may induce greater male bias, as the return to getting a high-probability high quality candidate increases. 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 forgo social benefits and refer even more men than we observed in our experiment.

Table 3 looks at the sub-components of the overall score. It shows that men referred by men under performance pay do statistically significantly better on the computer knowledge part of the exam, on feedback points¹⁸ 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.

3.3.1 Interpreting Attrition

In section 3.2.1, we saw that male CAs were more likely to make a referral in the presence of fixed fees than performance pay.¹⁹ 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. We note that this interpretation would remain consistent with the conclusions of this study, including the potential importance of differential information about men and women suggested by

¹⁸Feedback points are a subjective measure on a scale of 1 to 10 of how well the candidate did on the practical component of the test, as judged by the supervisor who was conducting the practical test.

¹⁹In Section 3.4, we will also note that female CAs responded similarly.

the model in the appendix.²⁰ That said, Figure 8 suggests that selective attrition is not the only factor behind the results in table 2. Though we do not have statistical power to detect differences across the distribution, Figure 8 provides suggestive evidence that among men, there is assortative matching in ability between the CA and their referral under fixed fees. If selective attrition were fully responsible for the results in Table 2, we might expect to see differential selection among CAs in terms of what ability men make a referral but not among the referrals made by men of a particular ability level. However, Figure 8 also suggests that the performance premium exists throughout the entire distribution of CA test scores.²¹

3.4 Women CAs' Behavior

Figure 2 showed that women refer other women about 43% of the time, which is statistically indistinguishable from the 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, it is possible 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

²⁰If 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 under performance pay, 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.

²¹Since the referral patterns are similar across the entire distribution of CA ability, we also have more confidence that the results extrapolate to other contexts where only existing employees make referrals.

than either the pool of traditional applicants or men’s referrals. Figure 5 and Table 1 reveal that, on average, women refer people who are thirteen percentage points less likely to qualify for a position than men’s referrals, a difference which is marginally statistically significant. Figure 6 plots the raw data to transparently show that the people referred by women are less likely to qualify than the people referred by men across the range of CA abilities.

Figure 7 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 stochastically dominates the distribution of referred women, with the Kolmogorov-Smirnov test rejecting the distributions being the same at the 10% 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 4 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 5 shows referral performance disaggregated by component for women CAs. When we provide performance pay, women refer women with better English skills and who solve more ravens matrices correctly, 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, including being less likely to have tertiary education. Women appear to respond to performance pay and 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).²² A plausible alternative explanation is that women need a larger performance pay premium in order for them to refer higher quality candidates. We have no direct evidence on this possibility, but there is 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.

4 Competition

Another possible reason women refer low ability individuals is aversion to competition (despite the firm's motivation of wanting to hire more women) as suggested in Flory

²²Appendix figure A3 suggests that there is little evidence of female CAs responding to the performance pay incentive at any point in the CA performance distribution, though we do not have power to perform valid statistical tests.

et al. (2010) and Niederle and Vesterlund (2007).²³ 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)²⁴, competition on the job may actually be stronger.

Nevertheless, if women CAs are concerned about the competitive threat their referrals pose, they may choose to either forgo the finder's fee (and not make a referral) or refer someone who is unlikely to qualify. We do not observe the former, as the referral rate is almost identical among women CAs and male CAs. However, the latter is consistent with the results presented in Table 4: in unrestricted treatments, women refer poor quality men and women. However, several additional pieces of evidence seem inconsistent with the competition aversion hypothesis. Figure 6 shows suggestive evidence 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. Second, Tables 4 and 5 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.

In order to directly look at the role of competition in referral decisions, we also experimentally varied how salient competition was to CAs. CAs were told the qualifi-

²³Niederle 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.

²⁴On 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.

cation 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 6 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, it provides suggestive evidence that, on average, competition at the qualification stage is unlikely to be driving our main results.

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 examine the possibility that women need not always shy away from competing with men in particular as in Niederle and Vesterlund (2007) and may be more averse to competition with women in some settings.

5 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) 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.²⁵ Using an experiment designed around a recruitment drive

²⁵Moreover, 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

for real-world jobs, we provide evidence that the use of referral systems can put 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 is not driven solely by not knowing other women or knowing only low-quality women. We also document that in this context at least, using women to make referrals is similarly unsuccessful at identifying high ability female workers. While women CAs in our experiment refer women more often than men CAs, they refer people (and particularly women) who are not very likely to qualify for positions. This result suggests that the role of job networks in the labor market could contribute to persistent gender gaps in labor market outcomes.

As with any 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 very different from 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)²⁶ and of the mechanisms creating this segregation (Fernandez and Sosa, 2005; Tassier and Menczer, 2008).

²⁶In that context, men are the disadvantaged group, who are similarly less likely to receive referrals.

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). The disadvantage of women in referral-based hires we observe with one firm in Malawi may well exist in these other contexts as well.

References

- Altonji, J. and R. Blank (1999). Race and gender in the labor market. *Handbook of Labor Economics* 3, 3143–4259.
- Arbache, J. S., A. Kolev, and E. Filipiak (Eds.) (2010). *Gender Disparities in Africa's Labor Market*. World Bank Publications.
- Beaman, L. and J. Magruder (2012). Who gets the job referral? Evidence from a social networks experiment. *American Economic Review* 102, 3574–3593.
- 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.
- Brown, M., E. Setren, and G. Topa (2015). Do informal referrals lead to better matches? Evidence from a firm's employee referral system. Forthcoming.
- Burk, S., B. Cowgill, M. Hoffman, and M. Housman (Forthcoming 2015). The value of hiring through employee referrals. *Quarterly Journal of Economics*.
- Calvo-Armengol, A. and M. O. Jackson (2004). The effects of social networks on employment and inequality. *American Economic Review* 94, 426–454.
- Doeringer, P. B. and M. J. Piore (1971). *Internal Labor Markets and Manpower Analysis*. M.E. Sharpe, Inc.
- Dustmann, C., A. Glitz, and U. Schoenberg (2015). Referral-based job search networks. *Review of Economics Studies* Forthcoming.

- 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.
- Flory, J., A. Leibbrandt, and J. List (2010). Do competitive workplaces deter female workers? A large-scale natural field experiment on gender differences in job-entry decisions. *NBER Working Paper 16546*.
- Godlonton, S. (2014). Employment risk and job-seeker performance. *Working Paper, IFPRI Discussion Paper #1332*.
- 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.
- Heath, R. (2015). Why do firms hire using referrals? Evidence from Bangladeshi garment factories. *Mimeo, University of Washington*.
- 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.
- Kugler, A. (2003). Employee referrals and efficiency wages. *Labour Economics* 10, 531–556.
- 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 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.
- Pallais, A. and E. Sands (2014). Why the referential treatment? Evidence from field experiments on referrals. Working paper, Harvard University.
- 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 (2011). *World Development Report 2012: Gender Equality and Development*. World Bank Publications.
- World Bank Group and others (2010). *The Education System in Malawi*. World Bank Publications.

Table 1: Gender Distributions of CAs and Referrals

	(1)	(2)	(3)	(4)
	All CAs	Male CAs	Female CAs	Diff: <i>p</i> value
<u>A. OP Characteristics</u>				
Fraction of OPs	1	62%	38%	
OP is Qualified	52%	54%	48%	0.074
N	816	506	310	
<u>B. OP Characteristics: Either Gender Treatments</u>				
Fraction of OPs	1	60%	40%	
OP is Qualified	55%	59%	49%	0.091
N	266	159	107	
<u>C. Referral Characteristics: Either Gender Treatments</u>				
Referral is Female	30%	23%	43%	0.002
Referral is Qualified	48%	53%	40%	0.072
Referral is Qualified Male	34%	41%	23%	0.005
Referral is Qualified Female	14%	11%	17%	0.209
N	222	133	87	
<u>D. Referral Characteristics, Fixed Fee Treatments</u>				
Referral is Female	32%	23%	43%	0.017
Referral is Qualified	50%	57%	39%	0.042
Referral is Qualified Male	34%	43%	21%	0.010
Referral is Qualified Female	16%	14%	18%	0.580
N	133	77	56	
<u>E. Referral Characteristics, Perf Treatments</u>				
Referral is Female	29%	21%	42%	0.043
Referral is Qualified	45%	46%	42%	0.691
Referral is Qualified Male	34%	39%	26%	0.210
Referral is Qualified Female	10%	7%	16%	0.192
N	87	56	31	

Table 2: Male CA's Referral Choices

	Made a Referral		Referral Qualifies		
	(1)	(2)	(3)	(4)	
Female Treatment	-0.004 (0.038)	-0.004 (0.050)	-0.022 (0.060)	0.085 (0.078)	
Either Gender Treatment	0.014 (0.040)	-0.052 (0.052)	0.054 (0.062)	0.202 (0.080)	**
Performance Pay		-0.148 (0.056)	***	0.252 (0.090)	***
Perf Pay * Female Treatment		0.004 (0.076)		-0.265 (0.122)	**
Perf Pay * Either Treatment		0.152 (0.079)	*	-0.362 (0.125)	***
Observations	506	506	429	429	
Mean of excluded group	0.840	0.892	0.485	0.398	

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	-0.030 (0.066)	0.022 (0.072)	0.003 (0.135)	-0.109 (0.199)	-0.086 (0.186)	0.075 (0.358)	1.023 (0.620)	3.130 (0.978)
Either Gender Treatment	0.025 (0.068)	0.028 (0.075)	-0.028 (0.140)	0.074 (0.206)	0.062 (0.192)	0.574 (0.370)	0.973 (0.644)	1.679 (1.015)
Performance Pay	0.082 (0.076)	0.039 (0.084)	0.134 (0.157)	0.042 (0.231)	0.153 (0.216)	0.915 (0.415)	0.284 (0.714)	2.056 (1.126)
Perf Pay * Female Treatment	-0.071 (0.103)	0.060 (0.113)	-0.313 (0.213)	-0.079 (0.313)	-0.186 (0.293)	-0.815 (0.562)	-1.094 (0.962)	-2.702 (1.517)
Perf Pay * Either Treatment	-0.175 (0.106)	-0.066 (0.116)	-0.069 (0.218)	-0.198 (0.322)	-0.171 (0.300)	-0.788 (0.577)	-1.370 (0.991)	-2.972 (1.568)
Mean of Dep Variable	0.235	0.692	2.033	7.748	1.524	4.914	15.687	27.109
SD	(0.425)	(0.462)	(0.867)	(1.285)	(1.220)	(2.361)	(3.954)	(6.160)
Observations	425	429	429	429	429	429	422	421

Notes

- 1 The dependent variable is listed in the column heading.
- 2 All specifications include CA visit day dummies.

Table 4: Female CA's Referral Choices

	Made a Referral		Referral Qualifies	
	(1)	(2)	(3)	(4)
Female Referral Treatment	-0.055 (0.054)	-0.042 (0.074)	-0.214 *** (0.079)	-0.219 ** (0.106)
Either Gender Treatment	0.017 (0.055)	-0.024 (0.071)	-0.214 *** (0.078)	-0.244 ** (0.100)
Performance Pay		-0.113 (0.080)		-0.018 (0.117)
Perf Pay * Female Treatment		-0.013 (0.111)		0.012 (0.162)
Perf Pay * Either Treatment		0.086 (0.110)		0.082 (0.161)
Observations	310	310	254	254
Mean of Excluded Group	0.821	0.852	0.590	0.609

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 5: Screening of Female CAs on Different Characteristics

	Survey Experience (1)	Tertiary Education (2)	Math Score (3)	Language Score (4)	Ravens Score (5)	Computer Score (6)	Practical Exam Score (7)	Feedback Points (8)
Female Referral Treatment	0.034 (0.085)	0.187 * (0.101)	-0.401 ** (0.198)	-1.124 *** (0.314)	-0.584 ** (0.251)	-0.491 (0.497)	0.803 (0.894)	2.318 * (1.256)
Either Gender Treatment	0.054 (0.080)	0.038 (0.096)	-0.298 (0.188)	-0.184 (0.298)	-0.384 (0.238)	-0.126 (0.471)	0.147 (0.846)	1.468 (1.188)
Performance Pay	0.266 *** (0.094)	0.146 (0.113)	-0.458 ** (0.220)	-0.439 (0.349)	-0.249 (0.279)	0.416 (0.553)	2.111 ** (1.008)	1.956 (1.416)
Perf Pay * Female Treatment	-0.317 ** (0.129)	-0.343 ** (0.155)	0.497 (0.304)	1.200 ** (0.481)	0.717 * (0.384)	-0.096 (0.761)	-1.695 (1.382)	-2.250 (1.942)
Perf Pay * Either Treatment	-0.226 * (0.129)	-0.062 (0.155)	0.457 (0.302)	0.559 (0.479)	-0.042 (0.382)	-0.253 (0.758)	-1.697 (1.375)	-4.310 ** (1.932)
Mean of Dep Variable	0.194	0.642	1.831	7.524	1.614	4.724	16.093	27.528
SD	(0.396)	(0.480)	(0.932)	(1.508)	(1.196)	(2.393)	(4.048)	(5.812)
Observations	253	254	254	254	254	254	248	248

Notes

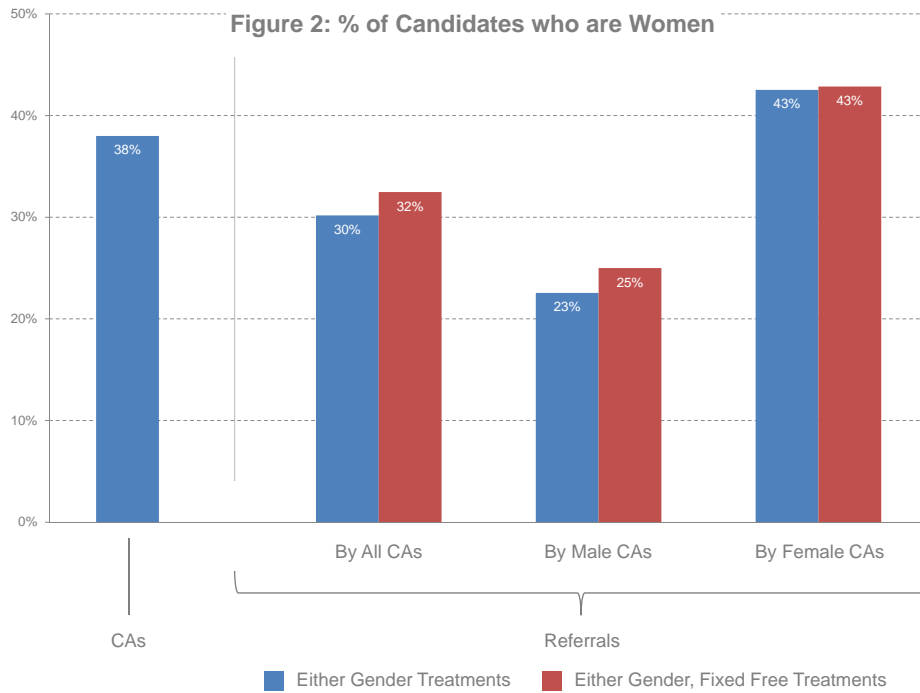
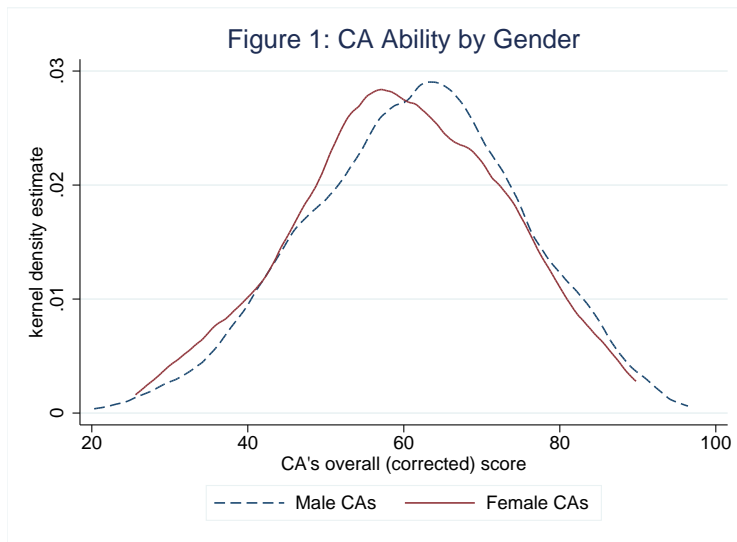
- 1 The dependent variable is listed in the column heading.
- 2 All specifications include CA visit day dummies.

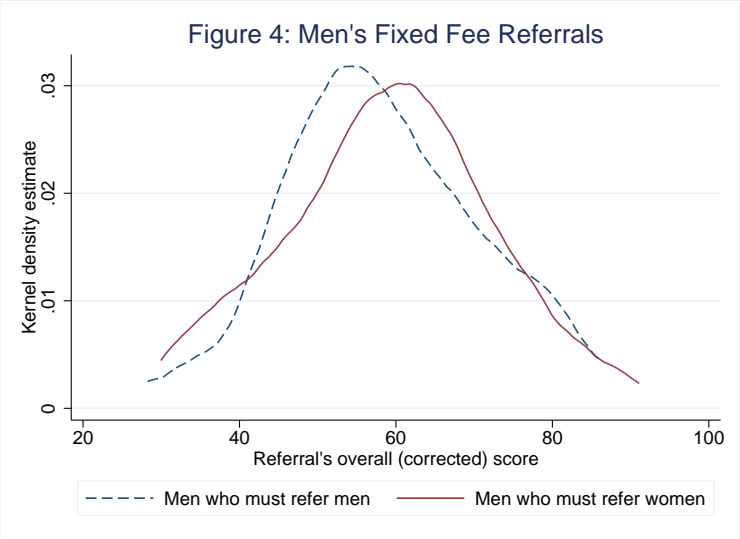
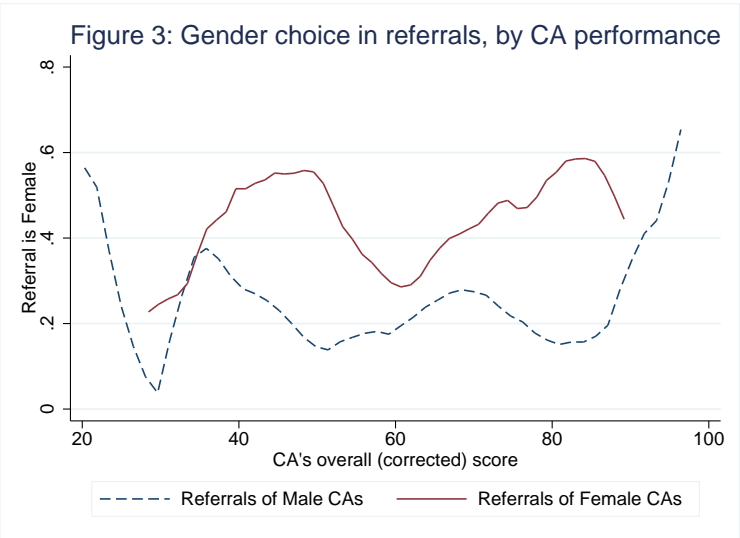
Table 6: Competition Incentives in the Fixed Fee Treatments

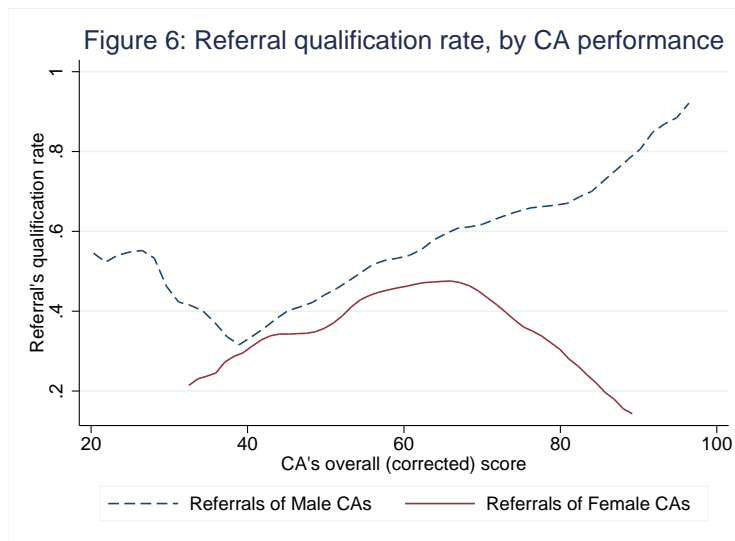
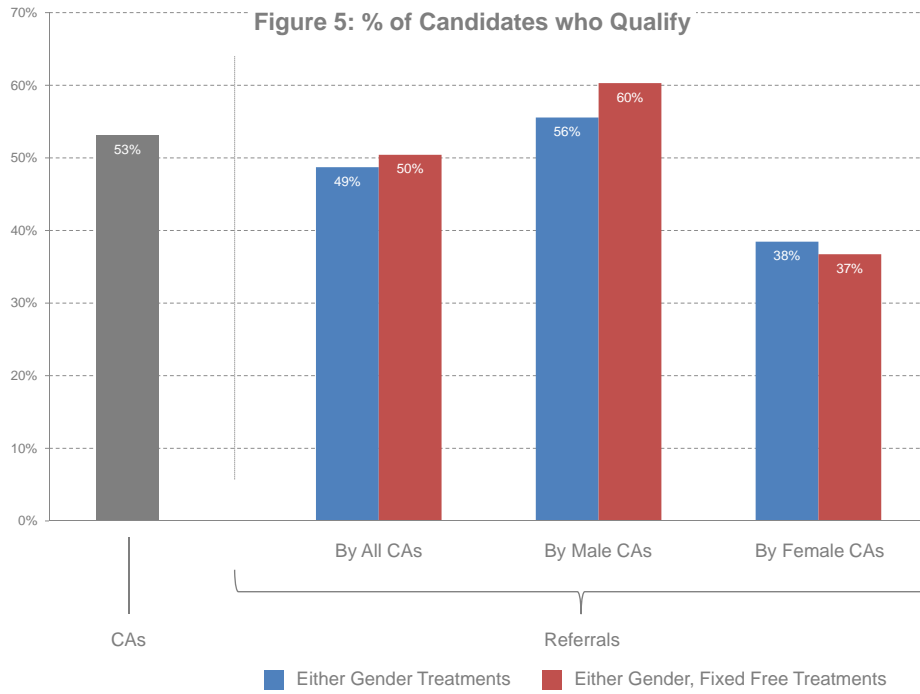
	CA Qualifies (1)	Referral Qualifies (2)	Referral Qualifies (3)	CA Qualifies (4)	Referral Qualifies (5)	Referral Qualifies (6)
Competitive Treatment	-0.073 (0.060)	0.065 (0.065)	0.043 (0.117)	0.019 (0.083)	0.047 (0.089)	0.092 (0.152)
Female Treatment			0.099 (0.113)			-0.132 (0.168)
Either Treatment			0.176 (0.119)			-0.267 (0.153) *
Competitive * Female Treatment			0.040 (0.159)			-0.138 (0.221)
Competitive * Either Treatment			0.053 (0.166)			0.029 (0.210)
Observations	288	255	255	176	151	151
CA Gender	Men	Men	Men	Women	Women	Women

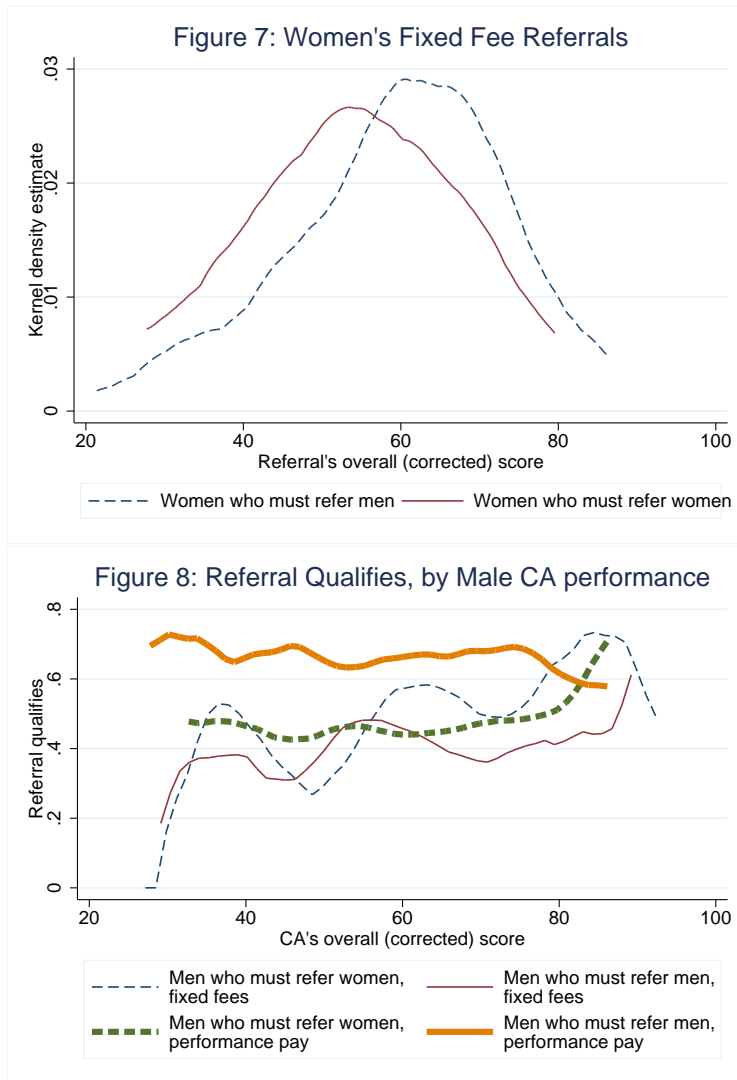
Notes

- 1 The dependent variable is indicated in the column heading.
- 2 All specifications include CA visit day dummies.









A Appendix

A.1 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 $N(0, \sigma_\varepsilon^g)$, 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 its relationship to Q_j or Y_j . Each potential referral of gender g is independently

drawn from a joint distribution $f^g(\alpha, Q)$. In addition to social payments, CAs may also consider ambient incentives to refer a high quality worker ($E[R(Y)|Q]$), which perhaps derive from reputational effects, as well as any direct financial incentives provided by the firm ($E[P_i(Y)|Q]$). $R(Y)$ is presumed to be increasing in Y . For simplicity, we consider contracts of the form $F_i + P_i I(Y_j > c)$, 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 CAs can choose from the entire network \mathcal{N} , CAs solve

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

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 referable at contract (F_i, P_i) if the CA can expect positive profits from referring j at that contract, that is, if $E[R(Y_j)|Q_j] + \alpha_j + E[P_i(Y_j)|Q_j] + F_i > 0$. If no one in the network is referable, then the CA declines to make a referral.

In this framework, men may be systematically chosen as referrals for four 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. Third, the distribution of social benefits, α , may differ across genders. Finally, the accuracy of quality signals, which may interact with the firm incentives and social payments to refer more men or women, may differ across male and female network members. We consider the implications for each of these in turn.

A.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 $P(j = N_{P_i}^* | j \in \mathcal{M}_i) = P(j = N_{P_i}^* | j \in \mathcal{F}_i)$*

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 referable women are much more scarce in CA networks than referable 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.

A.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 $E[Y_{j_N}^*] \geq E[Y_{j_M}^*] > E[Y_{j_F}^*]$.²⁷ Comparing quality distributions of referrals made under various gender restrictions and contract types allows a direct test of this hypothesis.

A.1.3 Social benefits

Proposition 1 $E[Y_{G_{P_i}^*}]$ is non-decreasing in P_i . $P(Y_{G_{P_i}^*} > Y_{G_0^*})$ 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 networks²⁸; and (iii): $\sigma_\varepsilon^g < \infty$. If any of conditions (i), (ii), or (iii) fail then $P(Y_{G_{P_i}^*} > Y_{G_0^*}) = 0$.

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 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.

²⁷Note 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.

²⁸"Not that much lower" depends on how much higher quality the person could be. The specific condition is $\int_{Q_0}^{\infty} \int_{\alpha_0 + E[R(Y_0) - R(Y) | Q_0, Q]}^{\alpha_0 + E[R(Y_0) - R(Y) | Q_0, Q] + P_i} \left(\Phi\left(\frac{c - Q}{\sigma_\varepsilon^g}\right) - \Phi\left(\frac{c - Q_0}{\sigma_\varepsilon^g}\right) \right) f^g(\alpha, Q) d\alpha dQ > 0$

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 1 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.

A.1.4 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 2 *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 $P[Y_{N_P^*} > c] \geq P[Y_{G_P^*} > c] \forall G$. If information is finite and the same between men and women ($\sigma_\varepsilon^F = \sigma_\varepsilon^M < \infty$) 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 ($\sigma_\varepsilon^F > \sigma_\varepsilon^M$), 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 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 2 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.²⁹ 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

²⁹This could either be 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.

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.

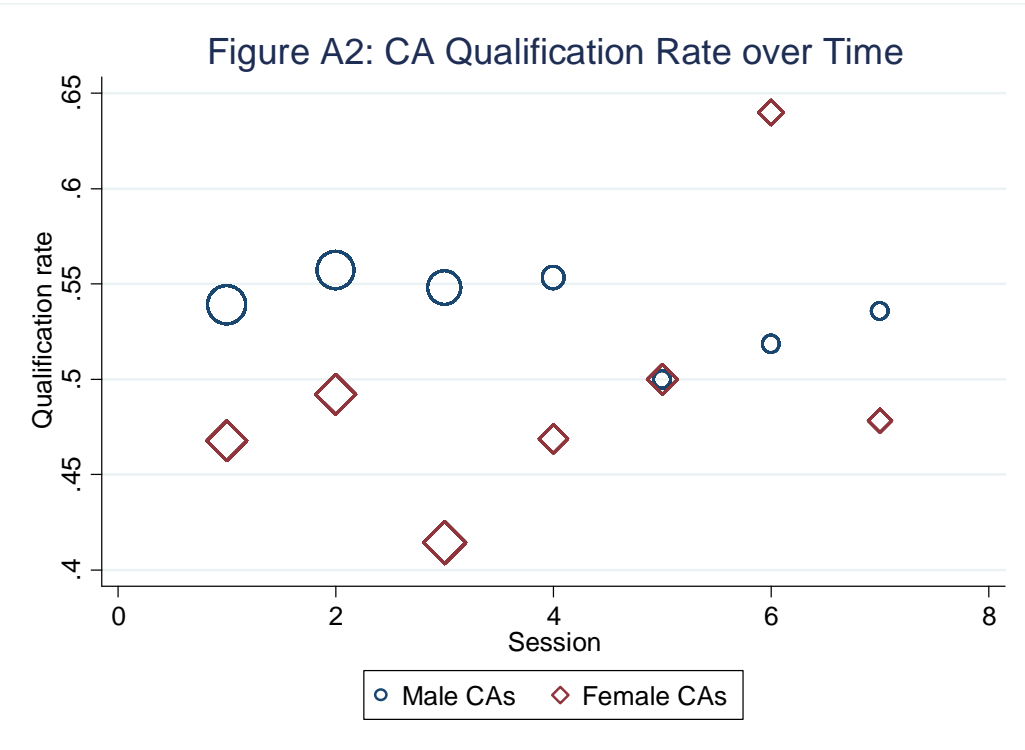
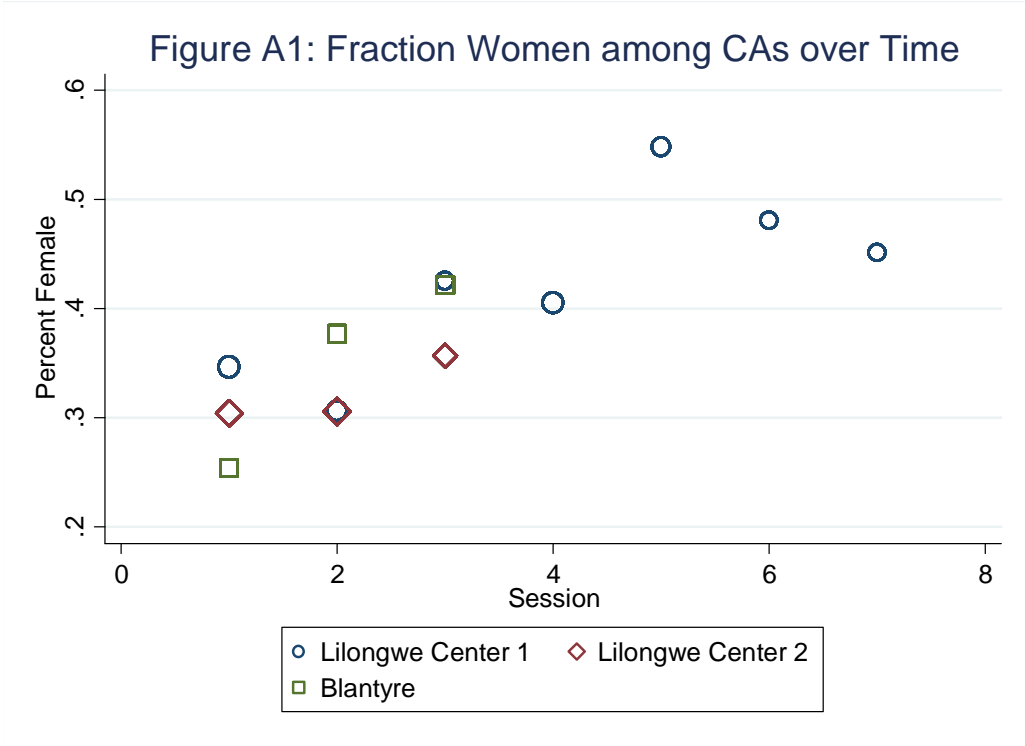
A.2 Appendix Tables and Figures

Appendix Table 1: Summary Statistics and Randomization Check

Dependent Variable	Mean and SD: Male	p value of joint test of treatments	N	Mean and SD: Female	p value of joint test of treatments	N
	(1)	(2)	(3)	(4)	(5)	(6)
CA Age	25.52 [3.90]	0.610	469	24.47 [4.55]	0.775	290
CA qualified	0.54 [0.50]	0.354	506	0.48 [0.50]	0.641	309
CA Overall Test Score (corrected)	61.31 [13.52]	0.567	506	59.68 [13.38]	0.121	309
CA Has Previous Survey Experience	0.30 [0.46]	0.524	505	0.27 [0.44]	0.120	308
CA Has Tertiary Education	0.68 [0.47]	0.274	506	0.76 [0.43]	0.165	309
CA MSCE Math Score	5.66 [2.30]	0.814	443	6.81 [1.81]	0.054	261
CA MSCE English Score	5.68 [1.50]	0.787	459	5.76 [1.39]	0.394	275
CA Job Comprehension Score	0.79 [0.40]	0.840	506	0.80 [0.40]	0.380	310
CA Math Score	0.21 [0.10]	0.263	506	0.18 [0.09]	0.327	310
CA Ravens Score	0.61 [0.39]	0.213	506	0.55 [0.40]	0.666	310
CA Language Score	0.15 [0.03]	0.250	506	0.14 [0.03]	0.785	310
CA Practical Component Z-score	-0.11 [1.02]	0.070	502	0.19 [0.89]	0.057	306
CA Computer Score	0.44 [0.21]	0.843	506	0.42 [0.20]	0.444	310
CA Feedback Points	25.84 [7.27]	0.018	500	27.97 [6.21]	0.151	306

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.



Notes: Both figures contain only data on CAs. Session is equal to 1 on the first day we were interviewing in a given center: either Lilongwe Center 1, 2 or in Blantyre. The size of the circles reflect the relative size of the sample at each session / training centre.

Figure A3: Referral Qualifies, by Female CA performance

