

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

Lori Beaman

Northwestern University

Niall Keleher

University of California, Berkeley and Innovations for Poverty Action

Jeremy Magruder

University of California, Berkeley

December 2016

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. Additionally, women did not refer enough high quality women 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 Group, 2011). One 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.

A large fraction of jobs - up to 50% - are attained through informal channels, including employee referrals (Bewley, 1999; Ioannides and Loury, 2004) and many - if not most - firms in the U.S. have programs to encourage employee referrals (CareerBuilder, 2012). In principle, it is unclear whether the use of job networks should benefit or harm women. An extensive literature in sociology (reviewed in McPherson, Smith-Lovin, and Cook (2001)) suggests that networks, particularly workforce networks, are quite gender homophilous. If men and women have distinct networks, and women are employed less frequently, than theory developed in Calvo-Armengol and Jackson (2004) would suggest that the use of networks exacerbates existing inequality.¹ On the other hand, if networks have access to information about hard-to-observe productive characteristics of women, then network screening could help women overcome discrimination deriving from the

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.

¹Mortensen 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. Moreover, Galenianos (2016) demonstrates that referrals reinforce existing inequalities.

“invisibility hypothesis” (Milgrom and Oster, 1987).

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-Malawi² advertised positions for survey enumerators using a traditional method of posting flyers. 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.

This recruitment drive allows a few types of analyses. We are able to observe referral choices for everyone. This contrasts with administrative data, where we would typically only observe referrals and conventional applicants who were hired into the firm. This allows us to examine the causal implications of hiring methods for the applicant pool without the potential for confounding variation driven by firm or worker choices to hire a particular person’s referral, or to make a referral in the first place.

To measure causal mechanisms in the referral decision, we cross-randomized the referral process along three dimensions. First, candidates were either told that they may refer a woman, that they may refer a man, or that they may refer a person of either gender. Second, their finder’s fee was randomly selected to be a fixed fee or a performance contract with a bonus paid if the referral attained a certain threshold.³ The full performance

²IPA-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.

³The fixed fee was randomized to be either 1000 or 1500 Malawi Kwacha (MWK; \$1=153 MWK,

incentive is approximately a day's wage for an enumerator.⁴ 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). As a result, we observe which type of people are chosen as referrals when men and women have constrained options, and under different incentive environments.

In our setting, qualified female candidates are disadvantaged by the use of referrals in this hiring drive. When conventional applicants (CAs) 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 (38%).⁵ The low number of women referred is driven 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, capable of referring women. Men choose to make a referral at identical rates when required to refer either women or men (84%). Moreover, women referred by men who are constrained to refer only women 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 are capable of referring women, but usually choose not to.

In contrast to men, women CAs do not exhibit a strong preference for either gender: 43% of women with unrestricted referral options choose to refer other women, which is (statistically) indistinguishable from the rate at which women apply themselves

2011). The performance contract was structured to guarantee 500 MWK with the potential to earn an additional 1300 MWK, for a total of 1800 MWK, if the referral attained a certain threshold.

⁴The 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.

⁵The Malawi Labor Force Survey of 2013 (NSO, 2013) indicates that 30% of all non-agricultural wage employees in Malawi are women.

(38%). However, since women's referrals also tend to be (weakly) less likely to qualify, generating referrals from women results in similar numbers of qualified women applicants as conventional application processes. We conclude that - at least in this context - recruiting through women's referrals would not lead to an advantage in encouraging qualified female applicants, and would come at a cost in terms of the quality of the overall applicant pool.

Our experimental design allows us to start to understand how the referral contract may interact with both the gender and quality of the applicant pool. 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 that prioritizes high quality workers over diversity may then prefer hiring male referrals. 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. While these incentives do not substantially affect the gender distribution of referred workers, we document that they do change referral choices on a number of different dimensions. In particular the firm got a more qualified referral pool when men are required to refer other men with a performance incentive. However, the performance incentive does not always lead to higher quality referrals, suggesting that (again, in

the range of contracts we consider) financial incentives may not always improve hiring outcomes for a firm. We discuss possible explanations for these two results in section 3.3.

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 creates a threat to external validity. In many cases in the real world, however, employees will compete with their own referrals. Some employees work directly with the people they refer as in Heath (2016), and many employees refer individuals who will work at the same level in the company as they do (Brown, Setren, and Topa, 2016). These employees compete with their referrals for promotions. Our experiment also created exogenous variation in the salience of competition between applicants using the relative versus absolute qualification threshold. 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.

As with any experiment, there is a risk that results would not generalize to other contexts. Should we expect network members in Malawi to distribute referrals in a very different way from other parts of the world? There are a few reasons to be concerned about external validity. First, we may expect that gender relations in Malawi, where women are less likely to finish secondary school than men, are very different than in other parts of the world, particularly in the U.S and Europe. In Calvo-Armengol and Jackson (2004), the key primitives of the model that generates the prediction that “inequality

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

begets inequality” are that (i) networks are homophilous and (ii) one group is initially disadvantaged in the labor market. We observe these two features in labor markets worldwide: gender homophily in social networks is found in rich and poor countries, and women earn less than men globally. As a practical example of this homophily, Caetano and Maheshri (2015) show that even within neighborhoods in the U.S., men visit establishments where they are much more likely to encounter other men than women. This pattern of social interactions could easily generate our findings from the experiment: if CAs tell the next person they interact with about the new opportunity, they will more likely refer a man. This is of course only one of many possible explanations, but it highlights that our finding need not be specific to a developing country context.

Even if networks share some important structural features across contexts, we may remain concerned that referrals happen in a very different way in firms other than our partner, in industries other than survey enumeration, or in contexts where more experienced workers are making most referrals. Yet, both this experiment and the theoretical models of networks and inequality are motivated by broad trends in data from rich countries which are consistent with similar mechanisms operating elsewhere. The idea that women are ill-served by networks is one of the “stylized facts” about job information networks presented by Ioannides and Loury (2004). This stylized fact is supported by observational studies from a wide variety of contexts, mostly in rich countries, suggesting that a broad range of women gain less from networks. For example, Lalanne and Seabright (2016) provide suggestive evidence that female CEOs have smaller and less effective networks than their male counterparts, while Bortnick and Ports (1992) document that unemployed women also find job search less effective through their networks. There

is also evidence that jobs found through female connections may be less desirable: Loury (2006) using the National Longitudinal Survey of Youth (NLSY) found that male workers referred by women get lower on average wages than those who applied through formal channels. Women in the US appear to internalize that they are underserved by networks and are less likely to report informal contacts as a method of job search (Bradshaw, 1973; Ports, 1993). The results are also consistent with the finding from observational data from a call-center in Fernandez and Sosa (2005).⁷ While it remains possible that the experiment would deliver different results with a different firm, or in a different country, it is striking that the bottom line from our experiment is so consistent with what has been documented for women so many times elsewhere in less tightly controlled contexts. 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.

Finally, we note that this paper contributes to two large literatures in labor economics. The literature on gender disparities in economics has largely 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).⁸ We find that another channel may be at play: how firms make hiring decisions. The literature on

⁷In the context of Fernandez and Sosa (2005), men are the disadvantaged group, who are similarly less likely to receive referrals. Using qualitative approaches, other researchers have attributed the lack of women in upper tier positions more generally to two factors. First, a failure of networks at the top Fawcett and Pringle (2000) (e.g. Fawcett and Pringle 2000; Holgersson 2013). Second, differences in how men and women form their networks: for example, Seabright (2012) 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.

⁸Additional 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).

gender has many similarities with the broader literature on disadvantaged groups. The model in Calvo-Armengol and Jackson (2004) was in fact motivated by the black-white wage gap in the U.S. The survey by Ioannides and Loury (2004) highlights as another stylized fact that informal search appears to be less effective for blacks than whites. Related work has looked at how different hiring methods impact the recruitment of minority workers, including the race of the manager in Giuliano, Levine, and Leonard (2009) and the use of formal job testing in Autor and Scarborough (2008).

This paper is also relevant to the broader literature on hiring. Oyer and Schaefer (2011) argue in their handbook chapter that there is too little work on hiring decisions of firms. There has been a recent resurgence of research on employee referrals, related to the seminal work by Granovetter (1973) and Montgomery (1991). Burk et al. (2015) uses data from 9 firms in the U.S. to demonstrate that employee referrals can benefit firms in terms of a higher recruitment rate and lower turnover. The recent literature has also sought to understand why firms use referrals; most of the work has focused on asymmetric information such as screening applicants (Brown et al. (2016); Dustmann, Glitz, and Schoenberg (2016); Hensvik and Skans (2016); Pallais and Sands (Forthcoming); Beaman and Magruder (2012)) and inducing effort on-the-job (Heath (2016); Kugler (2003)). Our paper highlights that while referrals may be helpful in reducing labor market frictions, it may come at a significant cost in terms of access to opportunities for initially-disadvantaged groups.

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, Kolev, and Filipiak, 2010). Malawi is not an exception to this trend. A recent survey of Malawian households suggests that less than one-third of women participate in the formal labor force, while nearly 58% of men do so (World Bank Group, 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%) (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.⁹ 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).¹⁰ For this experiment,

⁹See Godlonton (2014) for a comprehensive 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.

¹⁰Informal interviews with qualified female applicants suggest that one reason qualified female applicants were hard to find was that there are gender differences in willingness to travel regularly and for several weeks at a time in Malawi, which is necessary to work as a survey enumerator.

we 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) across 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 over time. In some recruitment sessions, we interviewed both CA and referral applicants. However, CAs 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 school 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.¹¹ 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 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.¹² 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.¹³

¹¹The 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.

¹²All 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.

¹³The 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 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.

As a final component to the practical test, IPA asked the experienced enumerators to provide a feedback score for participants. Since there is the potential for any biases on the part of the enumerators to affect this component, we remove it from our calculations in overall qualification measures in the analysis in this paper. However, we show how feedback points varied across treatment groups in Tables 3 and 5.

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 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 instructed CAs to refer a man, a woman, or we relaxed the gender restriction. 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.¹⁴ Most CAs did not know their score or whether they qualified before making their referral.

¹⁴To 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.

2.4 Internal Validity and CA Characteristics

Appendix Table 1 displays summary statistics for the sample of CAs, for men and women separately. It also shows that the randomization produced balance along most characteristics.

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. 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.¹⁵ 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 CAs.¹⁶ 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 Panel C also shows that women’s referrals are 10 percentage points less likely to qualify than men’s, though this difference is not statistically significant.¹⁷ Considering both the gender composition and qualification effects reveals that women CAs are more likely to refer qualified women than men CAs are (18 percent versus 11 percent, though only marginally significant). We discuss these results 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. Therefore the same trend in getting fewer women through referrals than through the traditional recruitment method continues if you look at only qualified applicants.

¹⁵Table 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.

¹⁶In both cases, the sample is restricted to CAs who have the choice of which gender to refer.

¹⁷This difference grows to 18 percentage points and becomes statistically significant, if feedback scores from enumerators are incorporated in the qualification measure (which was how IPA-Malawi actually determined qualification).

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 during the middle period of recruitment. By contrast, the qualification rate among men is largely flat. Overall, there is little evidence then that qualified women were unavailable after the initial interview session.

3.2 Are Qualified Women Absent from 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 present 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 present 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. (2016) 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, relative to fixed fees. 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.

First, the best applicants come from male CAs referring other men when offered performance pay. 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 20 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.¹⁸

Second, the performance incentive does not improve the quality of referrals among CAs who were asked to refer women or who were free to choose anyone. 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

¹⁸This 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 Bandiera, Barankay, and Rasul (2009) and Beaman and Magruder (2012), who both argue that social incentives can skew the CA's behavior away from what would benefit the firm most. 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. It is worth noting that Column (2) of Table 2 showed that fewer CAs made referrals in the performance pay-male referral treatment. There therefore 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.

is essentially zero). Simple descriptive statistics demonstrate clearly that among male CAs in performance pay treatment, the referred men outperform the referred women: 62% of referrals qualify in the male-only treatment vs 41% in the female-only treatment. A number of possible mechanisms may underlie this trend: 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). 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, and we do not want to over-interpret it. We offer a model in the appendix which provides a potential explanation.¹⁹

Taken together, the firm gets the highest quality candidates by asking male CAs to refer other men and providing a performance incentive. This implies that there is the potential for firm incentives to increase bias against women, by allowing more referrals from 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, contracts which put more emphasis on the

¹⁹In the model, a CA is maximizing two distinct objects: (i) the firm-provided benefit, which may depend on the ability of the worker, and (ii) a social benefit they get from the network member they refer and a benefit they get from the firm. 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. If CAs get a noisy signal of the ability and a precise signal of social benefits of each network member, the model shows that 'surprising' results can occur. CAs may 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.

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 points²⁰ 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

Attrition in this study is driven by CAs choosing not to make a referral. One striking trend from section 3.2.1 was that CAs restricted to refer a particular gender chose to make a referral at extremely similar rates regardless of whether they were restricted to refer men or women. This suggests that any differences between referrals restricted to be male or restricted to be female may be unlikely to be attributable to differences in attrition. However, section 3.2.1 also revealed that contracts affect attrition where gender restrictions did not: male CAs were more likely to make a referral in the presence of fixed fees than performance pay.²¹ In principle, these differential return rates mean that we can't attribute changes in referral characteristics under different contract types to either the choice of who to bring in, or to changes in the composition of which CAs make a referral. 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

²⁰Feedback 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. These are not included in the main measure of qualification.

²¹In Section 3.4, we will also note that female CAs responded similarly.

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 the model in the appendix.²²

3.4 Women CAs' Behavior

Figure 2 shows 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. Figure 5 and Table 1, however, suggest a need for some caution with this interpretation: while average qualification differences between referrals of men and women are not quite significant, the point estimates are fairly large: women's referrals on average qualify ten percentage points less often than men's ($p = 0.14$).²³ That said, women are more likely to refer qualified women than men are (18 percent versus 11 percent of unrestricted referrals). These numbers are, however, clearly still low and not an improvement over the traditional recruitment method (19 percent of CA applicants are qualified women). Figure 6 reports a non-parametric plot of

²²If 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.

²³This difference becomes marginally statistically significant when feedback points are incorporated into the measure of qualification.

referral ability against CA ability and observes that the ability of men’s referrals weakly dominates that of women’s across the CA skill distribution, with particularly large differences for highly skilled men and women. From this, we infer that these patterns would remain if only qualified CAs were eligible to make referrals.²⁴

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 0.46 of a standard deviation below the CA mean, while men referred by women CAs perform 0.09 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 Raven’s 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

²⁴Examining only qualified CAs, unrestricted referrals of female, qualified CAs are 14 percentage points less likely to qualify than referrals of qualified male CAs ($p = 0.12$).

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).²⁵ There are a number of plausible alternative explanations. First, women may struggle to know which characteristics make a candidate qualified. Second, women may 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.

²⁵Appendix 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.

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, Leibbrandt, and List (2014) and Niederle and Vesterlund (2007).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. In our setting, the referral only marginally affects the likelihood of qualifying or getting called for a job (given the large number of recruits).²⁷

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

²⁶Niederle 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.

²⁷On 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.

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 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 6 shows that referrals, both men and women, are if anything more likely to qualify when CAs are directly competing with their referrals (significantly so, for male CAs, but always with a positive point estimate). 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.²⁸ Using an experiment designed around a recruitment drive 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.

References

- Altonji, Joseph G. and Rebecca M. Blank. 1999. Race and gender in the labor market. In *Handbook of labor economics*, vol. 3, Part C. Elsevier, 3143 – 3259.
- Arbache, Jorge S., Alexandre Kolev, and Ewa Filipiak. 2010. *Gender disparities in Africa’s labor market*. World Bank Publications.
- Autor, David H and David Scarborough. 2008. Does job testing harm minority workers? Evidence from retail establishments. *The Quarterly Journal of Economics* 123, no. 1:219–277.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2009. Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica* 77, no. 4:1047–1094.

²⁸Moreover, occupational segregation is commonly cited as a source of income disparity across gender (Blau and Kahn (2000); Arbache et al. (2010)). The use of employee referrals may be one of the mechanisms creating this segregation (Fernandez and Sosa, 2005; Tassier and Menczer, 2008).

- Beaman, Lori and Jeremy Magruder. 2012. Who gets the job referral? Evidence from a social networks experiment. *American Economic Review* 102, no. 7:3574–3593.
- Bertrand, Marianne. 2011. New perspectives on gender. In *Handbook of labor economics*, vol. 4, Part B. Elsevier, 1543 – 1590.
- Bewley, Truman F. 1999. *Why wages don't fall during a recession*. Harvard University Press.
- Black, Sandra E. and Elizabeth Brainerd. 2004. Importing equality? The impact of globalization on gender discrimination. *Industrial and Labor Relations Review* 57, no. 4:540–559.
- Black, Sandra E. and Philip E. Strahan. 2001. The division of spoils: Rent-sharing and discrimination in a regulated industry. *American Economic Review* 91, no. 4:814–831.
- Blau, Francine D. and Lawrence M. Kahn. 2000. Gender differences in pay. *Journal of Economic Perspectives* 14, no. 4:75–99.
- Bortnick, Steven M. and Michelle H. Ports. 1992. Job search methods and results: Tracking the unemployed, 1991. *Monthly Labor Review* 115, no. 12:29–35.
- Bradshaw, Thomas F. 1973. Jobseeking methods used by unemployed workers. *Monthly Labor Review* 96, no. 2:35–40.
- Brown, Meta, Elizabeth Setren, and Giorgio Topa. 2016. Do informal referrals lead to better matches? Evidence from a firm's employee referral system. *Journal of Labor Economics* 34, no. 1:169–209.
- Burk, Stephon, Bo Cowgill, Mitchell Hoffman, and Michael Housman. 2015. The value of hiring through employee referrals. *Quarterly Journal of Economics* 130, no. 2:805–839.
- Caetano, Gregorio and Vikram Maheshri. 2015. Homophily and sorting within neighborhoods. Mimeo, University of Rochester.
- Calvo-Armengol, Antoni and Matthew O. Jackson. 2004. The effects of social networks on employment and inequality. *American Economic Review* 94, no. 3:426–454.
- CareerBuilder. 2012. Referral madness: How employee referral programs turn good employees into great recruiters and grow your bottom line. CareerBuilder e-book.
- Dustmann, Christian, Albrecht Glitz, and Uta Schoenberg. 2016. Referral-based job search networks. *Review of Economics Studies* 83, no. 2:514–546.
- Eagley, Alice and Maureen Crowley. 1986. Gender and helping behavior: A meta-analytic review of the social psychological literature. *Psychological Bulletin* 100, no. 3:283–308.
- Fawcett, Rebecca and Judith K. Pringle. 2000. Women ceos in new zealand: Where are you? *Women in Management Review* 15, no. 5-6:253–260.

- Fernandez, Roberto M. and M. Lourdes Sosa. 2005. Gendering the job: Networks and recruitment at a call center. *American Journal of Sociology* 111, no. 3:859–904.
- Flory, Jeffrey A., Andreas Leibbrandt, and John A. List. 2014. Do competitive workplaces deter female workers? A large-scale natural field experiment on gender differences in job-entry decisions. *The Review of Economic Studies* 81, no. 1:122–155.
- Galenianos, Manolis. 2016. Referral networks and inequality. Mimeo, Royal Holloway, University of London.
- Giuliano, Laura, David I Levine, and Jonathan Leonard. 2009. Manager race and the race of new hires. *Journal of Labor Economics* 27, no. 4:589–631.
- Godlonton, Susan. 2014. Employment risk and job-seeker performance. Tech. Rep. 1332, IFPRI Discussion Paper.
- Goldin, Claudia and Lawrence F. Katz. 2002. The power of the pill: Oral contraceptives and women’s career and marriage decisions. *Journal of Political Economy* 110, no. 4:730–770.
- Granovetter, Mark. 1973. The strength of weak ties. *American Journal of Sociology* 78, no. 3:1360–1380.
- Heath, Rachel. 2016. Why do firms hire using referrals? Evidence from Bangladeshi garment factories. Mimeo, University of Washington.
- Hensvik, Lena and Oskar Nordstrom Skans. 2016. Social networks, employee selection and labor market outcomes. *Journal of Labor Economics* 34, no. 4:825–867.
- Ioannides, Yannis M. and Linda D. Loury. 2004. Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature* 42, no. 4:1056–1093.
- Kugler, Adriana. 2003. Employee referrals and efficiency wages. *Labour Economics* 10, no. 5:531–556.
- Lalanne, Marie and Paul Seabright. 2016. The old boy network: The impact of professional networks on remuneration in top executive jobs. Tech. Rep. 123, SAFE Working Paper.
- Loury, Linda D. 2006. Some contacts are more equal than others: Informal networks, job tenure, and wages. *Journal of Labor Economics* 24, no. 2:299–318.
- McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27, no. 1:415–444.
- Milgrom, Paul and Sharon Oster. 1987. Job discrimination, market forces, and the invisibility hypothesis. *The Quarterly Journal of Economics* 102, no. 3:453–476.
- Montgomery, James D. 1991. Social networks and labor market outcomes: Toward an economic analysis. *American Economic Review* 81, no. 5:1408–1418.

- Mortensen, Dale T. and Tara Vishwanath. 1994. Personal contacts and earnings: It is who you know! *Labour Economics* 1, no. 2:187–201.
- Niederle, Muriel and Lise Vesterlund. 2007. Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics* 122, no. 3:1067–1101.
- NSO. 2013. Malawi labor force survey report. Tech. rep., National Statistics Office, Zomba, Malawi.
- NSO and ICF Macro. 2011. *Malawi Demographic and Health Survey 2010*. National Statistics Office and ICF Macro.
- Oyer, Paul and Scott Schaefer. 2011. Personnel Economics: Hiring and Incentives. In *Handbook of labor economics*, vol. 4. Elsevier, 1769–1823.
- Pallais, Amanda and Emily G. Sands. Forthcoming. Why the referential treatment? Evidence from field experiments on referrals. *Journal of Political Economy*.
- Ports, Michelle H. 1993. Trends in job search methods, 1970-92. *Monthly Labor Review* 116, no. 10:63–67.
- Seabright, Paul. 2012. *War of the sexes*. Princeton University Press.
- Tassier, Troy and Filippo Menczer. 2008. Social network structure, segregation, and equality in a labor market with referral hiring. *Journal of Economic Behavior and Organization* 66, no. 34:514–528.
- World Bank Group. 2010. *The education system in Malawi: Country status report*. World Bank Publications.
- . 2011. *World Development Report 2012: Gender Equality and Development*. World Bank Publications.

Table 1: Gender Distributions of CAs and Referrals

	(1)	(2)	(3)	(4)
	All CAs	Male CAs	Female CAs	Diff: p value
<u>A. CA Characteristics</u>				
Fraction of CAs	100%	62%	38%	
CA is Qualified	53%	55%	49%	0.088
N	816	506	310	
<u>B. CA Characteristics: Either Gender Treatments</u>				
Fraction of CAs	100%	60%	40%	
CA is Qualified	57%	60%	51%	0.148
N	266	159	107	
<u>C. Referral Characteristics: Either Gender Treatments</u>				
Referral is Female	30%	23%	43%	0.002
Referral is Qualified	49%	53%	43%	0.144
Referral is Qualified Male	35%	42%	24%	0.006
Referral is Qualified Female	14%	11%	18%	0.097
N	220	133	87	
<u>D. Referral Characteristics, Fixed Fee Treatments</u>				
Referral is Female	32%	23%	43%	0.017
Referral is Qualified	50%	56%	41%	0.094
Referral is Qualified Male	35%	43%	23%	0.019
Referral is Qualified Female	15%	13%	18%	0.442
N	133	77	56	
<u>E. Referral Characteristics, Perf Treatments</u>				
Referral is Female	29%	21%	42%	0.043
Referral is Qualified	47%	48%	45%	0.788
Referral is Qualified Male	36%	41%	26%	0.158
Referral is Qualified Female	11%	7%	19%	0.089
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)	-0.067 (0.060)	0.033 (0.078)	
Either Gender Treatment	0.014 (0.040)	-0.052 (0.052)	0.019 (0.062)	0.136 (0.080)	*
Performance Pay		-0.148 (0.056)	***	0.202 (0.090)	**
Perf Pay * Female Treatment		0.004 (0.076)		-0.248 (0.122)	**
Perf Pay * Either Treatment		0.152 (0.079)	*	-0.287 (0.125)	**
Observations	506	506	429	429	
Mean of excluded group	0.840	0.892	0.485	0.398	

Table 3: Screening of Male 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	-0.030 (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

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)	-0.245 *** (0.079)	-0.272 ** (0.106)
Either Gender Treatment	0.017 (0.055)	-0.024 (0.071)	-0.208 *** (0.078)	-0.232 ** (0.100)
Performance Pay		-0.113 (0.080)		0.013 (0.118)
Perf Pay * Female Treatment		-0.013 (0.111)		0.056 (0.162)
Perf Pay * Either Treatment		0.086 (0.110)		0.071 (0.162)
Observations	310	310	254	254
Mean of Excluded Group	0.821	0.852	0.590	0.609

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

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	-0.053 (0.060)	0.136 (0.064)	0.194 (0.116)	0.047 (0.082)	0.134 (0.088)	0.117 (0.149)
Female Treatment			0.111 (0.112)			-0.266 (0.165)
Either Treatment			0.174 (0.118)			-0.310 (0.150)
Competitive * Female Treatment			-0.085 (0.158)			-0.012 (0.217)
Competitive * Either Treatment			-0.068 (0.165)			0.107 (0.206)
Observations	288	255	255	176	151	151
CA Gender	Men	Men	Men	Women	Women	Women

Figure 1: CA Ability by Gender




Figure 2: % of Candidates who are Women




Figure 3: Gender choice in referrals, by CA performance





Figure 4: Men's Fixed Fee Referrals








Figure 7: Women's Fixed Fee Referrals



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 $N(0, \sigma_\varepsilon^2)$, 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 $E[Y_{j_N}^*] \geq E[Y_{j_M}^*] > E[Y_{j_F}^*]$.²⁹ 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 $E \left[Y_{G_{P_i}^*} \right]$ 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

²⁹Note 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.

in expectation than the person who is being referred under fixed fees and whose social payments are not much lower in gender G networks³⁰; and (iii): $\sigma_\varepsilon^g < \infty$. If any of conditions (i), (ii), or (iii) fail then $P\left(Y_{G_{P_i}^*} > Y_{G_0^*}\right) = 0$.

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 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.

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 $P\left[Y_{N_P^*} > c\right] \geq P\left[Y_{G_P^*} > c\right] \forall G$. If information is finite and the same between men and women ($\sigma_\varepsilon^F = \sigma_\varepsilon^M < \infty$) then Proposition 2 applies to unrestricted choices and performance premia will be positive unless condition (ii) fails*

³⁰ "Not that much lower" depends on how much higher quality the person could be. The specific condition is $\int_{Q_0}^{\infty} \int_{\alpha_0 + E[R(Y_0) - R(Y)|Q_0, Q]}^{\alpha_0 + E[R(Y_0) - R(Y)|Q_0, Q] + P_i} \left(\Phi\left(\frac{c-Q}{\sigma_\varepsilon^g}\right) - \Phi\left(\frac{c-Q_0}{\sigma_\varepsilon^g}\right)\right) f^g(\alpha, Q) d\alpha dQ > 0$

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.³¹ 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.

³¹This 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.

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	0.55 [0.50]	0.280	506	0.49 [0.50]	0.497	309
CA Overall Test Score (corrected)	61.88 [14.32]	0.555	506	59.65 [14.39]	0.175	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

Fig A1: Fraction women among CAs over time




Fig A2: CA qualification rate over time




Figure A3: Referral Qualifies, by Female CA performance

