Competitive Job Seekers: When Sharing Less Leaves Firms at a Loss *

Gaurav Chiplunkar
University of Virginia

Erin M. Kelley
World Bank

Gregory Lane
University of Chicago

September 2023

Abstract

We use a randomized control trial to study how young job-seekers share information about jobs within their social network, and its implications for firms. When competition for a job is made salient, we find that job-seekers are: (i) less likely to share information about the job with their peers; and (ii) choose to share it with fewer higher ability peers. This lowers the size and quality of the applicants a firm sees and the hires they make. These results suggest that firms who rely heavily on social networks to disseminate information about jobs may see lower quality applicants than they expected for positions that are deemed more competitive.

*Chiplunkar: ChiplunkarG@darden.virginia.edu Kelley: erinmkelley@worldbank.org Lane: laneg@uchicago.edu

We are grateful to Fiona Burlig, Rob Garlick, Sylvan Herskowitz, and Jeremy Magruder for their helpful comments and suggestions. Funding for this project was graciously provided by the IGC. We would like to thank Prachi Shukla for excellent research assistance. We would also like to thank Christy Tharaniyil and Mahika Mehta, who provided invaluable assistance implementing this project. AEA RCT identification number: 0007564. This project received IRB approval from The University of Virginia (IRB# 3463).

1 Introduction

Labor markets rely on social networks to spread information about jobs. Firms use these networks to identify and screen the most promising candidates, while job-seekers rely on their peers to learn about new job opportunities.[1] Social networks are particularly relevant in labor markets where information about jobs is either difficult to come by, or difficult to parse through. While social networks are typically considered to be an asset for transmitting information about a new opportunity, recent work suggests that their efficiency depends on whether individuals are in direct competition with one another and so face 'strategic disincentives' to share information (Cai and Szeidl, 2018; Hardy and McCasland, 2021). In labor markets, jobs are often 'rival'– hiring one person prevents someone else from getting the job– and individuals may choose to withhold information about a new labor market opportunity if they feel that doing so will enhance their own chances of employment. This has implications for firms. Specifically, companies that depend on social networks or referrals for filling competitive positions might face limitations in accessing the most qualified talent if job-seekers decide to withhold information from their highly skilled peers to avoid being in direct competition with them. This might also explain why referral systems can make the recruitment process less effective (Fafchamps and Moradi, 2015; Beaman and Magruder, 2012).

In this paper, we designed a randomized control trial to investigate how information about jobs flows through social networks and how the rival nature of this information impacts 1) the probability the job is shared, 2) the characteristics of who receives the information about the job, 3) how it affects the pool of applicants, and 4) who is ultimately hired for the job. To do so, we partnered with six colleges in Mumbai, India. Each college had multiple programs of study (hereby referred to as a 'batch') including commerce, finance, IT, management, and HR. Every week, each batch was assigned to receive information about a particular job designed by the research team. Moreover, 20% of students within the batch were selected to hear about this job opportunity. Students were labelled 'entry-points' if they heard about the job directly from us, and 'students' otherwise. Because we randomly selected a new group to hear about the jobs each week (for a month) the set of entry-points changed from one week to the next. The jobs we shared were randomly assigned to be 'rival' or 'non-rival' in nature. A 'rival' job meant that students we informed about the job (the entry-points) would have to apply and compete with their peers for these positions. A 'non-rival' job meant that entry-point students were guaranteed the position, but

[1] See Beaman and Magruder (2012); Dustmann, Glitz, Schönberg and Brücker (2016); Pallais and Sands (2016); Kramarz and Skans (2014) for some examples, and Trimble and Kmec (2011) for a review.

could still share the information with their peers. The assignment of whether a job was categorized as rival or non-rival was done randomly at the batch-week level. In other words, all participating batches were randomly allocated to receive information about a rival or non-rival job each week, and 20% of students within the batch were selected to hear about the job directly from the research team.

We are interested in whether the share and composition of students who hear and apply to jobs differs when we make it clear to entry-points that they will have to compete (rival), or not (non-rival), for the job. We rely on four sources of data. First, a baseline survey, where we captured detailed information about students' social networks. We use this to assess how being directly connected to entry-points affects whether the information is received. Second, we conducted weekly surveys with all students in our sample (entry-points and not) to capture whether they heard about the job. We were able to capture application data as we required all interested job-seekers to submit an online application to the job. Finally, we applied a simple hiring rule by selecting the candidates with the highest GPA, and any entry-point applicants in the non-rival batches who were guaranteed the job.

We document four main findings. First, we find that job-seekers were more likely to share information about jobs when they did not have to worry about competing for the job. On average, students in non-rival batches were 5 p.p. (30%) more likely to hear about the job we shared with entry-points relative to students in rival batches. Moreover, students with no direct connection to entry-point students were 3 p.p. more likely to hear about the job when it was non-rival (the probability that a student who was connected to an entry-point heard about the job increased by 25 p.p. regardless of whether the job is rival or not). This indicates that job information was more likely to spread beyond immediate connections to the entry-points when it was "non-rival" in nature.

Next, we investigate three factors that may exacerbate (or mitigate) the nature of competition between job-seekers. We see that job-seekers are more strategic about *who* they share information with when the job is rival and they know they have to compete for it. First, we explore how job-seekers' ability affects outcomes. We find that students were 8.5 p.p. *more* likely to hear about a job from lower ability entry-point peers (relative to higher or same ability peers) when the job was non-rival. Conversely, they were 7.5 p.p. *less* likely to hear about a job from lower ability entry-point peers (relative to higher or same ability peers) when the job was rival. This difference is statistically significant, and suggests that entry-points were taking the relative ability of their peers into account when deciding whether to share information or not. Second, we investigate the strength of close connections. Students who reported a close connection to an entry-point at baseline (as compared to those who did not) are more likely to hear about this information, and this holds whether the job is rival or not. Students are 11 p.p. more likely to hear about a job when they are closely connected to the entry point when the job is rival, and by 9 p.p. more likely to hear about it when the job is non-rival group. This suggests that closer social bonds can mitigate competitive concerns. Finally, we investigate the role of homophily. We find that when entry-points were guaranteed a job (non-rival), they were 10 p.p. more likely to share job information with another student of the same gender as compared to when the position was not guaranteed (rival). This suggests that students perceived greater competition from others of the same gender, and that this outweighs any social preferences for same gender connections. Lastly, these effects on ability and homophily are driven by men, who appear less likely to share information with other high-ability, male peers when the job information is rival.

Third, we establish that these results have implications for the quality of candidates that apply, and are hired for a job, as measured by their academic performance (GPA scores). Our detailed data collection allows us to examine how the pool of applicants differs between rival and non-rival batches along each step of the hiring process i.e., who heard about, applied to, and was hired for the job. We find that the GPA of students who *heard* about the job was 0.08 standard deviations higher when the job was non-rival. This translates into higher ability candidates who *applied* and were ultimately *hired* for the position as well: the GPA of applied (hired) candidates was 0.13 (0.38) standard deviations higher when the job was non-rival.

Our results indicate that firms advertising competitive positions may miss out on high-ability candidates if they rely heavily on social networks to disseminate information about jobs. To attract better talent, firms could employ the conventional strategy of making jobs more desirable by offering higher wages. If higher quality candidates demand higher compensation, then offering higher wages is essential to attracting such candidates (Dal Bó et al., 2013). However, increasing the wage changes two margins for job-seekers considering whether or not to share information about the position: the *competition* channel, whereby the cost of sharing job information increases as informed job-seekers lose more if those they share information with get the job; and (ii) the *altruism* channel, whereby the benefits of sharing the job increase as informed job-seekers derive greater satisfaction from helping their friends access better job opportunities. Since these channels work in opposite directions, if the competition channel is stronger than the altruism channel, job-seekers may be less likely to share the job with their high-ability peers. As a result, employers may find themselves with a lower quality candidate pool than they initially expected despite offering higher wages. To investigate this further, we conduct a sub-experiment to test whether job seekers propensity to share information changes when firms advertise jobs with higher wages, all else equal. To do so, we experimentally vary whether the job offers a high wage (10 USD) relative to the status-quo (5 USD). We find that doubling the wage among rival jobs attracts better quality hires (0.08σ) relative to the status-quo rival job. This confirms that higher wages can indeed attract better talent. However, an advantage of our setting is that we can identify how much stronger candidates ability would have been if higher wages could be offered without triggering a competitive response among job-seekers i.e., if the jobs were non-rival in nature. We find that the ability of hired candidates among high-wage *non-rival* jobs increases by an additional 0.35σ relative to high-wage *rival* jobs. Put differently, firms would have to increase the wage by almost four times (rather than doubling it, like we did for the experiment) to get the same increase in ability induced by eliminating strategic disincentives.

Taken together, our results suggest that job-seekers are strategically sharing less information about jobs when they are concerned about having to compete for them. In particular, they share less with students they perceive to be higher ability than themselves, and thus a greater competitive threat. This behavior has significant implications for the labor market, as it leads to a reduction in the overall quality of applications and hired candidates that firms receive. While firms can try and compensate for this by offering more attractive jobs, it is difficult to completely shut down the competition channel that drives job-seekers to share less. These results highlight the potential drawbacks for firms of relying solely on social networks and referrals to disseminate job information. They also motivate interventions that facilitate a broader dissemination of job information without relying exclusively on social networks.[2]

This study makes three primary contributions. First, we contribute to a large literature on the importance of referral networks in the labor market that follows the seminal work of Granovetter (1973) and Montgomery (1991). Empirical work in this space has investigated whether referrals improve the quality of hires or not. On the one hand, if individuals have information about their peers that is difficult for firms to observe, relying on referrals can reduce screening costs and improve hiring outcomes. Recent work has demonstrated how referrals can reduce asymmetric information (Beaman and Magruder, 2012; Brown et al., 2016; Pallais and Sands, 2016; Dustmann et al., 2016) and induce effort on the job (Kugler, 2003; Heath, 2018). On the other hand, if individuals lack pertinent information about their peers that employers seek, or if they prioritize recommending friends regardless of their quality, referral-based hiring can distort the recruitment process. Beaman and Magruder (2012) and Fafchamps and Moradi (2015) find evidence for this: the former reveals that referrals only en-

[2] Online job portals have the potential to fulfill this role, but their matching algorithms must be sophisticated enough to ensure that firms are not overwhelmed with irrelevant applications.

dorse higher-performing candidates when provided with financial incentives, while the latter demonstrates that high-ranked officers make lower quality referrals than lower-ranked soldiers because they are less concerned about promotion prospects tied to their referrals. Our work complements this literature by formally testing an important channel that affects the quality of referrals: namely competition. Our results could explain why the impact of referrals varies across contexts. When competition is minimal, as seen among full-time existing employees (Dustmann et al., 2016), referrals tend to be higher quality. However, in contexts where competition is more pronounced, such as among day laborers (Beaman and Magruder, 2012), the quality of referrals can decline. This may also explain why individuals need incentives to make better referrals.

Second, we contribute to a recent and growing literature that examines strategic incentives of sharing information within a social network, and how individual characteristics influence these decisions. Recent work has documented that factors ranging from political affiliation (Bandiera et al., 2023) to race (Miller and Schmutte, 2021) impact information sharing with important implications for the efficiency and fairness of information flows. Most related to our study, is a small literature which documents how competition can limit the transmission of information among firms (Cai and Szeidl, 2018; Hardy and McCasland, 2021) and individuals participating in a community activity (Vilela, 2019). We show that job-seekers share similar concerns, and communicate less about job information with their peers when they know they will have to compete with them for the job. We also show how behavior affects firms by lowering the quality of hires.

Finally, we contribute to a rich literature that examines labor market frictions in low-income countries. On the job seeker's side, evidence shows that the costs related to upskilling, and job searching, reduce the quantity and quality of job applications job-seekers make (Abebe et al., 2021; Franklin, 2017). Governments and institutions commonly employ strategies like on-the-job training, vocational training, and job-search assistance to mitigate these costs. Existing literature suggests that such measures can offer modest improvements in the short term. There is only one study to our knowledge that examines the implications of these programs within a social network: Caria et al. (2023) demonstrate that a job-search assistance intervention diminishes information sharing between program recipients and non-recipients. Treated job-seekers engage in less information exchange with their peers after directly receiving information from the program, and control job-seekers search less and have worse employment outcomes as a result. Our study complements Caria et al. (2023) by examining an additional mechanism that can influence the effectiveness of labor market interventions in the context of social network dynamics. In particular treated job-seekers engage in less information exchange with their peers if they are in direct competition with them.

Lastly, a smaller literature examines the implications of labor market frictions on the firm side. Broadly, there is growing evidence that such frictions limit firm profitability, particularly in the manufacturing sector (Crepon and Premand, 2019; de Mel et al., 2019; Alfonsi et al., 2020; Hardy and McCasland, 2023). We build on this existing body of literature by showing how job seekers' strategic considerations limit the transmission of job-related information within their social networks. Notably, information tends to be withheld from the most skilled job seekers. These strategic behaviors ultimately diminish the caliber of job applicants and subsequently affect the quality of hiring decisions made by firms.

The rest of this paper is organized as follows: Section 2 describes our setting, the, experiment and data collection, Section 3 lays out a conceptual framework behind competition and job sharing, Section 4 reports results, and Section 5 concludes.

2 Experimental Design and Data

2.1 Sample

We worked with six private colleges in Mumbai, India. These colleges cater to lower-income students across the city. Each college is divided into 'batches' or fields, which refers to the student's program of study (commerce, marketing, finance, HR). We worked with students who were about to complete their final year of college and intended to look for jobs once they graduated. Students from these colleges typically go on to work as BPO telecallers or back-office assistants. Unemployment rates are relatively high for students graduating from private colleges across India, and there is some debate as to the quality of education students receive at these institutions (Beniwal, 2023).

To recruit our sample within each college, we offered anyone who participated in the study a three-hour complementary "employability training" course. This course covered topics such as how to look for jobs using job-portals, how to build a professional CV, and how to get ready for an interview. Anyone who registered for the course became part of the sample we would subsequently engage for the next six weeks. We conducted a comprehensive baseline survey with the 496 students who registered for the course, out of a total of 2,834 students from all batches in these colleges. We collected detailed socio-demographic information, as well as English, logical, and quantitative abilities, information on students' social networks, who they talked to about employment opportunities, and the strength of their connection with their friends.

Students in our baseline sample were 20 years old on average, 60% were female, 82% were Hindu, and 60% came from the general castes (see Table A1). Students were typically from lower-middle income households with only 22% reporting a monthly family income exceeding INR 30,000 (USD 350). Their parents typically did not have higher education: only 13% of fathers and 6% of mothers had a college degree. Students reported speaking to their friends regularly about jobs. Just over half of the students reported having helped friends find a job in the past, 42% relied on friends to find jobs for themselves, and 86% discussed jobs more generally with their friends.

2.2 Experimental Design

Our experiment engaged this sample for another six weeks after the completion of the training program. Each week, we designed a small task (henceforth, our job) that required students to spend 45 minutes searching for five articles on a particular topic on Indian public policy that was of interest to a researcher at a renowned international institution. Students were tasked with finding relevant articles and summarizing them in a few sentences for the researcher. The topics changed weekly and covered issues in agricultural policy, women empowerment, education, etc. This job (and research topic) was the *same* across all college-batches in a particular week. Students were paid INR 500 (USD 6.5) for completing this task.

Second, we varied two aspects of the job before sharing with a randomly selected group students within a batch: (a) whether the job was "rival" or not (which we detail further below); and (b) whether it offered double the wage or not. Each batch in a week was then randomized to receive information about one of these four jobs: rival high-wage, rival normal-wage, non-rival high-wage, and non-rival normal wage. Our unit of randomization was at the batch-level. Within each batch we then randomly selected 20% of students each week to receive information about one of the four jobs as per the batch's treatment assignment. We call these students our 'entry-points'. To track the spread of information, entry-points received a unique referral code along with an application link via WhatsApp, which they could easily share and forward to other classmates.[3] Anyone who was not assigned to hearing about these jobs are

[3] This unique referral code (a 4 digit number) was created using a random number generator in R alongside the treatment assignment.

henceforth labelled as 'students'.[4] Since we selected a new group to hear about the job within a particular batch, the entry-points changed each week. Table A1 shows that the characteristics of these students were balanced across weeks.

The distinction between a rival job and a non-rival job influenced the strategic incentives that the entry-points faced when sharing information about the job opportunity within their social network. For a batch that was allocated to a rival job, the entry-points were invited to apply for the job, and we encouraged them to share this opportunity with other students in their batch. On the other hand, for the non-rival job, we told entry-point students that they were *guaranteed* a position should they want it, but there were additional slots available if their batch-mates wanted to apply as well. We were always prepared to hire more students than the number we informed each week, selecting the students with the highest GPA.[5] Students were unaware of the total number of slots that could be filled in any given week.

We collected all applications the day before the scheduled assignment. We ran our code to select which candidates to hire, and reminded selected students on the morning of the assignment that the work task was scheduled for 45 minutes later that afternoon. We invited students to connect to a Google link at the pre-determined time. We noted the time they connected, and the time they submitted the assignment. The students completed their assignment by searching for articles on their phones. Some students drafted their summaries of the articles directly on their phones and sent them to the team, while others submitted photographs of a handwritten summary and sent them via WhatsApp. Some students finished early, while others submitted their assignment after the allocated 45 minutes.

2.3 Data

We collect four datasets. First, a detailed baseline survey with 496 students who signed up for the employability training program, where we captured information about job-seekers' demographics (gender, GPA score, social norms) along with a list of their friends and the strength of their connection. This allows us to map out the social network.[6] Second, we conducted weekly midline surveys with our baseline sample to understand whether they had heard about our job opportunity, as well as

[4]Note that by definition, 'students' include *all* students in a batch who were not assigned to receiving the job information, including those who were not in our baseline.

[5]In rival batches we ranked all applicants by GPA and hired the most qualified set. In non-rival batches we hired the treated students, and the remaining hires were the applicants with the highest GPA.

[6]The survey asks respondents to list the friends they talk to about jobs.

who they heard it from. Third, we complement these data with information collected on applications. All applicants had to apply via google-forms so we could track applications. These forms asked for the applicant's name, gender, and a referral code. As mentioned earlier, the referral code was unique to each entry-point in a week, which enabled us to perfectly track which entry-points applicants heard the job from, thus fully characterizing the flow of information through the network. Fourth, we track who was hired each week. Note that both the data on applications and hires included information about applicants and hires regardless of whether we were able to collect any prior data on them (baseline or midline), thus enabling us to capture a comprehensive spread of information within the batch.[7]

The above datasets allow us to generate three key outcomes of interest, each of which is an indicator variable on whether a student: (i) heard about the job; (ii) applied to the job; and (iii) was hired for the job. The last outcome we are interested in is the GPA of these three groups of students.[8]

3 Conceptual Framework

We present a simple and stylized conceptual framework to guide our empirical analysis and interpret the results. This framework captures two key tradeoffs faced by individuals in deciding whether to share information with peers or not: first, competing with their peers for the job; and second, the utility gains derived from sharing information with peers. This provides with a parsimonious way to interpret our results.

Setup: Consider a pair of friends i and j who are indexed by a characteristic (such as gender, ability, etc.) X_i and X_j respectively. A job is indexed by a quality measure

[7] In practice, this dataset includes anyone from the group of 496 students we completed the baseline with, and an additional 147 students who never completed the baseline because they did not register for the initial employability program but subsequently heard about the job from an entry-point and chose to apply.

[8] Our trial was registered on the AEA RCT Registry (# AEARCTR-0007564). Although we did not create a pre-analysis plan (PAP), we identified a very parsimonious set of primary outcomes to investigate. First, we specified two primary outcomes of interest in the registry: hearing about job opportunities and actively applying for jobs. We expanded our analysis to include an investigation of who was hired, as this represents a natural extension of who applies. Second, we also specified two dimensions of heterogeneity (that we discuss in subsequent sections), namely ability and homophily. We expanded our analysis to include an investigation of the impact of being closely connected to a peer, as such connections are expected to reduce the impact of competition (in contrast to ability, which would intensify it). Following the guidance of Banerjee, Duflo, Finkelstein, Katz, Olken and Sautmann (2020), our readers may wish to interpret heterogeneity analysis on close connections as secondary analysis.

w (such as wage, amenities, etc.) so that a higher w implies a better quality job. We model the decision of an individual i who hears about a job w and has to decide whether to share this information with his/her friend j. The utility of individual i is given by:

$$\mathcal{U}_i = \Pr\Bigg(X_i, X_j, \mathbb{1}_i\{\text{Share}\}\Bigg)U(w) + \mathbb{1}_i\{\text{Share}\} \times \underbrace{\eta_{ij}\theta(w)}_{\text{Utility Channel}} \tag{1}$$

We define each term in turn. $U(w)$ is the utility derived by an individual i from working in a job w. $\Pr(X_i, X_j, \mathbb{1}_i\{\text{Share}\})$ is the probability that i is hired for a job w. This depends on the individuals own characteristics (X_i), the (endogenous) decision to share this information with j (denoted by $\mathbb{1}_i\{\text{Share}\}$) and if shared, the characteristics of j (X_j). $\Pr(X_i, X_j, \mathbb{1}_i\{\text{Share}\})$ is defined as follows:

$$\Pr\Bigg(X_i, X_j, \mathbb{1}\{\text{Share}\}\Bigg) = p(X_i) - \mathbb{1}_i\{\text{Share}\} \times \underbrace{\lambda(X_i, X_j)}_{\text{Competition Channel}} \tag{2}$$

Consider the case where there is no job sharing i.e., $\mathbb{1}_i\{\text{Share}\} = 0$. Then we denote the probability that an individual i is hired for a job w by $p(X_i)$, where $\partial p/\partial X_i \geq 0$ i.e., conditional on the job, individuals with "better" characteristics (higher ability for example) are more likely to be hired.[9]

Now consider the case where an individual decides to share information i.e., $\mathbb{1}_i\{\text{Share}\} = 1$. We assume (in a reduced-form way) that sharing information on jobs with friends might reduce the possibility that the individual gets the job. We term this the "competition channel". Moreover, the extent to which this competition matters depends on the characteristics of j relative to i. To put it more formally, we assume that sharing jobs reduces own-probability of getting a job by a function $\lambda(X_i, X_j)$, where $\partial\lambda/\partial X_j > 0$ i.e., conditional on X_i, a higher X_j would reduce i's probability of getting the job.

Lastly, we refer to $\eta_{ij}\theta(w)$ as a "utility channel". We assume that sharing information might have non-employment utility benefits for the individual, denoted by $\theta(w)$. η_{ij} captures how much an individual cares about sharing this information, which could for example be proxied by the strength of their connection or homophily considerations.

[9]There are two clarifications of note: first, we do not endogenously solve for $p(X_i)$ in equilibrium, but rather assume that it depends on the characteristics of an individual. Second, we do not distinguish between the probability of hearing and applying for the job. As we will show later, conditional on hearing about a job 75-80% of individuals apply for it, indicating that this is not an important margin.

Decision to share information: Given this setup, an individual i will share a job with his/her peer j as long as s/he receives higher utility from doing so i.e., $\mathcal{U}_i(\mathbb{1}_i\{\text{Share}\} = 1) \geq \mathcal{U}_i(\mathbb{1}_i\{\text{Share}\} = 0)$. From Equations (1) and (2), this implies:

$$\begin{aligned} \Delta\mathcal{U} &= \mathcal{U}_i(\mathbb{1}_i\{\text{Share}\} = 1) - \mathcal{U}_i(\mathbb{1}_i\{\text{Share}\} = 0) \geq 0 \\ &= \underbrace{\eta_{ij}\theta(w)}_{\text{Utility channel}} - \underbrace{\lambda(X_i, X_j)U(w)}_{\text{Competition channel}} \geq 0 \end{aligned} \tag{3}$$

4 Empirical Analysis

4.1 Does information flow?

We begin by examining whether information about the jobs we advertised was shared through social networks. There are two ways for students in our sample to hear about these jobs. Entry-point students could hear about the job directly from us, while other students could only hear about the job from their entry-point peers. Since our primary goal is to understand the flow of job information, we restrict our analysis to non-entry point students and examine whether they heard about the job each week. We begin by estimating the following regression specification:

$$Y_{ibt} = \alpha_b + \alpha_t + \beta_1 \text{Non-Rival}_{bt} + \gamma X_i + \varepsilon_{ibt} \tag{4}$$

where Y_{ibt} takes the value 1 if an individual i in batch b in week t hears about (or applies to) a job and 0 otherwise; Non-Rival_{bt} takes the value 1 if the job shared in individual i's batch (b) is non-rival in week t and 0 otherwise; α_b and α_t are batch and week fixed effects that we include to account for the stratification of treatment, and X_i controls for the number of friends individual i has. We cluster standard errors at the batch-week and individual level. The former is to account for how the treatment was administered at the batch-week level, while the latter allows for correlations within individual across weeks. From our theoretical framework (Equation 3) we anticipate that individuals will be more likely to share information about the job when the job is non-rival and the competition channel is shut down ($\beta_1 > 0$).

The results are reported in Panel A of Table 1. We see that the probability that a student hears about (Column 1) or applies (Column 2) to a job increases by 5.3 p.p. and 4.7 p.p., respectively if the job we advertise in their batch is non-rival relative to when it is rival. This represents a 30% increase. Furthermore, we can test whether the probability that a student hears about a job increases when they are directly connected to an 'entry-point' who receives this information from us, and how this varies with the "rivalness" of the job. Specifically, we estimate the following regression specification:

$$Y_{it} = \alpha_b + \alpha_t + \beta_1 \text{Non-Rival}_{bt} + \beta_{2A} \text{Rival}_{bt} \times \text{T}_{it} + \beta_{2B} \text{Non-Rival}_{bt} \times \text{T}_{it} + \gamma X_i + \varepsilon_{it} \tag{5}$$

where T_{it} takes the value 1 if at least one friend in i's social network was selected as the entry-point in week t and 0 otherwise and (Non)Rival$_{bt}$ takes the value 1 if the job shared in batch b was (non)rival and 0 otherwise.

The results are are reported in Panel B of Table 1. Column 1 shows that individuals are 25.5 p.p. (24.1 p.p.) more likely to hear about the job when they are connected to an entry-point and the job is non-rival (rival) respectively. This confirms that being directly connected matters regardless of whether the job is rival or not. Nevertheless, individuals with no connections to entry-points are 3 p.p. (25%) more likely to hear about the job when it is non-rival. This suggests that information disseminates more widely to non-connected peers when information is non-rival. We see a similar pattern emerge for applications (Column 2). While the number of application increases by 2 p.p for unconnected students when the job is non-rival, the point estimate is not statistically significant at conventional levels.

Taken together, these results highlight two key facts about how social networks affect the flow of information about jobs. First, we demonstrate the influence of strategic disincentives on sharing labor market opportunities. Specifically, jobs deemed "rival" are less likely to be shared within a group. That said, an individual's chances of hearing about a job are significantly improved when their friends know about it – both when the job is rival or non-rival.

4.2 Who Shares Information, and With Whom?

The results above show that individuals are more likely to share information about non-rival jobs than rival ones. Next, we want to understand whether these strategic decisions to share information are influenced by additional factors that could exacerbate or mitigate perceived competition for a job. First, we are interested in whether individuals share job information less with their higher-ability peers with whom they may feel more competitive. Second, we explore whether a (self-reported) measure of friendship can overcome job-seekers' tendency to withhold information about a job they will be competing for. Finally, we investigate the rate of information sharing between same-gender friends, building on a literature that suggests homophily matters.

We conduct our analysis at the pair-level (instead of at the individual level), where each pair consists of a non-entry point student (the respondent) i and their friend j, who we observe in a week t.

Ability of the individual: If individuals know they have to compete for a job they may be less likely to share information about it. One feature that could exacerbate this dynamic is if an individual j perceives their peer i to be of higher ability. Indeed, sharing information about a job with i means potentially competing with a stronger applicant pool, thus mechanically reducing j's chance of getting the job. Through the lens of our theoretical framework (Equation 3), $\partial \Delta \mathcal{U} / \partial X_j = -U(w) \partial \lambda / \partial X_j < 0$ i.e., individuals are less likely to share job information with their higher-ability peers.

We can test this hypothesis by looking at whether students are less likely to hear about the job from a lower ability entry-point peer when the job is rival as opposed to non-rival. Using a student's GPA score, we construct a binary variable for each pair ij that takes the value 1 if an individual i has a higher GPA score (and thus is defined to be higher ability) as compared to j. We then estimate the following specification for a pair of individuals ij in a week t:

$$\begin{aligned} Y_{ijt} = \alpha_b + \alpha_t + \beta_{1A}\text{Rival}_{jt} + \beta_{1B}\text{Rival}_{jt} \times 1(\text{Ability}_i > \text{Ability}_j) \\ + \beta_{2A}\text{Non-Rival}_{jt} + \beta_{2B}\text{Non-Rival}_{jt} \times 1(\text{Ability}_i > \text{Ability}_j) + \gamma X_{it} + \varepsilon_{ijt} \end{aligned} \quad (6)$$

where Y_{ijt} is a binary variable that takes the value 1 if the respondent i hears about a job from their friend j in a week t. (Non)Rival$_{jt}$ takes the value 1 if their friend j is an entry-point and receives information about a (non)rival job from us in week t i.e., it is a short-hand for (Non)Rival$_{bt} \times T_{jt}$. As before, α_b and α_t are batch and week fixed effects (the level of treatment stratification), while X_{it} controls for the number of friends for i as well as whether their batch b was rival in week t i.e., Rival$_{bt}$. Of particular interest are β_{1B} and β_{2B}. The coefficient β_{1B} measures the change in probability that an individual hears about a rival job when they have relatively higher ability than their entry-point friend as compared to when they have relatively lower ability. β_{2B} captures the same comparison for non-rival jobs. A key test of the significance of the competition channel (in line with the theoretcial framework) is if $\beta_{1B} < \beta_{2B}$ i.e., if rival jobs are less likely to be shared with high ability peers than non-rival ones.

The results are reported in Column 1 of Table 2. First, we find similar results to those discussed in Table 1: being connected to an entry-point increases the probability of hearing about a job in both rival and non-rival batches i.e., across all Columns in the table, $\widehat{\beta}_{1A} > 0$ and $\widehat{\beta}_{2A} > 0$. However, *who* receives this job information varies widely based on whether the job is rival or not. We see that a higher ability student is 7.5 p.p. (39.2%) *less* likely to hear about a rival job when their lower ability friend hears about it. On the other hand, they are 8.5 p.p. (51.5%) *more* likely to hear about it from their lower ability friend when the job is non-rival. We can comfortably reject the null hypothesis that $\hat{\beta}_{1B} = \hat{\beta}_{2B}$ (p-val: 0.02). This implies that the probability that a high-ability individual hears about a job from her low-ability friend is indeed different based on whether the job is rival or not.

Strength of the Connection: While perceptions of someone's higher-ability may mitigate information sharing, being closely connected to someone may have the opposite effect. There may be utility gains to sharing jobs with friends if individuals are altruistic and want to help their friends find jobs; if they believe that by sharing a job with their friends they are more likely to hear about an opportunity themselves in the future; or if they benefit from creating opportunities to interact with their friend by sharing job information. These channels could mitigate an individual's disutility from sharing competitive employment opportunities with his/her close friends. More formally through the lens of our theoretical framework (Equation 3), $\partial \Delta \mathcal{U} / \partial \eta_{ij} = \theta(w) \geq 0$.[10]

In our baseline survey we ask respondents to tell us for each friend, on a scale of 1 (Not Close) to 5 (Very Close), how frequently they talk to each other about employment and jobs. We then classify each pair as "close" if the respondent rates the frequency of interactions to be 4 or higher.[11] Similar to Equation (6), we estimate the following specification and report the results in Column (2) of Table 2:

$$\begin{aligned} Y_{ijt} = \alpha_b + \alpha_t + \beta_{1A}\text{Rival}_{jt} + \beta_{1B}\text{Rival}_{jt} \times \text{Close Friends}_{ij} \\ + \beta_{2A}\text{Non-Rival}_{jt} + \beta_{2B}\text{Non-Rival}_{jt} \times \text{Close Friends}_{ij} + \gamma X_{it} + \varepsilon_{ijt} \end{aligned} \quad (7)$$

We see that individuals are 10.8 p.p. and 9.3 p.p. more likely to share information on jobs with their close connections when the jobs are rival and non-rival respectively. That is, individuals are more likely to share jobs with their close connections regardless of the competition for the job ($\hat{\beta}_{1B} = \hat{\beta}_{2B}$, p-val: 0.76). This implies that having a strong connection with someone can counteract the tendency to share less due to a perceived sense of competition.

[10] This prediction assumes that being closely connected is uncorrelated with other characteristics X_j that matter for competition. If there is a correlation, then the effect of being a close connection on the probability of sharing is ambiguous when jobs are rival, but unambiguously positive when the job is non-rival.

[11] Our results are robust to alternate cutoffs (for example, 3) or whether we classify pairs as close if they report an above-median score of strength as opposed to below-median.

Same Gender: Lazarsfeld et al. (1954) coined the term "homophily" to capture the fact that socially connected individuals tend to be similar to one another. While a large literature has studied the causes and consequences of homophily in various contexts (McPherson et al., 2001; Jackson, 2021), how it affects information sharing about jobs in the presence of strategic disincentives is unclear (Batista et al., 2018). On the one hand, individuals that share an identity (gender in our case) may be able to relate more to one another, and may be more likely to share job information with each other – a "homophily channel" (a higher η_{ij} in our theory). On the other hand, job-seekers that share a certain characteristic may feel like they are more directly in competition with one-another for the job, which could reduce their propensity to share information – a "competition channel" (a higher λ in our theory). The probability of sharing information about a job with individuals of a similar identity ultimately depends on the strength of both these channels (which operate in different directions), and is therefore ambiguous when the job information is rival in nature. However, if a job is non-rival (i.e., $\lambda = 0$), then we should expect more information transmission across individuals under homophily (since $\partial \Delta \mathcal{U} / \partial \eta_{ij} \geq 0$).

We investigate this by defining a binary variable that takes the value 1 if both individuals in a ij pair are of the same gender and 0 otherwise. We then estimate the following specification for a pair of individuals ij in a week t:

$$\begin{aligned} Y_{ijt} &= \alpha_b + \alpha_t + \beta_{1A}\text{Rival}_{jt} + \beta_{1B}\text{Rival}_{jt} \times \text{Same Gender}_{ij} \\ &\quad + \beta_{2A}\text{Non-Rival}_{jt} + \beta_{2B}\text{Non-Rival}_{jt} \times \text{Same Gender}_{ij} + \gamma X_{it} + \varepsilon_{ijt} \end{aligned} \tag{8}$$

The results are reported in Column 3 of Table 2. We find that individuals are 5.3 p.p. (5.7 p.p.) less (more) likely to share a job with students of the same gender when the job is rival (non-rival). These magnitudes are statistically different from each other (p-val: 0.09). This suggests that when the competition channel is absent, homophily induces more sharing. However, the competition channel dominates homophily when jobs are rival, and same-gender friends are potentially less likely to share information.

Taken together our results indicate that the rival nature of job information can lead to certain types of job seekers being screened out of receiving job information from their peers. In particular, higher ability job-seekers, are less (more) likely to receive information on a job when it is rival (non-rival) by their relatively lower ability peers. Conversely, the strength of a friendship can mitigate these competitive effects of information sharing: individuals are more likely to share information on jobs with the friends they are closest to even if they have to compete with them for it. Lastly, homophily in information in sharing matters only to the extent that the information is non-rival in nature.

Gender Differences in Information Sharing: Extensive research indicates that women are often more hesitant to engage in labor market competition compared to men (Cashdan, 1998; Niederle and Vesterlund, 2011; Boudreau and Kaushik, 2023). Recent studies explore how this dynamic impacts career decisions (Buser et al., 2014), self-promotion behaviors (Exley and Kessler, 2022), and workplace outcomes (Flory et al., 2015). Our current context provides an opportunity to investigate a novel avenue that has yet to be explored in the literature: gender disparities in sharing competitive information. We examine this by conducting a distinct analysis for male and female job-seekers. The results, detailed in Panels A and B of Table A2, indicate that men are less likely to share job information with both high-ability peers (p-value = 0.00) and other men (p-value = 0.12) when the job is rival. We observe no such impacts among women: the estimated effects are small and statistically insignificant at conventional levels.

4.3 Impact on Application Pool and Hiring

Having established that strategic disincentives lead to certain individuals in a social network being excluded from hearing about job information, we now delve into the repercussions of this on the quality of applications received by firms. Recall from Table 1 that approximately 80% of individuals apply for a job conditional on hearing about it. This implies that the pool of applicants a firm receives is directly linked to who hears about it. We therefore focus on investigating how the composition of applicants and hires changes when the job is rival or not. We focus on one specific dimension of ability that is usually observable to employers and can affect hiring decisions, namely students GPA.[12]

Pooling applications across all batches b and weeks t, we first examine the entire ability distribution of applicants for non-rival and rival jobs. In Figure 1(a), we see that the distribution is shifted to the right for non-rival jobs relative to rival ones. We formalize this by re-estimating Equation (4) with applicants' standardized GPA as the outcome variable. As reported in Column (1) of Table 3, the ability of students who heard about the job is 0.08σ higher when the job is non-rival relative to rival. Considering the substantial conversion rate from learning about a job to applying for it, this translates into a notable increase in the quality of the applicant pool, by approximately 0.13σ, when the job is non-rival compared to rival (Column 2).

Our hiring rule was straightforward: we ranked our applicants according to their GPA and hired them until all the slots for the position were filled. Figure 1(a) shows

[12] To ease interpretation, we standardize GPA scores to have mean 0 and standard deviation 1.

a similar rightward shift in the ability distribution among hires when the job is non-rival.[13] More formally, in Column (3) of Table 3, we observe that the ability of hires is 0.38σ higher when the job in non-rival relative to rival. Taken together these results confirm that the strategic disincentives in information sharing meaningfully impact the quality of applications that a firm receives, and hence the hires it makes.

4.4 Can wages help?

In the previous sections, we demonstrate that job-seekers share less information with their peers in a competitive setting, and this has consequences for labor market hiring. Specifically, it suggests that firms that rely heavily on social networks to spread information about job opportunities might end up with lower quality applicant pools, and hires, than they expected. To mitigate this effect and motivate high-quality candidates to apply for their job openings, firms could enhance the job's appeal by increasing the wage.

While conventional labor supply models suggest that increased wages should attract higher quality candidates, these models do not consider the distinct dynamics that come into play within social networks. In particular, increasing the wage makes a job more appealing, but this could elicit two distinct responses from a job-seeker thinking about whether or not to share the job opportunity. First, there is the competition channel: a higher paying job is less enticing to share because the cost of losing the job to a potential competitor has increased. Second, there is the altruism channel: a higher paying job is more attractive to share because the warm glow from sharing a better job with a friend goes up. Captured more formally through the lens of our model (Equation 3), $\partial \triangle \mathcal{U} / \partial w = \eta_{ij}\theta'(w) - \lambda(X_i, X_j)U'(w)$. Therefore, whether $\partial \triangle \mathcal{U} / \partial w$ is greater than or less than 0 depends on how strongly a change in wages impacts the competitive and altruism channels. For example, job-seekers will share less if the competition channel outweighs the altruism channel, thus resulting in a lower quality applicant pool for firms to choose from despite an increase in wages.

To investigate this further, we embedded a sub-experiment by cross-randomizing whether the job was rival or not with a high or normal wage. This meant that in some batch-weeks we doubled the wage to INR 1000 for 45 mins work (high-wage category), as compared to the status-quo "normal-wage" of INR 500. This created four types of jobs that could be shared in any given week: "high-wage, rival", "high-wage, non-rival", "normal-wage, rival", "normal-wage, rival".

[13] Having a better pool of applicants on average doesn't necessarily guarantee better hires. What matters is the quantity and quality of candidates at the top of the distribution, as those are the ones we ultimately hire.

Similar to Equation (4), we can then estimate the following regression specification:

$$Y_{ibt} = \alpha_b + \alpha_t + \beta_1 \text{High Wage}_{bt} + \beta_2 \text{Non-Rival}_{bt} + \beta_3 \text{High Wage}_{bt} \times \text{Non-Rival}_{bt} \\ + \gamma X_{it} + \varepsilon_{ibt} \tag{9}$$

where all the variables remain the same as in Equation (4). In addition, High-Wage$_{bt}$ takes the value 1 if the job shared in individual i's batch (b) was high-wage in week t, and and 0 otherwise.

The above experimental design provides us with multiple insights on how competition interacts with changes in the quality of the job (wages in our case) and subsequently impacts information sharing within the social network. First, β_1 estimates the causal impact of doubling the wage on information sharing within the social network (and subsequently on the quality of applicants and hires). Second, we can isolate whether competition dampens how much information about high-wage jobs is shared by comparing the information sharing of high-wage rival jobs (β_1) to high-wage non-rival jobs ($\beta_1 + \beta_2 + \beta_3$) i.e., we can test whether $\beta_2 + \beta_3 = 0$. Lastly, we can compare two reasonable strategies that firms usually follow to attract better talent: increasing wages (β_1) or using referrals to shut down competition in information sharing (β_2).

Table 4 shows that doubling the wage (for rival jobs) does not significantly increase the probability that individuals hear about it (Column 1) or apply for it (Column 2). Indeed the coefficients– $\hat{\beta}_1 = -0.009$ (Column 1), and $\hat{\beta}_1 = -0.007$ (Column 2) – are very small in magnitude and statistically insignificant at conventional levels. However, shutting down competition by making the high-wage job information non-rival significantly induces more sharing of information on this high-wage job. Specifically, the number of students who hear and apply to this job increases by 8.2 p.p (Column 1) and 7.9 p.p (Column 2), respectively. More formally as we show in Table 4, we can comfortably reject $\beta_2 + \beta_3 = 0$ at conventional levels (p-val: 0.02).

Lastly, we investigate how the quality of applicants and hires is affected in Table 5. We find that doubling the wage improves the quality of applicants and hires by around 0.1σ and 0.08σ respectively (Columns 2 and 3). However, this impact is much smaller than what we see when these high-wage jobs are also non-rival. Indeed, the quality of hires (applicants) improves by 0.354σ (0.037σ), when the high-wage job is non-rival relative to rival.We recover these estimates by adding $\beta_2 +$ $\beta_3 becausethereferenceisahighwagejobbecomingnon-rival.Togaugetheimportanceofthecompetitioncha$ $of - theenvelopecalculationcomparing\beta_1$ and β_2 suggests that to get the same increase in the pool of hires (applicants), a firm would have to increase wages by 4.8 (2.3 times).[14]

[14] To see this, note from Column (3) that the average quality of hires is 0.081σ higher when wages

5 Conclusion

Social networks are central to well functioning labor markets in low income countries. Firms rely on these networks to disseminate information about new job openings and attract high quality candidates. Any frictions that are created by job-seekers competing for jobs could have negative impacts on the quality of matches. We explore this phenomenon empirically with Indian college students about to enter the job market. We randomly seed their social networks with jobs that are either rival or non-rival. We find that when a job is rival, information about that job is less likely to travel in the network, and is less likely to reach high ability job seekers. This is especially true among men. We find that firms can offer higher wages to helps attract better quality candidates. However, firms should not anticipate achieving the same level of quality improvement as they might if competition were not causing job-seekers to share information less frequently.

These results might explain why the literature finds that the impact of referrals varies across contexts. Specifically, they suggest that whenever competition-related worries are prominent (as seen among day laborers, for instance), the quality of referrals might be lower compared to situations where job-seekers are less concerned about their future job prospects (such as among full-time employees). They also have important have implications for policy makers seeking to improve matches between employers and job-seekers: they suggest there is value in supporting technologies that allow job information to flow outside social networks — e.g., job portals, or information campaigns at universities — or making sure that referees are properly incentivized.

References

Abebe, Girum, A. Stefano Caria, and Esteban Ortiz-Ospina, "The Selection of Talent: Experimental and Structural Evidence from Ethiopia," *American Economic Review*, June 2021, *111* (6), 1757–1806.

Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali, "Tackling Youth Unemployment: Evidence From a Labor Market Experiment in Uganda," *Econometrica*, 2020, *88* (6), 2369–2414.

Bandiera, Oriana, Robin Burgess, Erika Deserranno, Ricardo Morel, Munshi Sulaiman, and Imran Rasul, "Social Incentives, Delivery Agents, and the Effectiveness

are doubled, and 0.391σ higher when information is non-rival. Therefore, to get the *same* increase in the average applicant quality (assuming linear treatment effects), wages would have to be 0.391/0.081 (β_2/β_1) i.e., 4.8 times higher.

of Development Interventions," *Journal of Political Economy Microeconomics*, 2023, *1* (1), 162–224.

Banerjee, Abhijit, Esther Duflo, Amy Finkelstein, Lawrence F. Katz, Benjamin A. Olken, and Anja Sautmann, "In Praise of Moderation: Suggestions for the Scope and Use of Pre-Analysis Plans for RCTs in Economics," *NBER Working paper*, 2020.

Batista, Cátia, Marcel Fafchamps, and Pedro C Vicente, "Keep It Simple: A field experiment on information sharing in social networks," Technical Report, National Bureau of Economic Research 2018.

Beaman, Lori and Jeremy Magruder, "Who Gets the Job Referral? Evidence from a Social Networks Experiment," *American Economic Review*, 2012, *102* (7), 3574–93.

Beniwal, Bibhudatta Pradhan and Vrishti, "Worthless Degrees Are Creating an Unemployable Generation in India," *The Economic Times*, April 2023.

Boudreau, Kevin J. and Nilam Kaushik, "Gender Differences in Responses to Competitive Organization? A Field Experiment on Differences Between STEM and Non-STEM Fields from an Internet-of-Things Platform," *Organization Science*, 2023.

Brown, Meta, Elizabeth Setren, and Giorgio Topa, "Do informal referrals lead to better matches? Evidence from a firm's employee referral system," *Journal of Labor Economics*, 2016, *34* (1), 161–209.

Buser, Thomas, Muriel Niederle, and Hessel Oosterbeek, "Gender, competitiveness, and career choices," *The quarterly journal of economics*, 2014, *129* (3), 1409–1447.

Bó, Ernesto Dal, Frederico Finan, and Martín A. Rossi, "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service*," *The Quarterly Journal of Economics*, 04 2013, *128* (3), 1169–1218.

Cai, Jing and Adam Szeidl, "Interfirm Relationships and Business Performance," *The Quarterly Journal of Economics*, 2018, *133* (3), 1229–1282.

Caria, Stefano, Simon Franklin, and Marc Witte, "Searching with Friends," *Journal of Labor Economics*, 2023, *forthcoming.*

Cashdan, Elizabeth, "Are men more competitive than women?," *British Journal of Social Psychology*, 1998, *37* (2), 213–229.

Crepon, Bruno and Patrick Premand, "Direct and Indirect Effects of Subsidized Dual Apprenticeships," *Working Paper*, 2019.

de Mel, Suresh, David McKenzie, and Christopher Woodruff, "Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises," *American Economic Journal: Applied Economics*, January 2019, *11* (1), 202–35.

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker, "Referral-Based Job Search Networks," *The Review of Economic Studies*, 2016, *83* (2), 514–546.

Exley, Christine L and Judd B Kessler, "The gender gap in self-promotion," *The Quarterly Journal of Economics*, 2022, *137* (3), 1345–1381.

Fafchamps, Marcel and Alexander Moradi, "Referral and Job Performance: Evidence from the Ghana Colonial Army," *Economic Development and Cultural Change*, 2015, *63* (4), 715–751.

Flory, Jeffrey A, Andreas Leibbrandt, and John A List, "Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions," *The Review of Economic Studies*, 2015, *82* (1), 122–155.

Franklin, Simon, "Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies," *The Economic Journal*, 11 2017, *128* (614), 2353–2379.

Granovetter, Mark S, "The strength of weak ties," *American journal of sociology*, 1973, *78* (6), 1360–1380.

Hardy, Morgan and Jamie McCasland, "It Takes Two: Experimental Evidence on the Determinants of Technology Diffusion," *Journal of Development Economics*, 2021, *149*, 102600.

_ and _ , "Are Small Firms Labor Constrained? Experimental Evidence from Ghana," *American Economic Journal: Applied Economics*, April 2023, *15* (2), 253–84.

Heath, Rachel, "Why do firms hire using referrals? Evidence from Bangladeshi garment factories," *Journal of Political Economy*, 2018, *126* (4), 1691–1746.

Jackson, Matthew O, "Inequality's Economic and Social Roots: The Role of Social Networks and Homophily," *Available at SSRN 3795626*, 2021.

Kramarz, Francis and Oskar Nordström Skans, "When Strong Ties are Strong: Networks and Youth Labour Market Entry," *Review of Economic Studies*, 2014, *81* (3), 1164–1200.

Kugler, Adriana D, "Employee referrals and efficiency wages," *Labour economics*, 2003, *10* (5), 531–556.

Lazarsfeld, Paul F, Robert K Merton et al., "Friendship as a Social Process: A Substantive and Methodological Analysis," *Freedom and Control in Modern Society*, 1954, *18* (1), 18–66.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook, "Birds of a feather: Homophily in social networks," *Annual review of sociology*, 2001, pp. 415–444.

Miller, Conrad and Ian M Schmutte, "The Dynamics of Referral Hiring and Racial Inequality: Evidence from Brazil," Working Paper 29246, National Bureau of Economic Research September 2021.

Montgomery, James D, "Social networks and labor-market outcomes: Toward an economic analysis," *The American economic review*, 1991, *81* (5), 1408–1418.

Niederle, Muriel and Lise Vesterlund, "Gender and competition," *Annu. Rev. Econ.*, 2011, *3* (1), 601–630.

Pallais, Amanda and Emily Glassberg Sands, "Why the Referential Treatment? Evidence from Field Experiments on Referrals," *Journal of Political Economy*, 2016, *124* (6), 1793–1828.

Trimble, Lindsey B and Julie A Kmec, "The role of social networks in getting a job," *Sociology Compass*, 2011, *5* (2), 165–178.

Vilela, Inês, "Diffusion of rival information in the field," Technical Report, Working Paper 2019.

Tables

Table 1: Heard about and Applied to a Job

	(1) Heard	(2) Applied
Panel A:		
Non-Rival$_{bt}$	0.053	0.047
	(0.020)**	(0.022)**
Control Mean	0.18	0.14
Observations	2535	2535
Panel B:		
Non-Rival$_{bt}$	0.030	0.020
	(0.017)*	(0.017)
Non-Rival$_{bt}$ $\times$ T$_{it}$	0.255	0.230
	(0.057)***	(0.054)***
Rival$_{bt}$ $\times$ T$_{it}$	0.241	0.188
	(0.048)***	(0.044)***
Control Mean	0.12	0.08
Observations	2388	2388

Notes: This table shows whether the rival/non-rival nature of the job affects the probability of hearing (Column 1) or applying (Column 2) to the job. The sample is restricted to non-entry point students in week t. The dependent variable in Column 1 takes the value 1 if i has heard about the job in week t and 0 otherwise. The dependent variable in Column 2 takes the value 1 if i has applied to the job in week t and 0 otherwise. In Panel B, we drop respondents who were not in the baseline sample but applied for a job. (Non)Rival$_{bt}$ takes the value 1 if batch b was assigned to the (Non)Rival treatment in week t and 0 otherwise. T_{it} takes the value 1 if at least one friend of individual i was an entry-point in week t and 0 otherwise. All regressions include batch and week fixed effects. Standard errors are clustered at the week-batch and individual level. $* \ p < 0.10$, $** \ p < 0.05$, $*** \ p < 0.01$

Table 2: Heard About and Job-seeker Characteristics

	(1) Ability	(2) Close Friends	(3) Same Gender
Rival_{jt} (β_{1A})	0.191	0.090	0.190
	(0.036)***	(0.022)***	(0.048)***
Rival_{jt} $\times$ X (β_{1B})	-0.075	0.108	-0.053
	(0.038)*	(0.029)***	(0.045)
Non-Rival_{jt} (β_{2A})	0.165	0.143	0.143
	(0.033)***	(0.033)***	(0.041)***
Non-Rival_{jt} $\times$ X (β_{2B})	0.085	0.093	0.057
	(0.061)	(0.045)**	(0.049)
$\beta_{1B} = \beta_{2B}$	0.02	0.76	0.09
Observations	2781	3470	3470

Notes: This table shows whether individual characteristics affect how information disseminates when a job is rival or not. The sample is restricted to ij pairs where individual i was assigned to the non-entry point group in week t. The dependent variable takes the value 1 if i heard about the job in week t from friend j, and 0 otherwise. (Non)Rival_{jt} takes the value 1 if friend j was assigned to the (Non)Rival treatment in week t and 0 otherwise. In Column (1), X is an indicator for 1(Ability_i >Ability_j), which takes the value 1 if individual i has a higher ability than j. In Column (2), X is an indicator for Same Gender_{ij}, which takes the value 1 if both i and j are of the same gender and 0 otherwise. Similarly, in Column (3), X is an indicator for Close Friend_j, which takes the value 1 if both i and j are "close friends" and 0 otherwise. All regressions include batch and week fixed effects. Standard errors are clustered at the week-batch and individual level. $* \ p < 0.10$, $** \ p < 0.05$, $*** \ p < 0.01$

Table 3: Ability of Students

	(1) Heard	(2) Applied	(3) Hired
Non-Rival_{bt}	0.083	0.130	0.381
	(0.037)**	(0.045)***	(0.082)***
Control Mean	0.08	0.07	0.13
Observations	688	462	304

Notes: This table shows how the ability of students who hear (Column 1), apply (Column 2) and are hired (Column 3) changes when a job is rival or not. The sample is restricted to respondents assigned to control group in week t and respondents who were in the rival treatment group in week t. In Column (1), the sample is restricted to students who heard about the job, in Column (2) the sample is restricted to students who applied for the job, and in Column (3) the sample is restricted to students who were hired. The dependent variable is the respondent's standardized GPA score. Non-Rival_{bt} takes the value 1 if the batch b was assigned to the non-rival treatment in week t and 0 otherwise. High-wage_{bt} takes the value 1 if the batch b was assigned to the high-wage treatment in week t and 0 otherwise. All regressions include batch and week fixed effects. Standard errors are clustered at the week-batch and individual level. $* \ p < 0.10$, $** \ p < 0.05$, $*** \ p < 0.01$

Table 4: Heard about and Applied to a Job (High-Wage)

	(1) Heard	(2) Applied
High-wage$_{bt}$ (β_1)	-0.009	-0.007
	(0.024)	(0.025)
Non-Rival$_{bt}$ (β_2)	0.022	0.013
	(0.024)	(0.027)
High-wage$_{bt}$ $\times$ Non-Rival$_{bt}$ (β_3)	0.060	0.066
	(0.039)	(0.044)
$\beta_1 = \beta_2$	0.13	0.45
$\beta_2 + \beta_3 = 0$	0.01	0.02
Control Mean	0.19	0.14
Observations	2535	2535

Notes: This table shows whether the rival/non-rival/high-wage/normal-wage nature of the job affects the probability of hearing (Column 1) or applying (Column 2) to the job. The sample is restricted to respondents assigned to control group in week t. The dependent variable in Column 1 takes the value 1 if i has heard about the job in week t and 0 otherwise. The dependent variable in Column 2 takes the value 1 if i has applied to the job in week t and 0 otherwise. High-wage$_{bt}$ takes the value 1 if the batch b was assigned to the high-wage treatment in week t and 0 otherwise. All regressions include batch and week fixed effects. Standard errors are clustered at the week-batch and individual level. $*\ p < 0.10$, $**\ p < 0.05$, $***\ p < 0.01$

Table 5: Ability of Students (High-Wage)

	(1) Heard	(2) Applied	(3) Hired
High-wage$_{bt}$ (β_1)	-0.036	0.100	0.081
	(0.051)	(0.028)***	(0.047)*
Non-Rival$_{bt}$ (β_2)	0.117	0.232	0.391
	(0.080)	(0.045)***	(0.035)***
Non-Rival$_{bt}$ $\times$ High-wage$_{bt}$ (β_3)	-0.059	-0.195	-0.037
	(0.103)	(0.028)***	(0.144)
$\beta_1 = \beta_2$	0.06	0.05	0.00
$\beta_2 + \beta_3 = 0$	0.41	0.42	0.01
Control Mean	0.03	0.00	0.04
Observations	688	462	304

Notes: This table shows how the ability of students who hear (Column 1), apply (Column 2) and are hired (Columne 3) changes when a job is rival/high-wage or not. The sample is restricted to respondents assigned to control group in week t and respondents who were in the rival treatment group in week t. In Column (1), the sample is restricted to students who heard about the job, in Column (2) the sample is restricted to students who applied for the job, and in Column (3) the sample is restricted to students who were hired. The dependent variable is the respondent's standardized GPA score. Non-Rival$_{bt}$ takes the value 1 if the batch b was assigned to the non-rival treatment in week t and 0 otherwise. High-wage$_{bt}$ takes the value 1 if the batch b was assigned to the high-wage treatment in week t and 0 otherwise. All regressions include batch and week fixed effects. Standard errors are clustered at the week-batch and individual level. $*\ p < 0.10$, $**\ p < 0.05$, $***\ p < 0.01$

Figures

Figure 1: Ability Distribution of Applicants and Hires

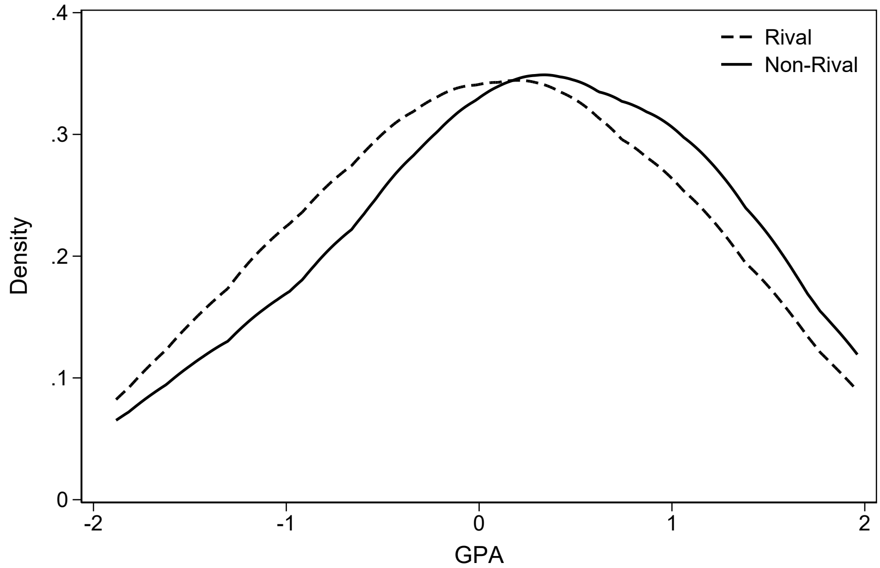

(a) Job Applicants

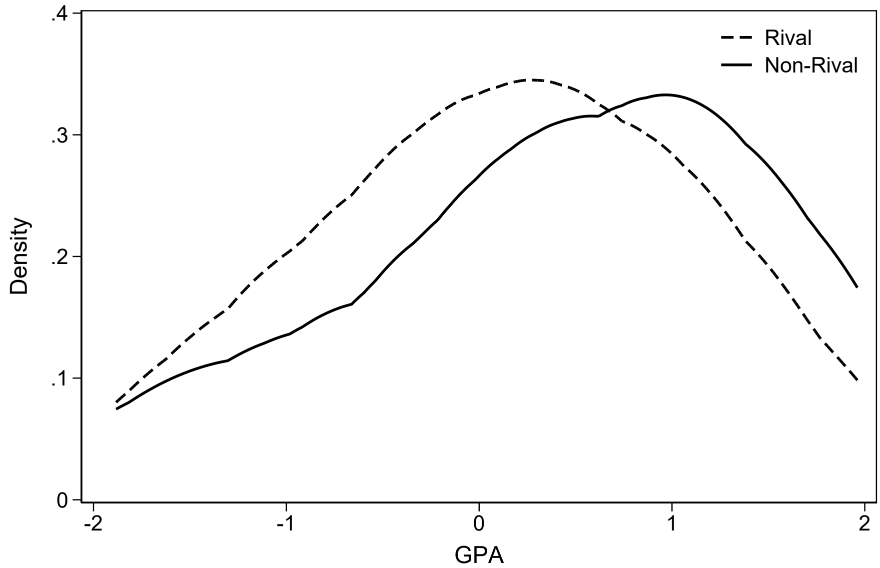

(b) Job Hires

Appendix

Table A1: Balance across entry-points and control students, all weeks

	Control Students	Entry-Points	p-value	N
Age	20.5	20.4	0.10	2976
Female (%)	57.7	59.4	0.48	2976
GPA	6.9	6.9	0.71	2964
Relgion: Hindu (%)	81.2	84.5	0.06*	2976
Caste: General (%)	61.5	61.4	0.97	2898
Mother completed college (%)	5.9	6.5	0.64	2976
Father completed college (%)	12.9	13.8	0.59	2976
Parents' monthly income > INR 30000 (%)	22.8	20.6	0.29	2604
Ever helped friend find jobs? (%)	53.6	56.0	0.30	2976
Rely on friends to find a job? (%)	41.8	41.7	0.97	2976
Ever talk to friends about jobs? (%)	86.6	85.9	0.63	2976
Speak to classmates about jobs? (%)	64.1	64.3	0.92	2976

Notes: This table pools all individuals across weeks and checks the balance across characteristics of entry-, for all weeks. The study sample is included. $* \ p < 0.10$, $** \ p < 0.05$, $*** \ p < 0.01$

Table A2: Job Sharing and Gender

	(1) Ability	(2) Close Friends	(3) Same Gender
Panel A: Males			
Rival_{jt}	0.229	0.132	0.238
	(0.047)***	(0.038)***	(0.056)***
$\text{Rival}_{jt} \times$ X	-0.142	0.062	-0.107
	(0.047)***	(0.047)	(0.059)*
Non-Rival_{jt}	0.090	0.111	0.132
	(0.040)**	(0.038)***	(0.061)**
$\text{Non-Rival}_{jt} \times$ X	0.117	0.102	0.036
	(0.075)	(0.057)*	(0.068)
$\beta_{1B} = \beta_{2B}$	0.00	0.51	0.12
Observations	1111.00	1491.00	1491.00
Panel B: Females			
Rival_{jt}	0.166	0.057	0.119
	(0.040)***	(0.022)**	(0.075)
$\text{Rival}_{jt} \times$ X	-0.033	0.142	0.019
	(0.059)	(0.041)***	(0.077)
Non-Rival_{jt}	0.213	0.175	0.156
	(0.046)***	(0.043)***	(0.063)**
$\text{Non-Rival}_{jt} \times$ X	0.073	0.081	0.073
	(0.070)	(0.061)	(0.081)
$\beta_{1B} = \beta_{2B}$	0.22	0.39	0.59
Observations	1670.00	1979.00	1979.00

Notes: This table shows whether individual characteristics affect how information disseminates when a job is rival or not separately for males (Panel A) and females (Panel B). The sample is restricted to ij pairs where individual i was assigned to the non-entry point group in week t. The dependent variable takes the value 1 if i heard about the job in week t from friend j, and 0 otherwise. (Non)Rival_{jt} takes the value 1 if friend j was assigned to the (Non)Rival treatment in week t and 0 otherwise. In Column (1), X is an indicator for 1(Ability_i >Ability_j), which takes the value 1 if individual i has a higher ability than j. In Column (2), X is an indicator for Same Gender_{ij}, which takes the value 1 if both i and j are of the same gender and 0 otherwise. Similarly, in Column (3), X is an indicator for Close Friend_j, which takes the value 1 if both i and j are "close friends" and 0 otherwise. All regressions include batch and week fixed effects. Standard errors are clustered at the week-batch and individual level. $* \; p < 0.10$, $** \; p < 0.05$, $*** \; p < 0.01$.