

NBER WORKING PAPER SERIES

COMPETITIVE JOB SEEKERS:
WHEN SHARING LESS LEAVES FIRMS AT A LOSS

Gaurav Chiplunkar
Erin M. Kelley
Gregory V. Lane

Working Paper 32171
<http://www.nber.org/papers/w32171>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2024

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, and 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). The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Gaurav Chiplunkar, Erin M. Kelley, and Gregory V. Lane. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Competitive Job Seekers: When Sharing Less Leaves Firms at a Loss
Gaurav Chiplunkar, Erin M. Kelley, and Gregory V. Lane
NBER Working Paper No. 32171
February 2024
JEL No. L14,M51,O12

ABSTRACT

We study how job-seekers share information about jobs within their social network, and its implications for firms. We randomly increase the amount of competition for a job and find that job-seekers are: (i) less likely to share information about the job with their peers; and (ii) choose to selectively share it with fewer higher ability peers. This lowers the quality of applicants received by firms, subsequent hires made, and performance on the job — suggesting that firms who rely on social networks to disseminate job information may see lower quality applicants than expected for their most competitive positions. While randomly offering higher wages attracts better talent, it is not able to fully overcome these strategic disincentives in information sharing

Gaurav Chiplunkar
University of Virginia
100, Darden Blvd.
Charlottesville, VA 22905
chiplunkarg@darden.virginia.edu

Erin M. Kelley
World Bank
erinmkelley@worldbank.org

Gregory V. Lane
Harris School of Public Policy
University of Chicago
1307 E. 60th Street
Chicago, IL 60637
and NBER
laneg@uchicago.edu

A randomized controlled trials registry entry is available at
<https://www.socialscienceregistry.org/trials/7564>

1 Introduction

A well-functioning labor market relies on firms' ability to identify and match with the most qualified candidates. Firms aim to attract the best candidates by promoting their job vacancies (through advertising or referral networks), and enhancing their jobs' appeal (by offering higher wages or improved benefits). However, the success of these strategies depends on how well this information reaches suitable candidates (Caria et al., 2024). Since jobs are competitive, information is unlikely to flow smoothly within social networks – when a job-seeker shares details about a job with their peers, it reduces their own chance of securing the position (Beaman et al., 2018). This means that the effectiveness of firms' hiring strategies depends on job-seekers' strategic response to the perceived competition for a job. While accurately measuring how firm strategies and job-seeker behaviors interact is crucial for understanding the efficiency of job matches and its resulting impact on the labor market, it is also challenging to do, since both firm and job-seeker responses are usually determined in equilibrium.

In this paper, we focus on disentangling these dynamics by isolating how firms' hiring strategies and job-seekers' strategic behavior affect the dissemination of job information and subsequent labor market outcomes. Specifically, we design a randomized control trial to achieve three goals – first, to document whether job-seekers exhibit “strategic disincentives” in information sharing i.e., share less information about jobs with specific types of peers when the competition for these jobs is made salient to them. Second, we explore how these strategic disincentives impact the talent pool of applicants received by a firm, the hires that they make, and their performance on the job. Finally, we assess how these strategic disincentives interact with a common strategy firms use to attract better talent – offering higher wages.

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, the research team created a short-term job opportunity where students were given 45 minutes to search and summarize articles on a particular topic of interest for a renowned international institution. Each student had to complete their task individually, and there was no interaction between them on the job. Information on this job opportunity was shared with a randomly selected 20% of students within each batch. Students who heard about the job directly from us were labelled as ‘entry-points’. Since we randomly selected a new group of students to hear about the jobs each week (for a month) the set of entry-points changed from one week to the next. Each job we shared was randomly assigned to be either ‘rival’ or ‘non-rival’ across batches. 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 could still share the information with their peers (who had to submit an application). To further test whether the dynamics we observe change as the quality of the job improves, we also experimentally varied whether the job offered a high wage (INR 1,000 or 12 USD) relative to the status-quo (INR 500 or 6 USD). The assignment of whether a job was categorized as rival/non-rival, and high wage or not, was done randomly at the batch-week level. In other words, all participating batches were randomly allocated to receive information about a rival or high-wage job each week, and 20% of students *within* the batch were selected as entry-points and heard about the job directly from the research team.

We are interested in how information about jobs flows through social networks and how the rival nature of this information impacts: (i) the probability this information is shared; (ii) the characteristics of who receives the job information; (iii) the pool of applicants received for the job posting; (iv) who is ultimately hired for the job; and (v) their performance on 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 and who they heard it from. Third, all interested job-seekers were required to submit an online application to the job, thus allowing us to capture data on applicants. Finally, we were willing to hire multiple students for the job and applied a simple hiring rule: applicants with the highest GPA, and any entry-point applicants in the non-rival batches, were invited to complete the job. Students submitted their work, which the research team graded anonymously on a scale from 1 to 10.

We provide a stylized conceptual framework to motivate our analysis. Specifically, we show that an individual’s decision to share information is determined by two channels: a “competition channel”, where sharing job information with friends reduces one’s own probability of getting the job; and a “utility channel”, where sharing information with friends confers utility (from feelings of reciprocity or warm glow for example). The amount of information sharing depends on which of these two channels dominates. This in turn depends on whether the job is rival or not, and on the characteristics of the individuals sharing the job relative to their peers (their ability, how closely connected they are, their homophily, etc.).

We document four main findings. First, we find that job-seekers were more likely to share information about a job when they did not have to worry about competing for it. 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. (25%) more likely to hear about the job when it was non-rival relative to rival. This indicates that job information was more likely to spread beyond immediate connections to the entry-points when it was non-rival in nature.

Second, we investigate three factors that may exacerbate (or mitigate) the nature of competition between job-seekers. We find that job-seekers were strategic about *who* they shared information with when the job was rival and they knew they had to compete for it. First, 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 friendship connections. Students who reported a close connection to an entry-point at baseline (as compared to those who did not) were more likely to hear about this information, regardless of whether the job was rival or not. Specifically, students were 10.8 p.p. more likely to hear about a rival job when they were closely connected to the entry point, and 9.3 p.p. more likely to hear about a non-rival job. 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 11 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 outweighed any preferences for same gender connections. Lastly, these effects on ability and homophily were driven entirely by men, who appear less likely to share information with other high-ability, male peers when the job information was rival. This result is consistent with a larger literature that finds more competitive behavior among men in the labor market (Cashdan, 1998; Niederle and Vesterlund, 2011; Boudreau and Kaushik, 2023).

Third, we establish that the above results have implications for the quality of candidates (as measured by their GPA scores) that apply and are hired for a job, as well as their performance on the job. 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 a systematic rightward shift of the ability distribution of individuals who applied to, and were ultimately hired for, the job when it was non-rival. Specifically, the

GPA of students who *heard* about the job was 0.08σ higher on average when the job was non-rival relative to rival. This translated into an increase in the ability of the applicant pool (0.13σ), and of those who were hired (0.38σ). We can also show how this resulted in improved job performance using several indicators. Students hired from non-rival batches (as compared to rival ones) were 11.5% more likely to show up for the job, and 32% more likely to finish it within the stipulated time allotted (45 minutes). Moreover, they took lesser time, and had better quality submissions as well.

These 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, when social network dynamics are present, increasing the wage amplifies both the competition and utility channels presented in the conceptual framework above. On the one hand, the cost of sharing job information (competition channel) increases as informed job-seekers lose more if their peers secure the higher-paying job. On the other, the benefits of sharing the job (utility channel) also increase as informed job-seekers derive greater satisfaction from helping their friends access a higher-paying job. Since these channels work in opposite directions, it is theoretically ambiguous whether increasing wages overcomes the strategic disincentives of sharing job information within the social network. If the competition channel is stronger than the utility channel, job-seekers may be less likely to share the job with their peers, and employers may find themselves with a lower quality candidate pool than they initially expected despite offering higher wages.

Fourth, we find that doubling the wage among rival jobs attracted a better pool of applicants (0.10σ) and hires (0.08σ), relative to the status-quo rival job. These hires also had a 14.5% higher evaluation score on their submission as well. This confirms that higher wages can indeed attract better talent. However, a key advantage of our setting is that we can also identify how much better candidates' ability would have been, if higher wages could be offered without triggering a competitive response among job-seekers i.e., if the job information was non-rival in nature. We find that the ability of hired candidates among high-wage *non-rival* jobs and the evaluation score on their submissions, increases by an additional 0.35σ and 19.1% respectively, relative to high-wage *rival* jobs. Put differently, firms would have to increase the wage by 3-6 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 strategically shared less information about jobs when they were concerned about having to compete for them. In

particular, they shared less with peers they perceived to be higher ability than themselves, and thus a greater competitive threat. This behavior has significant implications for the labor market, as it led to a reduction in the overall quality of applications and hired candidates that firms received. While firms could in principle compensate for this by offering higher paying, more attractive jobs, our results suggest that this increase in wages would have to be substantial to attract higher quality candidates in the presence of strategic disincentives. These results highlight the potential drawbacks for firms relying solely on social networks to disseminate job information, and motivate interventions to facilitate a broader dissemination of job information.¹

This study contributes to multiple strands of the literature. First, 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 that 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 have to compete with them for the job. We also show how this behavior affects firms by lowering the quality of their hires. Importantly, our design allows us to disentangle equilibrium dynamics: by varying both the wage and the level of competition for a job, we can examine how both sides of the labor market (job-seekers and firms) interact.

Second, we contribute to a rich literature that examines labor market frictions in low-income countries. On the job-seekers' side, evidence shows that the costs related to upskilling, and job searching, reduce the quantity and quality of job applications job-seekers make ([Abebe, Caria and Ortiz-Ospina, 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, Franklin and Witte \(2023\)](#) 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

¹Online 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.

search less and have worse employment outcomes as a result. Our study complements [Caria et al. \(2023\)](#) by examining another important 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.

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.

Finally, we contribute to a large literature on the importance of social networks in referring job-seekers in the labor market that follows the seminal work of [Granovetter \(1973\)](#) and [Montgomery \(1991\)](#). 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](#); [Fafchamps and Moradi, 2015](#)). Our results might help 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.

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 to formalize the role of competition in impacting job sharing, while Section 4 reports results. Section 5 concludes.

2 Experimental Design and Data

2.1 Sample Description

We worked with six private colleges in Mumbai, India. These colleges cater to lower-income students across the city. Each college consists of students across different programs of study (such as commerce, marketing, finance, HR, etc.), which we term as a ‘batch’. 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 our sample, and we subsequently engaged them for the next six weeks.² There were 496 students in these batches who registered for the training course out of a total of 2,834 students. We conducted a comprehensive baseline survey with these 496 students, where 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). Compared to a nationally representative sample of the 68th Round of the National Sample Survey (NSS) from 2011-12 as in [Banerjee and Chiplunkar \(2023\)](#), our sample was slightly younger, with a lower concentration from scheduled castes/ scheduled tribes/ other backward classes. Students in our sample were 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

²The reason for offering this training course was twofold. First, the participating colleges wanted us to help their students improve their job search process; and second, the training course helped us get students who were actually interested in looking for a job and for whom the intervention (described later) would be relevant.

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 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. Each student was tasked with finding relevant articles and summarizing them in a few sentences for the researcher. Each student had to complete their task individually and there was no interaction between them on the job. 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.

Randomization Procedure: We varied two aspects of the job before sharing it 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 i.e., INR 1000, or not. Everything else about the job (such as the topic students had to summarize) remained the same across all batches during a given week. The experimental design then followed a two-stage process. In the first stage, for a given week, a batch was randomly allocated to receive information about one of the four categories of jobs: (i) rival high-wage; (ii) rival normal-wage; (iii) non-rival high-wage; and (iv) non-rival normal wage. In the second stage, for a given week, we then randomly selected 20% of students within each batch to receive information about this job. 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 henceforth labelled as ‘students’.⁴ Note that since a new group of students within a batch was selected to hear about the job, the entry-points changed every week. Table A1 shows that the characteristics of these entry-points relative to the other students was balanced across weeks.

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

⁴Note 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.

Rival and Non-Rival Jobs: We now elaborate on the distinction between a rival and non-rival job. 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, entry-points were invited to apply for the job and would be notified if they were selected for it. For a non-rival job on the other hand, entry-points were *guaranteed* the job if they applied for it. In both cases (rival and non-rival), entry-points were encouraged to share this job opportunity with their peers since we were prepared to hire multiple students from a batch for the job.

Hiring Criteria and Job Activity: Our selection criteria from the application pool was straightforward: in rival batches, we ranked all applicants by GPA and hired the most qualified set. In non-rival batches, we hired all entry-point students (if they applied), along with other students with the highest GPA. While students were aware of the nature of the job (summarizing articles), they were unaware of the selection criteria or the total number of slots available in any given week. The selected students joined via a Google link at a pre-determined time to complete the job. While some students completed their task online and submitted their summaries to the team, others submitted handwritten summaries via WhatsApp. These summaries were anonymously graded by the research team on a scale from 1-10.

2.3 Data

We collected four datasets. First, a baseline survey with 496 students captured detailed information about job-seekers' demographics (gender, GPA score, social norms), along with a list of their friends and the strength of their connection. This allowed us to map out their social network.⁵ Second, we conducted weekly midline surveys with our baseline sample to understand whether they had heard about our job opportunity, as well as 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. 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 within the network. Fourth, we tracked who was hired each week as well as several measures of their performance on the job (such as time taken, quality of submission, etc.). Note that both the data on applications and hires

⁵The survey asked respondents to list all the friends in their class who they talk to about jobs.

included information about applicants and hires regardless of whether they were in our baseline sample or not.⁶ This enabled us to capture a comprehensive spread of information within the batch as opposed to just within our sample.

These datasets allow us to generate three key outcomes of interest: whether a student (i) heard about the job; (ii) applied to the job; and (iii) was hired for the job. We also examine the average GPA of the students in each of these three groups.⁷ Lastly, it enables us to measure various indicators of job performance, including attendance, task completion, completion time, and submission quality.

3 Conceptual Framework

Before discussing the results, we present a simple and stylized conceptual framework to help guide our empirical analysis and interpret the results. The aim of this framework is to highlight two key tradeoffs faced by individuals in deciding whether to share job information with peers or not: first, competing with their peers for the job; and second, the utility gains derived from sharing this job information with peers. This provides us with a parsimonious way to interpret the 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 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 their friend j . The utility of individual i is

⁶In practice, this dataset included 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.

⁷Our 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.

given by:

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

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 individual's 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). 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. Furthermore, $\Pr(X_i, X_j, \mathbb{1}_i\{\text{Share}\})$ is defined as follows:

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

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.⁸

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.

Decision to share information: Given this setup, an individual i will share a job with their 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:

⁸There 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.

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

i.e., the utility benefits of sharing the job outweigh the costs of competing for it.

4 Empirical Analysis

Given the conceptual framework outlined above, we will now explore the effects of our intervention on information sharing. In Section 4.1, we start by investigating whether information circulated within the social network and how the competitive nature of the job affected its dissemination. Next, in Section 4.2, we delve into the role of strategic disincentives in information sharing. Finally, in Section 4.3, we analyze how information sharing affected the quality of applicants and hires.

4.1 Does Information Flow?

There were two ways for students in our sample to hear about these jobs. Entry-points heard about the job directly from us (by design), while other students could only hear about the job from their entry-point peers (or from someone connected to an entry-point). Since our primary goal is to understand the flow of job information and how it changes with the rival nature of the job information, we restrict our analysis to the non-entry point students. We begin by estimating the following regression:

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

where Y_{ibt} takes the value 1 if an individual i in batch b in week t heard 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) was 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, while the latter allows for correlations within individual across weeks. From our conceptual framework (Equation 3) we anticipate that individuals would 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 heard about (Column 1) or applied to a job (Column 2) increased by 5.3 p.p. and 4.7 p.p. respectively, if the job we advertised in their batch was non-rival relative to when it is rival. This represents a 30% increase. Furthermore, we can test whether the probability that a student heard about a job increased when they were directly connected to an entry-point who received this information from us, and how this varied with the rival nature of the job information. Specifically, we estimate the following regression specification:

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

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. (Non)Rival $_{bt}$ takes the value 1 if the job we shared in batch b was (non)rival and 0 otherwise. X_i controls for the total number of friends that individual i reported.

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

Taken together, these results highlight two essential aspects of how the competitive nature of job information affects its dissemination within social networks. First, we illustrate the role of strategic disincentives in sharing labor market opportunities. Jobs classified as "rival" were less likely to be shared within a group. That said, an individual's chances of hearing about a job significantly improved when their friends got to know about it – both when the job was rival or non-rival.

4.2 Who Shares Information, and With Whom?

Next, we examine whether these strategic decisions to share information were influenced by additional factors that could exacerbate or mitigate perceived competition

for a job. First, we examine whether individuals shared job information less with their higher-ability peers when the job was rival versus non-rival. Second, we explore whether a (self-reported) measure of the closeness of their friendship was able to overcome job-seekers' tendency to withhold rival job information. Finally, we investigate information sharing between same-gender friends, building on a literature on homophily. 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 conceptual framework (Equation 3), $\partial \Delta U / \partial X_j = -U(w) \partial \lambda / \partial X_j < 0$ 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 were less likely to hear about the job from a lower ability entry-point peer when the job information was rival (as opposed to non-rival). Using a student's GPA score as measure of ability, 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 higher ability) relative to j . We then estimate the following specification for the pair ij in a week t :

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

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 . $(\text{Non})\text{Rival}_{jt}$ 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 $(\text{Non})\text{Rival}_{bt} \times T_{jt}$. 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 , the binary variable of relative ability, and 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 the 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

conceptual 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 increased the probability of hearing about a job in both rival and non-rival batches i.e., across all columns in the table, $\hat{\beta}_{1A} > 0$ and $\hat{\beta}_{2A} > 0$. However, *who* received this job information varies widely based on whether the job is rival or not. A higher ability student was 7.5 p.p. (39.2%) *less* likely to hear about a rival job when their lower ability friend heard about it. On the other hand, they were 8.5 p.p. (51.5%) *more* likely to hear about it from their lower ability friend when the job was non-rival. We can comfortably reject the null hypothesis that $\hat{\beta}_{1B} = \hat{\beta}_{2B}$ (p-value: 0.02). This implies that the probability that a high-ability individual heard about a job from her low-ability friend was indeed different based on whether the job was 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 their close friends. More formally through the lens of our conceptual framework (Equation 3), $\partial\Delta\mathcal{U}/\partial\eta_{ij} = \theta(w) \geq 0$.⁹

In our baseline survey, we asked respondents to tell us for each friend, on a scale of 1 (Not Close) to 5 (Very Close), how frequently they talked to each other about employment and jobs. We then classified each pair as “close” if the respondent rated the frequency of interactions to be 4 or higher.¹⁰ Similar to Equation (6), we estimate the following specification and report the results in Column (2) of Table 2:

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

Turning to the results, we see that individuals were 10.8 p.p. and 9.3 p.p. more likely to share information on jobs with their close connections when the jobs were rival

⁹This 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.

¹⁰Our results are robust to alternate cutoffs (for example, 3 or 5) as reported in Table A3.

and non-rival respectively. That is, individuals were more likely to share jobs with their close connections regardless of the competition for the job ($\hat{\beta}_{1B} = \hat{\beta}_{2B}$, p-value: 0.76). This implies that having a strong connection with someone could potentially counteract the tendency to share less due to competition.

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 :

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

The results are reported in Column (3) of Table 2. We find that individuals were 5.3 p.p. (5.7 p.p.) less (more) likely to share the job information with their same gender friends when the job was 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 less likely to share this rival job 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 were less (more) likely to receive information on a job when it was rival (non-rival) by their relatively lower ability peers. Conversely, the strength of a friendship could mitigate these competitive effects of information sharing: individuals were more likely to share job information with their closest friends even if they had to compete with them for it. Lastly, homophily can induce more information sharing only when the information shared is non-rival.

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.

4.3 Impact on Application Pool, Hiring, and Job Performance

Having established that strategic disincentives can systematically 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 applied for a job conditional on hearing about it. This implies that the pool of applicants we received was directly linked to who heard about the job. We therefore first focus on investigating how the composition of applicants and hires changed when the job was rival or not. For measuring ability, we focus on a students' GPA, which is an observable characteristic that employers routinely use to make hiring decisions.¹¹ Lastly, we use several measures to examine how this impacts students' performance on the job.

Pool of Applicants and Hires: 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 2, we see that the distribution is shifted to the right for non-rival jobs

¹¹To ease interpretation, we standardize GPA scores to have mean 0 and standard deviation 1.

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 was 0.08σ higher when the job was non-rival relative to rival. Considering the substantial conversion rate from learning about a job to applying for it, this translated into a 0.13σ increase in the quality of the applicant pool as well (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 2 shows a similar rightward shift in the ability distribution among hires when the job was non-rival.¹² More formally, in Column (3) of Table 3, we observe that the ability of hires was 0.38σ higher when the job was non-rival relative to rival.

Performance on the Job: We use several measures of job performance and re-estimate Equation (4) to examine differences between the job performance of students hired through rival and non-rival batches. First, we create binary variables that take the value 1 if a student was either: (i) present for the task; (ii) completed the task; and (iii) completed the task within the stipulated time (45 mins), and 0 otherwise. While 83% of students who were hired in rival batches showed up for the job (Table 4, column 1), students hired from non-rival batches were 9.6 p.p. (11.5%) more likely to be present. Students from non-rival batches were also 15.3 p.p. (21%) more likely to finish the task (Table 4, column 2), and 20.4 p.p. (32%) more likely to finish it within the stipulated time. Lastly, we use two other measures of job performance, namely the time taken (in minutes) to finish the job, and the quality of the submission, as measured by the score students received on their submission (1-10). Students from non-rival batches finished their jobs 3.6 minutes (9.1%) faster (column 4), and their scores were 22% higher as well (column 5).

Taken together, these results confirm that the strategic disincentives in information sharing can meaningfully impact the quality of applications that a firm receives, as well as the hires that the firm can make, which ultimately impacts the quality of the work being performed.

4.4 Can Offering Higher Wages Help?

In previous sections, we show that job-seekers share less information with their peers in a competitive setting, and this has consequences for labor market hiring. Specifi-

¹²Having 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.

cally, 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 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 would suggest that higher wages should attract higher quality candidates, these models do not consider the dynamics that come into play within social networks. In particular, while increasing the wage makes a job more appealing, this could elicit two distinct responses from job-seekers 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 fellow applicant has increased. Second, there is the utility channel: a higher paying job is more attractive to share because of the warm glow from sharing information about a better job with a friend. Captured more formally (Equation 3), $\partial \Delta \mathcal{U} / \partial w = \eta_{ij} \theta'(w) - \lambda(X_i, X_j) U'(w)$. Therefore, whether $\partial \Delta \mathcal{U} / \partial w$ is greater than or less than 0 depends on how strongly a change in wages impacts the competition and utility channels. It is entirely possible that job-seekers would share less if the competition channel outweighed the altruism channel, 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, non-rival”.

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

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

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 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 of a rival job (the status-quo), 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 shutting down competition in information sharing, for example by using referral channels (β_2).

Table 5 shows that doubling the wage (for rival jobs) did not significantly increase the probability that individuals heard about it (Column 1) or applied to 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, it is possible that shutting down competition by making the high-wage job non-rival could significantly increase its information-sharing. We find that the probability that a student heard and applied to this high-wage job when it was non-rival increased by 8.2 p.p (Column 1) and 7.9 p.p (Column 2) respectively.¹³ We can comfortably reject the null hypothesis that $\beta_2 + \beta_3 = 0$ (p-val ≤ 0.02).

Lastly, we investigate how the quality of applicants and hires was affected in Table 6. We find that doubling the wage improved the quality of applicants and hires by around 0.1σ and 0.08σ respectively (Columns 2 and 3), and their final evaluation score (ranging from 1-10) on the task by almost 1 point or 14.5% (Column 4).¹⁴ An advantage of our setting is that we can also 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. Specifically, the quality of hires (applicants) improved by 0.35σ (0.04σ) when the high-wage job was non-rival relative to rival and these candidates had a 1.3 points (19.1%) higher evaluation score on the task as well.¹⁵ Lastly, to gauge the importance of the competition channel, a simple back-of-the envelope calculation comparing β_1 and β_2 suggests that to get the same increase in ability among the pool of hires (applicants), a firm would have to increase wages by 5.8 (3.3 times).¹⁶ Similarly, firms would increase wages by 2.7 times to maintain the same quality of job performance (Column 4).¹⁷ In other words, these results underscore the importance of strategic disincentives in informa-

¹³We recover these estimates by adding $\beta_2 + \beta_3$.

¹⁴The difference between this positive impact and the absence of an effect on the likelihood of hearing about a job can be clarified by the observation that high-ability students were more inclined to apply when they heard about a high-wage job. Consequently, they were more likely to be hired compared to their lower-ability peers.

¹⁵We recover these estimates by adding $\beta_2 + \beta_3$.

¹⁶To see this, note from Column (3) that the average quality of hires is 0.08σ higher when wages are doubled, and 0.39σ 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 $1+0.391/0.081$ ($1+\beta_2/\beta_1$) i.e., 5.8 times higher.

¹⁷The results for other job performance indicators are reported in Table A4.

tion sharing. Specifically, they suggest that firms would have to offer substantially higher wages to attract the same pool of applicants they would have attracted in the absence of these disincentives.

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 help 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 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 of Development Interventions," *Journal of Political Economy Microeconomics*, 2023, 1 (1), 162–224.
- Banerjee, Abhijit and Gaurav Chiplunkar**, "How Important Are Matching Frictions in the Labor Market? Experimental & Non-Experimental Evidence from a Large Indian Firm," *Working paper*, 2023.
- , **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.
- , **Niall Keleher, and Jeremy Magruder**, "Do job networks disadvantage women? Evidence from a recruitment experiment in Malawi," *Journal of Labor Economics*, 2018, 36 (1), 121–157.
- 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, Kate Orkin, Rob Garlick, Niharika Singh, Rachel Heath, and Alison Andrews**, "Barriers to Search and Hiring in Urban Labour Markets," Technical Report, Vox Dev Literature 2024.
- , **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.

- 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.
- 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
<i>Panel A:</i>		
Non-Rival _{bt}	0.053 (0.020)**	0.047 (0.022)**
Control Mean	0.18	0.14
Observations	2535	2535
<i>Panel B:</i>		
Non-Rival _{bt}	0.030 (0.017)*	0.020 (0.017)
Non-Rival _{bt} × T _{it}	0.255 (0.057)***	0.230 (0.054)***
Rival _{bt} × T _{it}	0.241 (0.048)***	0.188 (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 (0.036)***	0.090 (0.022)***	0.190 (0.048)***
Rival _{jt} \times X (β_{1B})	-0.075 (0.038)*	0.108 (0.029)***	-0.053 (0.045)
Non-Rival _{jt} (β_{2A})	0.165 (0.033)***	0.143 (0.033)***	0.143 (0.041)***
Non-Rival _{jt} \times X (β_{2B})	0.085 (0.061)	0.093 (0.045)**	0.057 (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 $1(\text{Ability}_i > \text{Ability}_j)$, 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 (0.037)**	0.130 (0.045)***	0.381 (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 non-entry point group in week t and respondents who were entry-points when the job was rival 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: Performance of Students on the Job

	(1) Present	(2) Complete	(3) Complete Overtime	(4) Time Taken	(5) Eval. Score
Non-Rival _{bt}	0.096 (0.038)**	0.153 (0.041)***	0.204 (0.058)***	-3.576 (1.842)*	1.619 (0.395)***
Control Mean	0.83	0.74	0.64	39.21	7.26
Observations	304	304	304	304	304

Notes: This table shows how the performance of hired students changes when a job is rival or not. The sample is restricted to hired respondents who were assigned to non-entry point group in week t and respondents who were entry-points when the job was rival in week t . The dependent variables in Columns (1)-(3) are binary variables that take the value 1 if a student showed up for a task, submitted, and submitted within the stipulated time limit (45 mins), and 0 otherwise. The dependent variables in Columns (4) and (5) are the time taken (in minutes) to complete the task and the quality of submission, which is a score from 1-10. Non-Rival_{bt} takes the value 1 if the batch b was assigned to the non-rival 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: Heard About and Applied to a Job (High-Wage)

	(1) Heard	(2) Applied
High-wage _{bt} (β_1)	-0.009 (0.024)	-0.007 (0.025)
Non-Rival _{bt} (β_2)	0.022 (0.024)	0.013 (0.027)
High-wage _{bt} \times Non-Rival _{bt} (β_3)	0.060 (0.039)	0.066 (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 non-entry point 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

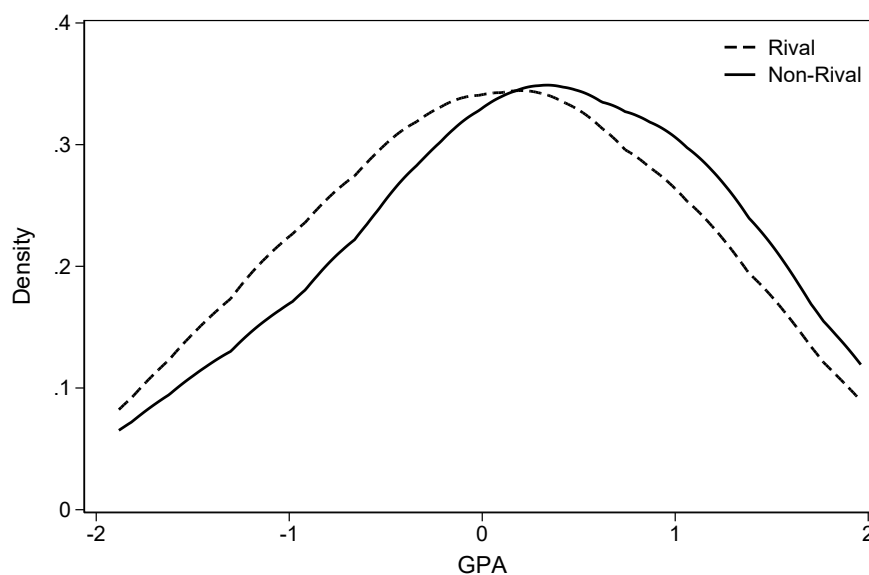
Table 6: Ability of Students (High-Wage)

	Std. GPA Score			Eval. Score
	(1) Heard	(2) Applied	(3) Hired	(4) Hired
High-wage _{bt}	-0.036 (0.051)	0.100 (0.028)***	0.081 (0.047)*	0.989 (0.320)***
Non-Rival _{bt}	0.117 (0.080)	0.232 (0.045)***	0.391 (0.035)***	1.723 (0.568)***
Non-Rival _{bt} × High-wage _{bt}	-0.059 (0.103)	-0.195 (0.028)***	-0.037 (0.144)	-0.407 (0.774)
$\beta_1 = \beta_2$	0.06	0.05	0.00	0.18
$\beta_2 + \beta_3 = 0$	0.41	0.43	0.01	0.01
Control Mean	0.03	0.00	0.04	6.80
Observations	688	462	304	304

Notes: This table shows how the ability of students who hear (Column 1), apply (Column 2), are hired (Column 3), and their evaluation score (Column 4) changes when a job is rival/high-wage or not. The sample is restricted to respondents assigned to non-entry point group in week t and respondents who were entry points when the job was rival 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 Columns (3) and (4), the sample is restricted to students who were hired. The dependent variable is the respondent's standardized GPA score in Columns (1)-(3) and the evaluation score (from 1-10) on the quality of submission in Column (4). 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

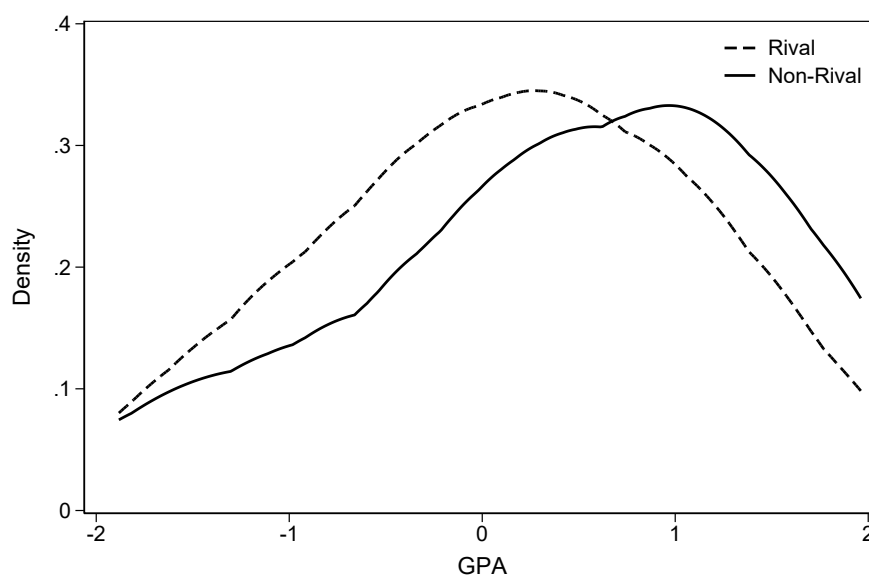
Figures

Figure 1: Ability Distribution of Applicants and Hires



[b]

Figure 2: Job Applicants



[b]

Figure 3: Job Hires

Notes: This figure plots the densities of students' standardized GPAs for those who applied for jobs in Panel (a) and those who were hired in Panel (b). In both panels, the distribution for rival jobs is shown by the dashed line and the distribution for non-rival jobs is shown by the solid line.

Appendix Tables

Table A1: Balance Across Entry-Point and Non-Entry Point 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
Religion: 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-points and non-entry points across 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
<i>Panel A: Males</i>			
Rival _{jt}	0.229 (0.047)***	0.132 (0.038)***	0.238 (0.056)***
Rival _{jt} × X	-0.142 (0.047)***	0.062 (0.047)	-0.107 (0.059)*
Non-Rival _{jt}	0.090 (0.040)**	0.111 (0.038)***	0.132 (0.061)**
Non-Rival _{jt} × X	0.117 (0.075)	0.102 (0.057)*	0.036 (0.068)
$\beta_{1B} = \beta_{2B}$	0.00	0.51	0.12
Observations	1111	1491	1491
<i>Panel B: Females</i>			
Rival _{jt}	0.166 (0.040)***	0.057 (0.022)**	0.119 (0.075)
Rival _{jt} × X	-0.033 (0.059)	0.142 (0.041)***	0.019 (0.077)
Non-Rival _{jt}	0.213 (0.046)***	0.175 (0.043)***	0.156 (0.063)**
Non-Rival _{jt} × X	0.073 (0.070)	0.081 (0.061)	0.073 (0.081)
$\beta_{1B} = \beta_{2B}$	0.22	0.39	0.59
Observations	1670	1979	1979

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 $1(\text{Ability}_i > \text{Ability}_j)$, 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 A3: Robustness by Definition of Close Friendship

	Threshold for Strength of Friendship		
	(1) = 3	(2) = 4	(3) = 5
Rival _{jt}	0.089 (0.035)**	0.090 (0.022)***	0.134 (0.022)***
Rival _{jt} × X	0.070 (0.039)*	0.108 (0.029)***	0.074 (0.042)*
Non-Rival _{jt}	0.125 (0.046)***	0.143 (0.033)***	0.171 (0.027)***
Non-Rival _{jt} × X	0.080 (0.053)	0.093 (0.045)**	0.101 (0.064)
$\beta_{1B} = \beta_{2B}$	0.88	0.76	0.69
Observations	3470	3470	3470

Notes: The outcome variable in the above table is a binary variable that takes the value 1 if a student hears about the job and 0 otherwise. Individuals were asked on a scale of 1 (not a lot) to 5 (a lot) how frequently they talked to their friend about other employment and job opportunities. Columns (1)-(3) then define a binary variable X that takes the value 1 if a respondent's frequency of interactions had a value of at least 3, 4, or 5 respectively. 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 A4: Job Performance Measures and High-Wage Jobs

	(1) Present	(2) Submitted	(3) Submitted Overtime	(4) Time Taken	(5) Eval. Score
High-wage _{bt}	0.122 (0.026)***	0.109 (0.036)***	0.120 (0.045)***	-6.509 (1.348)***	0.989 (0.320)***
Non-Rival _{bt}	0.142 (0.062)**	0.168 (0.062)***	0.178 (0.099)*	-4.388 (3.389)	1.723 (0.568)***
Non-Rival _{bt} × High-wage _{bt}	-0.109 (0.073)	-0.050 (0.083)	0.020 (0.120)	2.913 (4.230)	-0.407 (0.774)
$\beta_1 = \beta_2$	0.74	0.33	0.53	0.52	0.18
$\beta_2 + \beta_3 = 0$	0.44	0.03	0.00	0.46	0.01
Control Mean	0.77	0.69	0.59	41.98	6.80
Observations	304	304	304	304	304

Notes: The dependent variables in Columns (1)-(3) are binary variables that take the value 1 if a student showed up for a task, submitted, and submitted within the stipulated time limit (45 mins), and 0 otherwise. The dependent variables in Columns (4) and (5) are the time taken (in minutes) to complete the task and the quality of submission, which is a score from 1-10. Non-Rival_{bt} and High-wage_{bt} take the value 1 if the batch b was assigned to the non-rival or high-wage treatment in week t respectively, and 0 otherwise. All regressions include batch and week fixed effects. Standard errors are clustered at the week-batch and individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$