Complementarities in Labor Supply: How Joint Commuting Shapes Work Decisions

Aletheia Donald and Florian Grosset-Touba*

September 2024

Please click here for the most recent version

Abstract

Commuting is notoriously challenging in lower-income country cities, limiting firms' ability to attract workers from different neighborhoods and reducing the economic benefits of urban agglomeration. This paper examines whether decreasing the non-financial cost of commuting—by commuting with peers rather than alone—reduces this friction and increases labor supply. We conduct two field experiments in urban Côte d'Ivoire. In the first experiment, job seekers are 16pp more likely to accept a formal full-time factory job if their peers also receive a job offer, and 15pp more likely to remain in that job four months later—but only if they will be employed in the same shift (rather than different shifts). These effects are driven by workers with long commute times, who can commute to work together. Consistent with this channel, workers' own attendance and turnover are predicted by the attendance and quits of co-commuting peers. In a second field experiment with a different firm, we again randomize whether a worker's peers are offered a job and whether they could commute together. We also randomize job location—inducing exogenous variation in commute time. We replicate the finding of complementarities in labor supply, but only in the case of long assigned commute times. These findings indicate that commuting together serves as an important job amenity with large impacts on labor supply. Our results provide a novel explanation for key features of urban labor markets, including firms' widespread use of referrals for hiring and persistent gaps in employment across social groups.

*Donald: World Bank (adonald@worldbank.org); Grosset-Touba: CREST & ENSAE Paris (Corresponding author: florian.grosset@ensae.fr). Florian is indebted to his advisors: Eric Verhoogen, Kiki Pop-Eleches, Supreet Kaur, and Wolfram Schlenker. This paper greatly benefited from comments by Livia Alfonsi, Michael Best, Sandy Black, Laura Boudreau, Hannah Farkas, Rob Garlick, Louise Guillouet, Gabriel Kreindler, Suresh Naidu, Anna Papp, Tommaso Porzio, Jeff Shrader, Jack Willis, and Krzysztof Zaremba, and multiple seminar and conference participants. We gratefully acknowledge financial support from the World Bank's Umbrella Facility for Gender Equality; the Gender, Growth, and Labour Markets in Low Income Countries Programme at IZA/FCDO; and the Center for Development Economics and Policy at Columbia University. We thank Innovations for Policy Action (IPA), especially Augustin Kouadio, Bechir Ouedraogo, and Romuald Kouadio for assistance with implementation. We are grateful to the leadership and staff at our partner firms, who have been instrumental in the implementation success of this project. We thank Anika Fergusson for excellent research assistance. This project is a product of the World Bank's Africa Gender Innovation Lab. The findings, interpretations, and conclusions expressed in this paper 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. Research approved by Columbia University IRB and IPA IRB. AEA RCT Registry ID: AEARCTR-0006154. All errors and omissions are our own.

1 Introduction

Commuting in cities of lower-income countries is notoriously challenging. It is often long and unsafe due to inadequate transportation infrastructure, congestion, and poor urban planning (Akbar et al. 2023; Aksoy et al. 2023). These commuting costs constrain economic development, by reducing the economic benefits from urbanization: labor markets become fragmented within cities, as people prefer to work close to their homes, and firms struggle to hire and retain workers who do not live nearby.

This paper empirically examines the possibility that joint commute induces complementarities in labor supply. Does a person's desire to work at a given firm depend on her ability to commute there together with others from her social network?

If present, such complementarities have potentially broad implications for firms, labor markets, and urban development. They suggest that the organization and social composition of the workforce is an important job amenity. This could explain stylized facts about firm behavior——such as why firms across contexts hire by soliciting referrals (Afridi and Dhillon 2022; Topa 2019, 2011) and why they often establish close to each other within cities. It would also imply that policy-makers' interventions to help coordinate hires from the same social network across firms located in the same area could increase labor supply and facilitate hiring by each of these firms. Such complementarities also have potential aggregate consequences: for instance, they could generate multiple equilibria where social groups have either high or low labor supply to a given area in the city—with the potential to generate persistent differences in patterns of labor supply, employment, and earnings across groups.

We design our study to establish the empirical relevance of these ideas. We consider a set of job seekers (which we refer to as "focal") deciding whether to accept a job offer, and elicit their networks at baseline. We randomly vary whether individuals in these networks also receive a job offer at the same firm. We further vary whether the focal job seekers and their network members would be working at the same time and location at the firm—inducing random variation in their ability to commute together—and their commuting time from home to work. We test whether this changes the focal job seekers' labor supply, primarily their job take-up decisions. We also estimate effects on retention and productivity.

We implement this design with a set of two field experiments in Abidjan, Côte d'Ivoire. The first experiment is with a large multinational company operating a cashew-processing plant. We offer long-term factory jobs to 163 job seekers and elicit referrals for jobs at the same plant. Most of the individuals they refer are friends and live in the same neighborhood, so they could potentially commute together to work. All jobs offered are at the same plant, but work is organized into two non-overlapping shifts. We randomly assign the focal job seekers to one of three groups: "Control" (no network members receive an offer), "Same Shift" (up to three network members receive an offer to work in the same shift as their focal individual), and "Different Shift" (up to three network members receive an offer to work in

a different shift than the focal individual).

We find evidence of strong labor supply complementarities, driven by the ability to interact with network members. Focal individuals assigned to the Same Shift group are 16 percentage points more likely to take up the job than those assigned to Control (p-value: 0.093; a 63% increase from the take-up rate in Control). This effect persists over time: the share of the Same Shift group focal individuals working at the factory remains 10-18 percentage points higher than the Control group, whether we consider working at least one day, one month, or four months. By contrast, there is virtually no difference in average take-up and retention between the Control and Different Shift arms. This indicates that having network members employed, even at the same firm and plant, is insufficient to trigger an employment response. Direct interactions between the focal job seekers and their network members drive the effects and reflect the existence of complementarities in labor supply.

Though the value of interacting with network members could be due to a number of reasons, our factory data provides two types of evidence suggesting that the ability to commute together is an important driver of the observed complementarities. First, the effects of hiring network members to work the same shift are larger for focal job seekers with longer commutes. Second, non-experiment factory workers coordinate their absences and quit decisions when they commute together. Since complementarities in labor supply driven by interactions during commute should influence decisions to come to work beyond job take-up, we surveyed 760 non-experiment workers at the factory, elicited with whom they commute, and matched it to the firms' attendance records. We find that workers' own attendance and turnover are strongly predicted by the attendance and quits of their co-commuting peers. We use a rich set of fixed effects to rule out common shocks and identify the coordination of labor supply decisions as driver of these patterns.[1] These results indicate that complementarities in labor supply determine a wide range of work decisions and suggest that interactions during commute play an important role.

We design the second experiment to identify the role of interactions during commute as a driver of complementarities in labor supply, while also directly testing for alternative candidate mechanisms such as interactions at work. We partner with a marketing firm hiring for short-term street canvassing jobs. The jobs take place at two different worksites, each located in a different borough of Abidjan. We offer the jobs to 873 focal job seekers, and randomly assign them to one of three groups: "Control" (no network members receive an offer), "Same Site" (up to three network members receive an offer to work at the same site as their focal individual), and "Different Site" (up to three network members receive an offer to work at the other site than their focal individual). We also randomize each focal job seeker's assignment to one of the two potential worksites, inducing random variation in commuting

[1] We control for day-by-neighborhood and day-by-production unit fixed effects to account for potential common shocks linked to the workers' neighborhood (local disruption of transports, local festival, etc.) and production unit (such as production slowdowns with mandated absences).

time. This design allows us to test for causal heterogeneous treatment effects by commuting time.

In this second experiment, we identify interactions during commute as the primary driver of the complementarities. First, we replicate the main findings from the factory experiment in the different settings of sales jobs. We find that focal job seekers assigned to the Same Site treatment are 7 percentage points more likely to take up the job than those assigned to Control (p-value: 0.092; a 20% increase from the take-up rate in Control). Second, if interactions during commute are a driver of the complementarities, the treatment effects should be larger for the focal job seekers (i) who can commute with their network members and (ii) who have longer commutes. Our design exogenously varies these two dimensions across the focal job seekers. We find empirical support for this prediction: the Same Site treatment effects are heterogeneous by commuting time (p-value: 0.077), while the Different Site treatment effects are not. We rule out alternative explanations for this pattern, including the potential higher value of joint learning when working outside one's neighborhood. Third, we find that the Same Site treatment effects are concentrated among job seekers with longer commutes: there are virtually no treatment effects among focal job seekers with below-median commuting time. This indicates that complementarities in labor supply are primarily driven by interactions during commute, rather than by interactions while on the job.

Besides simply filling vacancies, firms are concerned with the productivity of their workers. Ex ante, the productivity consequences of hiring through networks are ambiguous. There can be direct effects, with ambiguous net impacts: workers who know each other can learn faster and motivate each other, but they can also chat and shirk more together. There can also be a selection effect: the marginal job seekers, who only take up the offered job when their network members also receive an offer, may have different productivity than the inframarginal ones. We therefore collect detailed productivity data for the sales jobs from two independent sources.[2] We find that workers in the Same Site group do not have lower productivity than workers in Control. This holds whether we consider workers with commuting times below the median (no treatment effects on job take-up; isolates the direct effects of working with network members on productivity) or above the median (combining the direct and selection effects on productivity). Further, the focal and referred workers have similar productivity. Overall, hiring through networks did not reduce the average productivity of the sales workers.

Building on these results, we next estimate the amenity value of being able to interact with network members during commute. We do so by combining the estimated treatment effects with a natural wage increase that happened when piloting the experiment. This calibration is not trivial since the treatment is implemented based on a fixed network, while changing the offered wage can affect both whether a job seeker provides a referral and takes

[2] At the time of the experiment, no productivity data was collected by the factory.

up the job. We build a model of referral and job take-up decisions to guide the calibration. We estimate that increasing job take-up by eight percentage points, corresponding to the effects of the Same Site treatment, would require increasing the offered wage by 13 to 20%.

Our study advances the literature on how social interactions influence work decisions. A strand of this literature considers how individuals choose the timing of their work to be able to spend leisure time with their spouse (Georges-Kot et al. 2022; Goux et al. 2014; Moghadam et al. 2023; Hamermesh 2000; Hallberg 2003) and relatives (Georges-Kot et al. 2017). We complement these studies by jointly examining different types of interactions that could give rise to complementarities, besides joint leisure outside work. These include interactions on the job and during commute. We primarily study effects on job take-up decisions, a dimension not yet explored in this literature with potentially large welfare consequences. We also consider interactions among a different social group than relatives: referrals for a job at the same firm.

A rich body of work studies interactions among workers within a firm. We consider individuals who would work at the same hierarchical level, and so remain distinct from studies on interactions between workers and managers (Ashraf and Bandiera 2018; Bandiera et al. 2009). The literature on interactions among workers of the same level has primarily focused on peer effects in productivity. Studies in this literature have found effects of different sizes and even directions, likely due to differences in incentives and the type of links among workers in the study samples (Mas and Moretti 2009; Bandiera et al. 2010; Herbst and Mas 2015; Cornelissen et al. 2017; Park 2019; Afridi et al. 2020a; Afridi et al. 2020b). Theoretically, workers could motivate each other and learn better together, but they could also chat and distract each other. We advance this literature by expanding the dimensions of work decisions considered, beyond productivity. Our sales experiment provides an additional data point to the literature: we find no negative productivity effects from working with the referred individuals. We also find that factory co-workers who commute together coordinate their absences and quit decisions. This outlines a potential trade-off for firms in their hiring and management decisions: even if peers at work have the potential to increase the level of labor supply to the firm, they can also increase its volatility.

Our results can help explain firm hiring behavior, particularly the extensive use of referrals. A large existing literature has studied this topic through the lens of information asymmetries: hiring through referrals can reduce search, screening, and monitoring costs.[3] From a worker's perspective, recent work has demonstrated that they value providing referrals as a way of being involved in their firm's decisions (Friebel et al. 2023) and because of expected future reciprocity (Witte 2021). We advance this body of work by showing that individuals from the same social group value being able to commute together to work and that hiring network members can accordingly represent an important non-wage job amenity. Firms can therefore use referral-based hiring to increase the level of labor supply they face

[3] Afridi and Dhillon (2022) and Topa (2019) and Beaman (2016) provide recent reviews of this literature.

without increasing the offered wage.[4] Our findings on the role of joint commute for labor supply are thereby consistent with US-based evidence indicating that commuting time and monetary cost represent one of the most important non-wage job amenities influencing work decisions (Dube et al. 2022; Mas and Pallais 2017).

Our focus on the role of joint commute for labor supply complementarities directly builds on the rapidly growing literature on commuting in lower-income countries. A robust set of studies has established the existence of major commuting costs in these settings, with broad economic consequences (Abebe et al. 2020; Franklin 2018; Borker 2021; Field and Vyborny 2022). We focus on their implications for labor supply and demonstrate how, given an existing transportation system, being able to commute together can stimulate employment.

Other potential avenues to reduce commuting costs include upgrading transportation infrastructure and work-from-home arrangements. However, while transportation infrastructure projects can have positive economic impacts in developing country cities (Tsivanidis 2023; Kreindler et al. 2023; Kreindler 2023; Velásquez 2023; Zárate 2022), they remain very costly and slow to implement. Likewise, despite their benefits in terms of reduced commute, work-from-home arrangements suffer from important downsides: they are not feasible for many types of occupations, the required digital infrastructure is often inadequate in lower-income settings, and promoting work-from-home may be undesirable for other social reasons (e.g., isolation of women in the home).

More broadly, our work builds on and advances the literature on the aggregate economic consequences of social interactions. In two studies, Glaeser et al. (1996) and Glaeser et al. (2003) formalize the idea that complementarities can give rise to multiple equilibria. They empirically apply it to the case of criminal activities but suggest potential applications to other domains, including labor. Alesina et al. (2005) build on this idea and discuss the possibility that the presence of complementarities in labor supply could have contributed to the dramatic reduction in hours worked starting in the early 1970s in Europe but not in the United States. However, they only offer anecdotal evidence that the complementarities required for this explanation exist. In a developing country urban setting, our study provides causal evidence of the existence of strong complementarities in labor supply, and therefore supports the possibility of multiple equilibria in employment. Specifically, our findings indicate that individuals may only supply labor to firms in a given area (such as an industrial district) if they have network members working there with whom they could commute. If they do not, they may refuse job opportunities, thus perpetuating low levels of formal employment in the network. Since network members often share similar characteristics, it implies that individuals from different genders, ethnic groups, or neighborhoods can

[4]Recent work has found negative impacts of job search assistance programs on the employment of the beneficiaries' network members: due to changes in information sharing and job search assistance among network members (Caria et al. 2022) and to a re-allocation of work among household members in a setting with conservative gender norms (Afridi et al. 2022). A key difference with our setting is that we offer jobs to several network members at the same firm, allowing them to commute together.

have persistently different patterns of labor supply, including formal wage employment.

In what follows, we present the context of the study (section 2). We then proceed with the empirical strategy and results for the factory experiment (section 3) and for the sales experiment (section 4). Section 5 presents the model and calibration used to estimate the amenity value, while section 6 summarizes and discusses the project.

2 Context: Labor Market and Commute in Abidjan

We work in Abidjan, the economic capital of Côte d'Ivoire. Abidjan has a dense urban labor market, characterized by long commutes. Firms in this setting often report difficulties in hiring and retaining blue-collar workers at prevailing wages (World Bank 2018, 2020). This is also the case for our partner companies. For instance, at our partner factory, only 25% of the individuals who reported being interested in a production job and received a job offer ever came to work, and 41% of those who started quit within 4 months. These difficulties are not specific to Abidjan but have been documented across lower-income countries (Abebe et al. 2020; Blattman and Dercon 2018; Donovan et al. 2023). By constraining firms' investment and growth, they represent an important barrier to economic development and poverty alleviation (Le Barbanchon et al. 2023; Bloom and Van Reenen 2011).

Commuting in Abidjan is typically long and unsafe due to a combination of poor urban planning, congestion, and a lack of adequate transportation infrastructure. Most opportunities for formal blue-collar jobs are concentrated in industrial areas located close to the commercial harbor and areas with high housing prices, but far away from the denser and poorer residential areas where most potential blue-collar workers live (World Bank Group 2019). Those who do work typically commute by a combination of walking and public transportation. Most trips with public transport are on informal minibusses, with safety risks from both harassment and road accidents. These characteristics are common across many cities in lower-income countries.[5]

3 Factory Experiment

Our first experiment is with a large multinational company hiring for full-time jobs at a cashew-processing plant in Abidjan. We test for the existence of complementarities in labor supply, isolate the role of direct interactions among job seekers and their network members, and examine the specific role of joint commuting.

[5]For instance, in Morocco, industry executives report unreliable transport services as a source of high employee turnover, and transportation to be a major concern for employees and labor union representatives (Minster et al. 2022). Akbar et al. (2023) estimate that urban travel speed is roughly 50% faster in rich countries than in poor countries, mostly due to different city structures.

This section describes the setting of the experiment (3.1), its design (3.2), and implementation (3.3). It then presents the available data and the associated estimation strategy (3.4), the sample of job seekers offered the job (3.5), and the results on job take-up and retention (3.6).

Complementarities in labor supply driven by joint commuting should influence work decisions beyond take-up and retention. We therefore complement the experimental results with evidence on the coordination of absences and quits among non-experiment factory workers commuting together (3.7). Our second experiment then confirms the role of joint commute as key driver of the complementarities (section 4).

3.1 Factory: Setting

Our partner company for the first experiment is Olam, a large multinational company. It is one of the largest employers in Côte d'Ivoire and has been operating in Abidjan and elsewhere in the country since 1994. We partnered as they expanded their cashew-processing activities in the country and opened a new factory in Anyama, on the outskirts of Abidjan. Situating their facility there, rather than the more remote and traffic-jammed industrial district where most competitors are, allowed them to locate close to the most densely populated borough of Abidjan and have ready access to the large potential blue-collar labor force. A large share of Olam's workers find their jobs through networks.

At the time of our experiment, Olam was hiring for production roles at the factory. These roles include the manual peeling and grading of nuts, an individual activity where each worker performs the same repetitive task, or the packaging of the processed nuts, an activity with a group of workers each doing their task along a production line. About 70-80% of these blue-collar roles are traditionally filled by women. For such jobs, skills are difficult to screen ex ante, and the eligibility criteria to become hired are essentially the legal requirements for formal work.[6] Contracts are open-ended—with the expectation of long-run employment relationships. Workers are paid 4,000 FCFA (∼6 USD) for each day of work, slightly above the national minimum wage.

3.2 Factory: Design

Our study has two main objectives. The first objective is to test for complementarities in labor supply. By definition, there exist complementarities in labor supply if an individual's utility from supplying labor depends on the labor supply of their network members. While we cannot directly observe individuals' utility, we can conduct a strong test for the existence of complementarities in labor supply: whether an individual's labor supply varies when the labor supply of their network members exogenously varies. The second objective is to shed

[6] These include being at least 18 years old and producing an identity document.

light on the drivers of the complementarities. To do so, we will test for the conditions under which the complementarities influence the labor supply decisions of the focal job seeker. Do they require direct interactions between the focal job seeker and their network members? If so, what type of interactions?

In both experiments, we study a set of job seekers. They all receive a job offer, and we study their take-up decisions. We randomly vary whether their network members also receive a job offer at the same firm, and whether they would be working at the same time and location. The specific design of the first experiment leverages an institutional feature of the factory. The factory operates in two non-overlapping shifts—a morning shift and an afternoon shift. These shifts are large, with over 100 workers in each. Workers are always assigned the same co-workers, but whether they work the morning or afternoon shift varies every week. We use this feature and randomly assign the focal job seekers to one of three groups. In the "Control'' group, the network members of the focal job seeker do not receive a job offer. In the "Different Shift'' group, the network members of the focal job seeker receive a job offer but will not work in the same shift as the focal job seeker. In the "Same Shift'' group, the network members of the focal job seeker receive a job offer and they will work in the same shift as the focal job seeker.

The design of our factory experiment allows us to advance towards the study's objectives. Comparing the labor supply decisions of the focal job seekers assigned to the Same Shift group relative to those assigned to Control provides the broadest test for the presence of complementarities. However, this comparison may not yield perfect identification: a confound would exist if focal job seekers change their labor supply in response to the job offer being made to their network members—irrespective of whether it changes their network members' labor supply decisions. This might happen, for instance, if individuals update their beliefs about the firm upon learning about these offers. Comparing the Same Shift group with the Different Shift group alleviates this concern.

Focal job seekers' network members receive a job offer in both the Same Shift and Different Shift groups; the difference is whether they would be able to interact with their network members. This comparison cleanly identifies the existence of complementarities in labor supply, driven by direct interactions. Such interactions could happen at the worksite, during commute, or even possibly outside work. In the factory setting, heterogeneity in treatment effects by commute times and patterns of attendance and quits among co-commuting non-experiment workers shed light on the role of commuting. Our second experiment is then specifically designed to causally identify the role of commuting as a mechanism.

A difference in the labor supply of focal job seekers across the Control and Different Shift groups could come from either of two sources. First, it could indicate that complementarities in labor supply do not need direct interactions at work to exist. For instance, this can be the case since having network members get a job can reduce redistributive pressure on the focal individuals, thereby stimulating their labor supply (Carranza et al. 2022a; Hoff

and Sen 2011). Once network members become employed, they can also change their attitudes towards unemployment or no longer be available to take care of the focal individual's children——which could influence the labor supply of the network members. Second, it can come from the simple effect of providing a job offer to network members, whether they accept or not, as discussed above.

3.3 Factory: Implementation

Across the two experiments, the network members consist of referrals elicited at the hiring stage from each focal job seeker. After the focal job seekers receive information about the job opportunity, it is announced that the firm is looking to hire a large number of workers, and they are asked if they know other individuals suitable and available for the job. The list of individuals referred by each focal job seeker constitutes their network members. To reduce strategic considerations in these referral decisions, it is made clear to the focal job seekers that the referred individuals will be placed in a different pool of applicants with whom they will not be competing and that the firm will not link their job offer and pay to the take-up and performance of those they refer. Because of these features, we can expect the focal job seekers to refer individuals with whom they may have complementarities in labor supply, rather than individuals whose employment provides no or negative value to them. As such, our results will likely be different than if we had made random job offers to a subset of, for instance, the focal individuals' friends or neighbors. Rather, they relate to complementarities in the type of network most relevant for firms and policy-makers considering using referrals to recruit workers or program participants.

The treatment assignment, including which network members will receive a job offer, is announced to the focal job seekers after referrals are elicited. During the treatment announcement, focal job seekers in the "Different Shift" or "Same Shift" groups learn which of their network members are receiving a job offer, whether they would be working in the same shift, and the specific work hours and location. In these two groups, up to three network members per focal job seeker receive a job offer. If the focal job seeker reported three names or less, all reported network members receive an offer. If more than three names were given, three of the reported network members are randomly chosen to receive the offer. The focal job seekers in the "Control" group learn the specific work hours and location.

We implement the two experiments as part of the hiring process of our partner firms. Usually, they hire as follows. When the firms have job openings, they assemble a list of job seekers who are potential candidates for the jobs. This list can come from a database of earlier candidates, hiring scouts, spontaneous applications, ads placed on social media, etc. The firm will then call the job seekers to provide them with information about the specific job openings and elicit their interest. If they are interested, they will invite them to work, usually starting with half a day of onboarding. Our experiments are embedded in this process.

Specifically, we obtain from the firms the list of job seekers who are potential candidates for the job. A short baseline survey is administered to them: they receive some information about the job opportunity, their network is elicited, basic additional socio-economic baseline characteristics are collected, their treatment assignment is announced, and their interest in the job is elicited. We then create a list containing (i) the initial job seekers who are interested and eligible, and (ii) the network members of the job seekers assigned to one of the two non-Control groups. We provide that list to the firm, who will continue the hiring process as usual: call them, provide them with more information, elicit interest, and invite them to work. This implementation process is presented visually in Appendix Figure A1.

To be eligible, the focal job seekers must satisfy the basic work requirements from the partner company[7] and have provided at least one referral—otherwise, there is no network member we could offer a job to, and the treatment cannot be implemented.

3.4 Factory: Data and Estimation

We obtain access to the factory records compiling each day the attendance of each worker. This allows us to measure job take-up by each focal job seeker and their network members. It also enables us to track workers over time at the factory and measure their retention. These records are used by the firm to calculate workers' pay and, therefore, have high accuracy.[8]

We estimate treatment effects on the job take-up and retention decisions of the focal job seekers. Our main specification estimates the Intent-to-Treat effects of being assigned to either the Same Shift or Different Shift group, relative to Control. We implement a simple comparison of means at the job seeker level and estimate:

$$y_i = \alpha^f + \beta_1^f \text{DifferentShift}_i + \beta_2^f \text{SameShift}_i + \gamma^f X_i + \epsilon_i^f, \tag{1}$$

where y_i is the outcome of interest for focal individual i, DifferentShift_i and SameShift_i are indicator variables for the job seeker's assignment into the "Different Shift" or "Same Shift" groups, and X_i a vector of controls. Our primary specification controls for the number of network members receiving a job offer.

Our main outcome of interest is job take-up: a binary measure indicating whether the focal individual ever came to work. Before being invited to come to work, prospective workers have to undergo a short training session implemented by the firms for all their workers. During this training, job seekers can ask the questions they still have about the job characteristics. This reduces the likelihood that workers have any important surprises on their first day of work that would lead them to quit and not come back, although our

[7] Being at least 18 years old and having a valid ID. In addition, the sales jobs require owning a smartphone and knowing how to use apps on it.

[8] At the time of the experiment, the factory was not collecting worker-level production data.

results are robust to defining job take-up as coming at least twice to work. Any potential surprises on the first day of work are not driving our results.

Our secondary outcome is the retention of the job by the focal job seekers. Since observing this secondary outcome is conditional on job take-up, the estimates will capture a mix of direct and selection effects. We still consider them, since they are of first-order relevance for firms.

3.5 Factory: Sample

For the factory experiment, we work with a sample of 163 eligible job seekers, described in Table 1. These are individuals with some experience but limited education: they are 29 years old on average, over half report having a partner, and over a third have a young child under their care. The average job seeker has a middle-school education. Consistent with the gender composition of the factory, 72% of the focal job seekers are women. The sample composition across the treatment arms is balanced over these and other characteristics.[9]

From a pecuniary perspective, the offered factory jobs are desirable relative to the outside option of the focal job seekers. Nearly a third of the focal job seekers in our sample haven't had any income-generating activity in the three months preceding their application. Even for the others, the offered factory jobs pay over twice their baseline average earnings, for similar hours worked.

[9] As expected when testing the equality of means across multiple groups, we do find a limited number of statistically significant differences. Specifically, workers in the Same Shift group are older than those in Control (p-value=0.090), while workers in the Different Shift group are less likely to have children than workers in Control (p-value=0.023), as reported in Appendix Table A1. Our main results are qualitatively similar whether we control for these two variables or not.

Table 1: Sample Description for the Factory Experiment

Variable	Mean	Std. dev.	Obs.
Is a woman	.72	.45	163
Age	29	6	163
Years of education	8	6	163
Has a partner	.52	.50	163
Any childcare	.36	.48	163
Income-Generating Activities			
Any IGA, past 3 months	.70	.46	163
Any earnings, past 30 days	.62	.49	163
Earnings, past 30 days (FCFA)	54760	44596	101
Any earnings, past 7 days	.48	.50	163

Notes: Childcare data from the past 7 days. Earnings in the past 30 days are conditional on having had any earnings. Exchange rate in December 2022: 620 FCFA = 1 USD.

The focal job seekers' networks are small and characterized by homophily. When solicited for referrals at the hiring stage, the eligible focal job seekers refer, on average, 2.9 network members. The number of referrals is balanced across the three treatment arms, as expected, since network members were elicited prior to random treatment assignment. The distribution of referrals, as well as its average by treatment arm, is represented in Appendix Figure A2. 69% choose to refer at most three network members: for them, all network members receive a job offer. These networks are characterized by strong homophily: the focal job seekers tend to refer individuals with similar socio-economic characteristics. Most network members are friends with the focal job seeker (as opposed to family members or acquaintances). Most are neighbors: if they are assigned to work at the same time and place, they have the possibility to commute together.

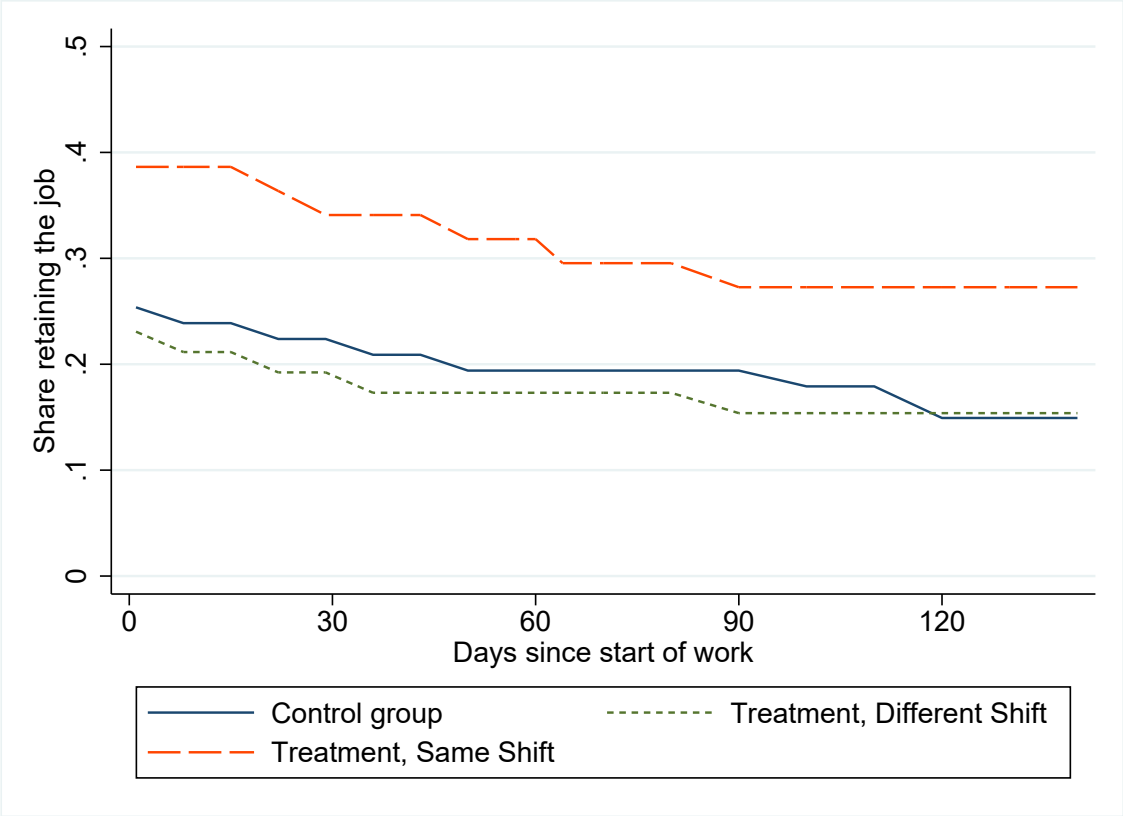
3.6 Factory: Treatment Effects

We find evidence for the existence of complementarities in labor supply, influencing job take up and retention when the focal job seekers and their network members can interact.

The labor supply of the focal job seekers assigned to the Control group illustrates the hiring and retention difficulties facing our partner company at prevailing wages. The focal job seekers assigned to the Control group receive an offer for a job paying over twice their baseline earnings for similar hours worked (for those who had any income-generating activity at baseline), and for which they had been pre-selected based on their potential interest. Despite this, only 25% decide to take up the job offer and come to work at least one day at the factory (Figure 1). Further, among those who do take up the offered job, 41% quit

within four months. This retention rate is on par with that of non-experiment workers at comparable factories. Such a stark pattern represents the hiring and retention difficulties reported by our partner companies, also documented by other firms across sub-Saharan Africa seeking to hire blue-collar workers.

Figure 1: Job Take-Up and Retention at the Factory



Notes: Share of the focal job seekers retaining the job after X days since the start of work, with the number of days on the horizontal axis. Raw data; the associated regression results are presented in Table 2.

Our intervention induces exogenous variation among focal job seekers in whether their network members work at the factory. For focal job seekers assigned to the Control group, none of their network members receive a job offer. As a consequence, none of their network members work at the factory. By contrast, up to three network members of each focal job seeker assigned to either the Different Shift or Same Shift treatment arm receive a job offer to work at the factory. These network members have a take-up rate of 8%: on average, 15% of the focal job seekers assigned to the Different or Same Shift groups have at least one network member coming to work at the factory (Appendix Figure A3A). The take-up rate is larger for network members of the focal job seekers assigned to the Different Shift group rather than the Same Shift group, but the difference is not statistically significant at conventional levels (Appendix Figure A3B).

Table 2: Treatment Effects from the Factory Experiment

Retain the job:	$\geq$ 1 day	$\geq$ 8 days	$\geq$ 60 days	$\geq$ 90 days	$\geq$ 120 days
	(1)	(2)	(3)	(4)	(5)
Different shift	-0.0189	-0.0218	-0.0149	-0.0361	0.00705
	(0.0803)	(0.0780)	(0.0721)	(0.0701)	(0.0671)
	[0.814]	[0.780]	[0.836]	[0.607]	[0.916]
Same shift	0.159*	0.177*	0.150*	0.103	0.143*
	(0.0941)	(0.0926)	(0.0876)	(0.0850)	(0.0830)
	[0.093]	[0.057]	[0.088]	[0.226]	[0.087]
Control mean	0.25	0.24	0.19	0.19	0.15
Obs.	163	163	163	163	163
p-val, equal coefs	0.060	0.033	0.065	0.103	0.112

Notes: Comparison of means, OLS. Robust standard errors are in parentheses, and p-values in brackets.
* denotes significance at the 10% level.

The labor supply of the focal job seekers is strongly influenced by their network members receiving a job offer to work the same shift. The take-up rate of the focal job seekers assigned to the Same Shift treatment arm is 16 percentage points higher than those assigned to Control (p-value: 0.093). Given the low take-up rate in the Control group, this represents a remarkable 63% increase. This effect is persistent over time: the share of the initial focal job seekers who are working at the factory for at least one week, one month, or four months, remains 10-18 percentage points higher in the Same Shift group relative to Control (Table 2).

Allowing for interactions during work-related time between the focal job seekers and their network members is necessary to induce these effects. There is virtually no difference in the labor supply of the focal job seekers in the Different Shift treatment arm relative to those in Control—whether we consider job take-up or retention (Table 2). The Same Shift effects relative to Control are thereby entirely driven by the fact that the focal job seekers and their network members can commute together and potentially interact at work (since they work at the same location and time). Simply offering a job to network members, to work at a different time, is not sufficient to induce a labor supply response. This pattern of results thereby indicates the presence of complementarities in labor supply driven by direct interactions.

The effects of the Same Shift treatment are concentrated among job seekers with longer commuting times to the factory. As shown in Appendix Figure A5, the focal job seekers in the Control and Same Shift groups with limited commuting time to the factory have a similar likelihood of taking up the job offer. The difference emerges when considering the

focal job seekers with longer commuting time: the focal job seekers in the Control group become less likely to take up the offered job when their commutes are longer, while the focal job seekers in the Same Shift group retain a similar take-up rate. However, this evidence is only suggestive. First, the limited sample size of the factory experiment makes heterogeneity analyses underpowered. Second, distance between home and work can be correlated with other characteristics influencing treatment effects. Nonetheless, it suggests a potential role for interactions during commute as a driver of complementarities. We assess this channel more rigorously through the design of the sales experiment and by considering interactions during commute on the intensive margins of labor supply at the factory.

3.7 Factory: Intensive Margins of Labor Supply

To provide direct evidence of the role of interactions during commute as a condition for the complementarities to influence labor supply decisions, we turn to the intensive margin of labor supply. Because the number of focal job seekers deciding to take up the offered job is too small to estimate their intensive-margin labor supply decisions at work, we consider non-experiment factory workers.

We conducted a network survey at the factory to test for the coordination of attendance and quits among co-commuters. We surveyed 726 factory workers, hired after our first experiment. We elicited from each of these workers the list of the co-workers with whom they commute, as well as whether they otherwise interact at or outside work and the type of interaction. This allows us to identify the specific role of interactions during commute as a mechanism for the complementarities in labor supply.

We seek to identify coordination in attendance and quits among co-commuters. We thus need to rule out common shocks that could confound our estimates by jointly affecting the labor supply of workers commuting together. We do so by estimating the following equation:

$$\text{notAtWork}_{it} = \omega + \rho \text{NumberCoCommutersNotAtWork}_{it} + \delta_t + \gamma_i + \chi_{it} + \zeta_{it} \tag{2}$$

Where notAtWork_{it} is a dummy variable indicating whether worker i is absent from the factory on day t (when testing for coordination in quits, that dummy equals one on the day at which the worker left the factory). $\text{NumberCoCommutersNotAtWork}_{it}$ is the number of the co-commuters of worker i who are absent on day t (or who left the factory on day t, when considering quits). All specifications include worker fixed effects γ_i and day fixed effects δ_t. When studying attendance, we flexibly control for the network composition of each worker to ensure that entry and exit of network members at the firm do not bias the estimates. To rule out common shocks as a driver of the estimated correlation in attendance and quits, and to isolate coordination decisions, we include a rich set of fixed effects in χ_{it}. We include day-by-neighborhood fixed effects to rule out common shocks at the neighborhood level, such

as weather shocks or transport disruptions that could affect all workers living in the same neighborhood. We also include day-by-production unit fixed effects to rule out common shocks at the production unit, such as a decision by managers to operate at lower capacity for a given day. In order to claim that the estimated ρ does not reflect the coordination in labor supply among co-commuters, one would need to imagine shocks that jointly influence the labor supply of co-commuters while not influencing the labor supply of workers living in the same neighborhood or working in the same production unit.[10]

Table 3: Coordination of Absences and Quits at the Factory

	(1)	(2)	(3)	(4)
Panel A: Is Absent from Work				
Nbr Co-Commuters absent	0.289***	0.276***	0.140***	0.128***
	(0.0241)	(0.0312)	(0.0159)	(0.0183)
	[0.000]	[0.000]	[0.000]	[0.000]
Mean: Covar = 0	0.15	0.14	0.15	0.14
N: Worker-days	73406	52340	72059	50874
N: Workers	247	171	247	171
Panel B: Is Leaving the Factory				
Nbr Co-Commuters leaving	0.183***	0.0672**	0.107***	0.0672**
	(0.0245)	(0.0279)	(0.0227)	(0.0279)
	[0.000]	[0.017]	[0.000]	[0.017]
Mean: Covar = 0	0.00	0.00	0.00	0.00
N: Worker-days	90231	62966	88937	62966
N: Workers	254	178	254	178
Day-neighborhood FE	No	Yes	No	Yes
Day-prod. unit FE	No	Yes	Yes	Yes

Notes: All columns control for day and worker FE; plus links FE in Panel A. A link FE is created for each co-commuter and equals 1 while the co-commuter works at the factory. Robust standard errors are in parentheses.

We find evidence of strong coordination in attendance and quits among co-commuters, underscoring the important role that joint commuting plays in generating complementarities in labor supply. A worker is 29 percentage points more likely to be absent on days when a co-commuter is absent from the factory than on days when all co-commuters are present at the factory (Table 3, Panel A, Col. 1). These results are not driven by common shocks at the neighborhood or production unit levels (Table 3, Panel A, Cols. 2-4). This coordination of labor supply among non-experiment factory workers extends beyond attendance decisions

[10] In the absence of an instrument for network members' labor supply in this sample, and given the reflection problem when identifying endogenous social effects (Manski 1993), we do not attempt to interpret the estimated coefficients as reflecting the magnitude of the causal effect of network members' labor supply on workers' labor supply. Rather, we seek to test for the presence of a social effect: absent direct complementarities in labor supply among co-commuters, the estimated effect would be zero.

and into quits. A worker is 18 percentage points more likely to quit the factory on a day when a co-commuter quits than on a day when no co-commuter quits the factory (Table 3, Panel B, Col. 1). As for attendance, these results are not fully driven by common shocks at the neighborhood or production unit levels (Table 3, Panel B, Cols. 2-4) and are robust to focusing on the set of co-commuters with whom the workers report having no interactions at work or outside work besides commuting together.

4 Sales Experiment

Our second experiment, implemented with short-term sales jobs, also tests for the existence of complementarities in labor supply and additionally isolates the specific role of interactions during commute as a mechanism. In what follows, we describe the setting of this second experiment, its design and implementation (4.2), the data and estimation strategy (4.3), and the sample of focal job seekers (4.4). The results from the experiment are then presented: the main tests for the existence of complementarities and the underlying mechanisms (4.5), and the implications for workers' productivity (4.6).

4.1 Sales: Setting

Our second experiment is with Coin Afrique, a firm working in the service sector. This allows us to show the relevance of our results across settings and in tertiary occupations, a major driver of economic development in sub-Sarahan Africa (Gollin et al. 2016; Rodrik 2016). Coin Afrique is the largest digital classified ads company—similar to Craigslist in the US—in francophone Africa, established in 2015.

Coin Afrique conducted a large-scale marketing campaign in Abidjan to increase the user base of their app. They hired a large number of street canvassers, whose primary role was approaching people, talking to them about Coin Afrique, and getting them to download the app on their smartphones. For such street canvassing campaigns, the gender composition is usually more balanced than the Olam factory jobs. Street canvassing for Coin Afrique also requires more specific skills than working at Olam, such as a basic knowledge of how to work a smartphone and ability to trouble-shoot potential technical issues on the fly. It also requires attention to work quality: the firm wants workers to sign up as many potential customers as possible, but also wants these potential customers to use the app in the future. As a result, the pool of potential workers recruited for the sales job mostly consisted of university graduates. The offered pay was also higher than for the factory jobs and set at 9,500 FCFA (∼15 USD) for each day of work. The campaign was implemented in two nearly identical waves, and the sales jobs offered 5 days of work.[11]

[11] We started with 5 days of work for the first wave of the intervention. Then, we were able to expand the sample size in a second wave of the exact same intervention. Due to budget constraints, we were only able

Coin Afrique subcontracted with a local marketing company to implement the campaign: Atom-BTL Marketing. Atom-BTL has extensive experience working with large formal companies in Côte d'Ivoire. They were in charge of finding and hiring the workers and managing the work. They usually hire through a combination of recommendations (including recommendations at hiring), social media ads, word of mouth, and hiring scouts—all common hiring practices in the study setting.

4.2 Sales: Design

The sales experiment has two objectives. First, to verify that the existence of complementarities in labor supply generalizes to other settings. Second, to separate the mechanisms of interactions during commute and interactions at work. The design of the sales experiment leverages an institutional feature of the offered sales jobs: work is performed at either of two worksites, which are located in different boroughs of the city. The distance is such that workers assigned to different worksites will not be able to interact during the day, including during the lunch break. It also reduces their possibility of commuting together to work.

We use this feature and randomly assign the focal job seekers to one of three groups. In the "Control'' group, the network members of the focal job seeker do not receive a job offer. In the "Different Site'' group, the network members of the focal job seeker receive a job offer but will not work at the same worksite as the focal job seeker. In the "Same Site'' group, the network members of the focal job seeker receive a job offer and they will work at the same worksite as the focal job seeker. For each focal job seeker, we cross-randomize this primary treatment assignment with their assignment to either of the two possible worksites. Conditional on their commuting time to each of the two potential assignments, this induces exogenous variation in the focal job seekers' commuting time. Appendix Figure A7 represents this design.

This design allows us to achieve the two objectives of the experiment. Comparing the labor supply decisions of the focal job seekers assigned to the Same Site group relative to those assigned to Control yields the broadest test for the presence of complementarities in labor supply. However, as discussed in section 3.2, this comparison will not exactly identify complementarities if the focal job seekers respond to the simple fact that their network members receive a job offer—even if the network members do not change their labor supply. To alleviate this concern and obtain a clean test for complementarities in labor supply, we can compare the Same Site and Different Site groups. This isolates the effects of exogenously varying interactions at work and during commute. In this experiment, the two worksites have the same work hours. As such, being assigned to either the Same Site or the Different Site

to offer the jobs for 4 days and no longer 5. Except for this difference, everything else is identical. Overall, 29% of the focal individuals in the final sample come from this second wave. We control for the intervention wave in our specifications to account for this.

group has no influence on the possibility of interacting with network members outside work and commuting.

The exogenous variation in commuting time enables us to isolate the role of interactions during commute from interactions at work. If complementarities influence labor supply decisions due to interactions during commute, we can expect that individuals with longer commutes will benefit more from their network members being hired, but only if they can commute together. This yields two predictions for the heterogeneity of treatment effects by the exogenously varied commuting time. First, we would expect larger Same Site treatment effects for focal job seekers with longer commuting time. Second, we would not expect the Different Site treatment effects to vary with the commuting time of the focal job seekers.[12] Our design allows us to go further than testing for the role of interactions during commute. It also enables us to compare the importance of interactions during commute with interactions at work. This can be done by comparing the Same Site treatment effects for individuals with exogenously short commute time—which would mostly reflect other mechanisms than commute—to the same effects for individuals with exogenously longer commute times.

The implementation of this design closely follows the procedures from the factory experiment. We therefore refer the reader to its description in section 3.3.

4.3 Sales: Data and Estimation

We obtain access to the firm records that compile each worker's attendance and production each day, as well as records from the e-commerce app providing quality-adjusted productivity measures. The firm records allow us to measure job take-up and daily productivity for each focal job seeker and their network members. Unlike the factory experiment in which jobs were offered with open-ended contracts, the jobs offered in the sales experiment are for 5 days of work and retention is not a first-order priority for the firm. Productivity, however, is. The firm records, for each worker and each day, the following four dimensions of sales: the number of people who (i) were approached by the worker during the day to download the app, (ii) accessed the app or associated website with the help of the worker, (iii) created an account with the help of the worker and (iv) turned on app notifications with the help of the worker.

We supplement the factory records with a second, independent, source of sales data. We worked with the firm to ensure that any download of the app during the street-canvassing campaign could be linked to the worker who induced the download. This is done by assigning

[12] Besides interactions during commute, heterogenous treatment effects could also arise if individuals value being at work with network members more when they work far from their home than if they work close to their home. This could happen, for instance, if working farther from home entails more learning—about the type of customers, about the best commuting route, etc.—thus increasing the value of joint learning. This should be especially important when the choice is between working in one's own neighborhood or outside. We therefore verify the robustness of the heterogeneity patterns to excluding the focal job seekers for whom one of the two worksites is located in their own neighborhood.

each worker an individual QR code and having potential customers download the app through the QR code. Besides providing an alternative independent measure of sales, this approach yields one major benefit: it allows us to track the activity of users on the app. The objective of the firm was to induce people to actively use the app for purchases and sales, rather than have merely have a downloaded but dormant app on their phone. For each user, we therefore measure, for 20 days post-download, (i) the number of times the customer has accessed the app and (ii) the number of times the user has attempted to contact a seller through the platform (specifically, clicked on the button to see the seller's contact information). This provides us with measures of quality-adjusted productivity for each worker. A downside is the presence of missing observations for 13% of the workers, driven either by workers not using their QR-code properly or by technical limitations.[13] In our analyses, we therefore present results for both types of productivity data.

The estimation strategy closely follows that of the factory experiment, described in section 3.4. Our main specification estimates the Intent-to-Treat effects on the job take-up and productivity decisions of the focal job seekers. We implement a simple comparison of means at the job seeker level and estimate:

$$y_i = \alpha^s + \beta_1^s \text{DifferentSite}_i + \beta_2^s \text{SameSite}_i + \gamma^s X_i + \epsilon_i^s, \tag{3}$$

where y_i is the outcome of interest for focal individual i, DifferentSite_i and SameSite_i are indicator variables for the job seeker's assignment into the "Different Site" or "Same Site" groups, and X_i a vector of controls. X_i includes an indicator for the wave of the campaign, the number of network members offered a job, and the share of neighbors among the network members offered the job.

Our main outcome of interest is job take-up: a binary measure indicating whether the focal individual ever came to work. Before being invited to come to work, prospective workers underwent a short training session during which they could ask lingering questions about the job. Similarly to the factory job, this reduces the likelihood that workers are surprised on their first day of work (leading them to quit and not come back). Only 2% of the focal job seekers who came to work at least one day didn't come back for a second day; any potential surprises on the first day of work are not driving our results.

Our secondary outcome is focal job seekers' productivity. Since observing productivity is conditional on job take-up, the estimates will capture a mix of direct and selection effects. We still consider them, since they are of first-order relevance for firms. The sales jobs are offered for a short-term campaign, so retention is not an outcome of interest.

When testing for heterogeneity in treatment effects by commuting time, we interact the indicator variables DifferentSite_i and SameSite_i with the relevant variable for commuting time: either a continuous measure of round-trip commuting time to the assigned worksite,

[13] Some types of phones were not supported by the system.

elicited at baseline from the focal job seekers, or an indicator for the commuting time to the assigned worksite being above the median. Since distance to work can be correlated with other worker characteristics, we isolate exogenous variation in commuting time by controlling for commuting times to the two potential worksites as well as their interaction.

4.4 Sales: Sample

Table 4: Sample Description for the Sales Experiment

Variable	Mean (1)	Std. dev. (2)	Obs. (3)
Demographics			
Is a woman	.37	.48	873
Age	25	5	873
Years of education	15	2	873
Has a partner	.1	.3	873
Has young children	.18	.38	873
Income-Generating Activities			
Any IGA, past 3 months	.62	.49	873
Any earnings, past 30 days	.38	.49	873
Earnings, past 30 days (FCFA)	104983	118149	336
Any earnings, past 7 days (FCFA)	.2	.4	873
Currently a student	.41	.49	873

Notes: Earnings in the past 30 days are conditional on having had any earnings. Exchange rate in December 2022: 620 FCFA = 1 USD.

For the sales experiment, we work with a sample of 873 eligible job seekers, described in Table 4. These job seekers tend to be young and educated and correspond to the job seekers expected to enter the labor market in large numbers across sub-Saharan Africa in the coming years. In our sample, the average job seeker is 25 years old and has an undergraduate education. Few have a partner or children. From a short-term pecuniary perspective, the offered sales jobs are desirable relative to the outside option of these focal job seekers. Only 40% had any income-generating activity in the past 30 days, and 20% in the past 7 days. Even for those with an income-generating activity, the offered jobs pay over twice their baseline average earnings. The sample is balanced across these characteristics, as reported in Appendix Table A2.

As in the factory experiment, the networks are small and characterized by considerable homophily. When solicited to provide referrals at the hiring stage, the eligible focal job seekers refer, on average, 1.8 network members. The number of referrals is balanced across the three treatment arms, as expected since it has been elicited at baseline. Its distribution

and its average by treatment arm are represented in Appendix Figure A8. 91% choose to refer at most three network members; for them, all network members will receive a job offer. The focal job seekers tend to report network members of the same gender and with similar ages and levels of education. At the time of the referral, only 11% of the network members had an income-generating activity, and 17% were studying. The others were either spending time with friends or at home.

4.5 Sales: Treatment Effects

The results from the sales experiment confirm the existence of complementarities in labor supply and identify interactions during commute as the main underlying mechanism. As in the factory experiment, focal job seekers are significantly more likely to take up the offered jobs when their network members receive a job offer to work in the same location and at the same time. Specifically, focal job seekers assigned to the Same Site group are 7 percentage points more likely to take up the offered sales jobs than those assigned to the Control group (p-value: 0.092, Table 5, Col. 1). This represents a 20% increase over the take-up rate in the Control group. There are no significant differences between the Same Site and Different Site groups, on average.[14]

The results from the heterogeneity analysis are directly consistent with complementarities influencing labor supply decisions due to interactions during commute. The Same Site treatment effects are significantly larger for job seekers with longer commutes, while the Different Site treatment effects are not (p-values: 0.077 and 0.486, Table 5, Col. 2). The results are similar when we use exogenous variation in commuting time to estimate the heterogeneity in treatment effects (Table 5, Col. 3). This indicates that commuting time drives the heterogeneity, and not characteristics of the focal job seekers that could be correlated with distance to work.

The results indicate that interactions during commute are the main mechanism through which the complementarities influence labor supply. The Same Site treatment effects are indeed concentrated among the focal job seekers with above-median commute time. For the focal job seekers with below-median commute time, the point estimate of the effect of being assigned to the Same Site group relative to Control is virtually zero (Table 5, Cols. 4-5). This pattern of results is confirmed when estimating the treatment effects separately by deciles of commute time: they are concentrated among job seekers assigned to the top four deciles of commute time (Appendix Figure A10). This directly points to

[14] A possible reason for the absence of significant differences is that while the focal job seeker may value working with their network members, these network members do not derive the same value from working with the person who referred them. Another comes from the fact that the network members only receive information from the firm about the job being offered to them. They are not told who referred them for the job, or whether they would be working at the same worksite. If the focal job seekers and their network members do not communicate together before the first day of work, then there would not be a difference in take-up rate by the network members.

interactions during commute—such as joint leisure or increased safety—as the key drivers of the complementarities in labor supply.

Qualitatively, study participants report both increased safety and more pleasant commutes as key benefits from having network members be hired and offered to work at the same worksite. Safety tends to be reported more often by women, while men more frequently report that commuting together is more pleasant. However, heterogeneous treatment effects by gender indicate that men and women have similar responses to the treatment. Results are thus not fully driven by gender-specific safety concerns, with co-commuting valued for different reasons by both men and women.

Table 5: Treatment Effects in the Sales Experiment

		Commuting time to assigned worksite			
		Hours (round-trip)		Is Above Median	
	(1)	(2)	(3)	(4)	(5)
Covariate		-0.0125	-0.00446	-0.0634	-0.0520
		(0.0234)	(0.0288)	(0.0648)	(0.0707)
		[0.592]	[0.877]	[0.328]	[0.462]
Different Site	0.0370	-0.00964	-0.00618	0.00108	0.00154
	(0.0403)	(0.0777)	(0.0780)	(0.0499)	(0.0499)
	[0.359]	[0.901]	[0.937]	[0.983]	[0.975]
Different Site X Covariate		0.0223	0.0205	0.108	0.105
		(0.0320)	(0.0322)	(0.0874)	(0.0875)
		[0.486]	[0.525]	[0.216]	[0.232]
Same Site	0.0679*	-0.0527	-0.0511	-0.00497	-0.00461
	(0.0403)	(0.0769)	(0.0774)	(0.0495)	(0.0496)
	[0.092]	[0.494]	[0.510]	[0.920]	[0.926]
Same Site X Covariate		0.0567*	0.0557*	0.202**	0.198**
		(0.0320)	(0.0324)	(0.0873)	(0.0877)
		[0.077]	[0.086]	[0.021]	[0.024]
Take-up rate in control	0.34	0.34	0.34	0.34	0.34
Controls: commute time	No	No	Interacted	No	Interacted
Mean covariate		2.06	2.06	0.34	0.34
Obs.	873	873	873	873	873

Notes: Unit of observation is job seeker. Dependent variable is 1 if the job seeker came to work at least one day. Commuting time self-reported at baseline by the job seekers, in hours, round-trip, between home and each of the two possible worksites. Each job seeker is randomly assigned to one of the two possible worksites. Covariate in Cols. (2)-(3) is the commuting time between home and the assigned worksite. Covariate in Cols. (4)-(5) is an indicator if the round-trip commuting time to the assigned worksite is greater than 2 hours (the median value). Cols. (3) and (5) control for the distance between home and each of the two potential worksites, and their interactions. Robust standard errors in parentheses, p-values in brackets.

4.6 Sales: Productivity

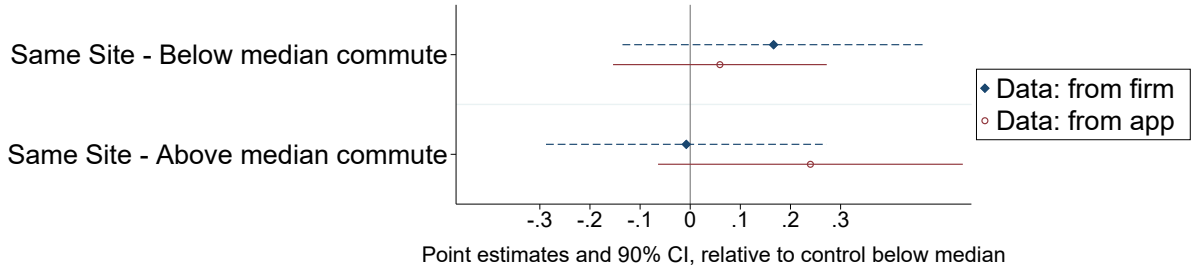
For firms, understanding the consequences of hiring network members on workers' productivity is of first-order importance. Two independent sources of productivity measures exist for the sales jobs, described in detail in section 4.3. Since each source measures productivity across multiple dimensions, we construct a standardized index of productivity for each source. This enables the direct comparison of results across the two measures and follows best practices to reduce the risk of p-hacking (Anderson 2008).

We estimate the difference in productivity across groups of workers. First, we compare the productivity of focal workers with below-median commute time assigned to the Control group to the productivity of focal workers also with below-median commute time but assigned to the Same Site group. Since there is no difference in job take-up across these two groups (section 4.5), the estimated difference in productivity can be interpreted as coming from enabling direct interactions at work between the focal job seekers and their network members. Second, we compare the productivity of focal workers with below-median commute time assigned to the Control group to the productivity of focal workers with above-median commute time assigned to the Same Site group. The difference comes from a combination of the potential direct effects of interactions between the focal job seekers and their network members with the selection effect of which focal job seeker decides to take up the job. Third, we compare the productivity of the focal job seekers who take up the job, across all groups, to the productivity of the network members who take up the job. Ex ante, the direction of each of these comparisons is ambiguous based on the academic literature reviewed in the introduction.[15] Discussions with staff at the partner firm also yielded ambiguous predictions, although their primary concern was whether the intervention would lead to lower productivity among the workers taking up the job.

We find that hiring through networks did not reduce the productivity of the sales workers. First, as reported in Figure 2, the focal workers assigned to the Same Site groups do not have significantly lower productivity at work. The point estimate from comparing focal job seekers with below-median commute time across the Same Site and Control group is positive, but not statistically significant. This result is consistent across the two sources of data. When considering the focal job seekers in the Same Site group with above-median commute time, for which the comparison to Control entails both a direct and selection effect, the point estimate is positive when using data from the app and negative but close to zero when using data from the firm. These results indicate that the direct and selection effects do not significantly reduce the productivity of workers relative to hiring individuals alone. Second, we also find that the focal workers and the network members have similar productivity at work. This result indicates that there is no productivity trade-off between those providing

[15] There could also be additional productivity gains from workers collaborating with peers sharing similar characteristics, including ethnicity (Hjort 2014).

Figure 2: Productivity Relative to Control



Notes: Estimated difference in productivity between focal workers in the Control group with below-median commute time and either the Same Site group with below-median commute time or the Same Site group with above-median commute. The point estimate and associated 90% confidence intervals are reported. The outcomes are the productivity indices from each of the two sources of productivity data (the firm's field managers and the app), expressed in standard deviations.

the referrals and those being referred.

5 Amenity Value

Next, we estimate the amenity value for the focal job seekers of offering their network members a job—specifically, one that allows them to interact at work and during commuting. We build a simple model to guide the estimation and calibrate it using natural wage variation. We find that wages would need to be increased by 13-20% to achieve the same increase in take-up as being assigned to the Same Site group relative to Control in the sales experiment.

We use natural variation from the sales jobs. Between the pilot of the sales experiment and its main implementation, the offered daily wage increased from 7,000 to 9,500 FCFA ($\sim$ USD 11 to 15). This wage increase was explicitly decided to increase job take-up. As such, the resulting take-up increase can be taken as an upper bound on the take-up response that would have been observed from a random wage increase. This implies that our calibration yields a lower bound on the amenity value of being hired with network members and the possibility of interacting.

We cannot directly use the outcomes from the natural wage variation to estimate the amenity value. On the one hand, we have an increase in job take-up estimated from the focal job seekers being assigned to the Same Site group relative to Control in the sales experiment. We seek the wage increase equivalent to the amenity of being assigned to the Same Site group rather than Control. Its effect on job take-up was estimated while keeping the underlying network fixed: the referrals were elicited from the focal job seekers at baseline, before treatment was announced, and are balanced across the Same Site and Control groups. On the other hand, we have an increase in job take-up resulting from a wage increase. This

is the effect that we want to use to estimate the amenity value in terms of wage increase. In this case, however, the network was not kept fixed: the wage increase was announced before referrals were elicited. As such, the wage increase may have changed both whether the focal job seekers provided any referrals and, if they did, whether they took up the job. We thus need to correct for the change in networks when comparing the job take-up responses.

We build a model to help address this challenge and guide the estimation of the amenity value. A brief overview is provided here; see Appendix Section A.3 for a detailed description. Since the endogeneity of the network elicitation causes the estimation challenge, we explicitly model it. Job seekers make two sequential decisions. In the first stage, they decide whether to make a referral. In the second stage, they decide whether to take up the offered job. If their referral also receives a job offer, they make a simultaneous decision and can communicate to coordinate. We solve for equilibrium decisions.

We calibrate the model and estimate an amenity value equivalent to a 13-20% wage increase. We assume that the cost of effort is heterogeneous across the focal job seekers. Given the available empirical evidence, we can estimate the amenity value for any distribution of the cost of effort that has at most two parameters. Assuming a uniform distribution yields the following estimating equation:

$$\delta = (y_{1t} - y_{0t})\, p_{0t} \frac{\Delta\omega}{(y_{1\omega} p_{1\omega} - y_{0\omega} p_{0\omega})} \tag{4}$$

Where δ is our object of interest: the wage increase necessary to increase take-up as much as the Same Site treatment does. The subscripts t and ω denote, respectively, the Same Site *t*reatment and the natural *w*age increase. As such, $y_{1t} - y_{0t}$ is the difference in job take-up between the Same Site treatment relative to Control (estimated on the set of all focal job seekers assigned to Same Site or Control, whether they provide a referral or not); $y_{1\omega}$ is the job take-up rate among all focal job seekers after the natural wage increase; $y_{0\omega}$ is the job take-up rate among all focal job seekers before the natural wage increase; $\Delta\omega$ is the natural increase in the wage, in FCFA; and $p_{1\omega}, p_{0\omega}, p_{0t}$ are the share of all focal job seekers who provide at least one referral, respectively after the natural wage increase, before the natural wage increase, and in the Control group of the experiment.

We obtain, under this uniform assumption, that increasing job take-up as much as the Same Site treatment effects (i.e., by 8 percentage points) requires increasing the offered wage by 20%. Assuming a normal distribution of the cost of effort across job seekers would instead imply an amenity value equivalent to a 13% wage increase.

6 Conclusion

In this paper, we demonstrate that complementarities in labor supply exist and that they meaningfully affect employment decisions. We also show that the ability to commute together is the primary driver underlying the observed labor complementarities. Our results have implications for our understanding of key features of labor markets, firms' personnel policies, and transportation programs.

Job seekers who are offered employment are significantly more likely to take it up when their network members are also offered a job, but only if they would be working at the same location and time. This result replicates across two different samples of job seekers, each offered a different type of job in Abidjan: long-term factory employment and short-term sales work. The increase in job take-up does not come at the expense of retention (evidence from the factory experiment) or productivity (evidence from the sales experiment).

The effects of job seekers being offered to work at the same place and time are concentrated among job seekers with longer commute times—including when commute times are exogenously determined, as in the sales experiment. Further, non-experiment factory workers who commute together coordinate their absences and quits, in a way that is not driven by common shocks in their neighborhood or production unit. Guided by a model and leveraging natural wage variation, we estimate the amenity value of being offered a job together with network members with the ability to interact during commute to be equivalent to an increase in wages by 13-20%.

An immediate implication of our results concerns firms' personnel policies. We demonstrate that eliciting and hiring referrals at the hiring stage and enabling them to interact during commute with the person who referred them can increase labor supply to the firm at prevailing wage, in a substantial and sustained way. This result dovetails with recent work showing that entrepreneurial training and clinic visit take-up and effectiveness can increase when potential participants are invited with network members, especially for women (Field et al. 2016; Anukriti et al. 2022).[16] Our results show that this network effect also applies to labor supply, and suggest that firms may be able to boost worker labor supply by choosing to hire through referrals from their incumbent workers.

Our results also underscore an important aspect of efforts to increase diversity in organizations. A strong focus, including in the academic literature, has been on expanding access to job opportunities beyond the networks of incumbent workers (Calvó-Armengol and Jackson 2004; Fernandez and Sosa 2005; Fernandez and Fernandez-Mateo 2006; Calvó-Armengol and Jackson 2007; Wahba and Zenou 2005; Rubineau and Fernandez 2013; Galenianos 2014,

[16] We bring this idea to the domain of labor supply and thereby add network-based hiring to a growing literature designing and evaluating active labor market interventions in lower-income settings. Previously-considered interventions include mentoring (Alfonsi et al. 2022), training (Alfonsi et al. 2020), matching (Bandiera et al. 2022), skills assessments (Carranza et al. 2022b; Kiss et al. 2023), job search training (Wheeler et al. 2022), and transport subsidies (Caria et al. 2022; Abebe et al. 2020).

2020; Bolte et al. 2020). The underlying rationale is that referral networks exhibit strong homophily, so that they cannot be relied upon to diversify the talent pool (Beaman et al. 2018; Jackson 2020). Our findings highlight that worker *take-up* of offered jobs is an important margin for diversifying workplaces. In particular, offering jobs to several members of the same network or social group can increase their acceptance of the job, leading to more diversification in practice. More broadly, considering the presence of multiple equilibria in labor supply for specific firms, sectors, or neighborhoods driven by complementarities in labor supply can be a fruitful lens to reduce the sorting of individuals from different gender or ethnic groups across different firms and sectors, and ultimately reduce existing gender and ethnic gaps.

When deciding whether to hire through networks, firms face a trade-off between the level and volatility of labor supply. Even if it does not necessarily come at a productivity cost, firms still face a trade-off when deciding whether to hire through referrals or on the open labor market. Increasing the level of labor supply can come at the expense of increasing its volatility if network members coordinate their absences and quits while at work. The level of inter-dependency in the production processes will likely influence which effect dominates. More broadly, how firms weigh the effects of referral-based hiring on the referred workers relative to the effects on those providing the referrals, and ultimately decide how to hire, remains an open question.

While our study is set in urban Côte d'Ivoire, the value workers place on interacting with others during commuting yields useful policy insights for both lower- and higher-income countries. Even in the United States, the average round-trip commute is 54 minutes long (Aksoy et al. 2023) and workers report it as a major disamenity (Dube et al. 2022). Encouraging carpooling or arranging company shuttles could be an effective way to stimulate labor supply and improve workers' retention—while also having notable environmental benefits. This could be especially useful for the large class of occupations, such as manufacturing and essential services, where working from home is not an option.

More broadly, the existence of strong complementarities in labor supply through interactions during commute implies a social multiplier that amplifies negative effects of disruptions to the transportation network. Heavy rains drastically increasing congestion and flash floods making parts of the transportation system inoperable will not only disrupt the trade of goods across firms, they can also have major effects on firms' workforce.[17] So will extreme heat events that exacerbate the disutility from time spent in (informal) public transport with limited air conditioning. The effects of these events on the workforce will be amplified by co-commuters influencing each other's labor supply—even when the disruptions affect a portion of the commute used by one network member only. This has implications for the design of resilient urban transport networks and the optimal hiring strategy of firms, which we leave for future work.

[17]For the effect of floods on supply chain networks, see for instance Balboni et al. (2023).

References

Abebe, Girum, Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn. 2020. "Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City." *The Review of Economic Studies* 88 (3): 1279–1310. (Cited on pages 5, 6, 27).

Afridi, Farzana, and Amrita Dhillon. 2022. *Social Networks and the Labour Market.* Working Paper, Discussion Paper Series 15774. IZA. (Cited on pages 1, 4).

Afridi, Farzana, Amrita Dhillon, Sherry Xin Li, and Swati Sharma. 2020a. "Using social connections and financial incentives to solve coordination failure: A quasi-field experiment in India's manufacturing sector." *Journal of Development Economics* 144:102445. (Cited on page 4).

Afridi, Farzana, Amrita Dhillon, Sanchari Roy, and Nikita Sangwan. 2022. *Social Networks, Gender Norms and Women's Labor Supply: Experimental Evidence Using a Job Search Platform.* Working Paper. (Cited on page 5).

Afridi, Farzana, Amrita Dhillon, and Swati Sharma. 2020b. *The Ties that Bind Us: Social Networks and Productivity in the Factory.* Working Paper, Discussion Paper Series 13711. IZA. (Cited on page 4).

Akbar, Prottoy, Victor Couture, Gilles Duranton, and Adam Storeygard. 2023. *The Fast, the Slow, and the Congested: Urban Transportation in Rich and Poor Countries.* Working Paper 31642. National Bureau of Economic Research. (Cited on pages 1, 6).

Aksoy, Cevat Giray, Jose Maria Barrero, Nicholas Bloom, Steven J Davis, Mathias Dolls, and Pablo Zarate. 2023. *Time Savings When Working from Home.* Working Paper 30866. National Bureau of Economic Research. (Cited on pages 1, 28).

Alesina, Alberto, Edward Glaeser, and Bruce Sacerdote. 2005. "Work and Leisure in the United States and Europe: Why So Different?" *NBER Macroeconomics Annual* 20:1–64. (Cited on page 5).

Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2020. "Tackling youth unemployment: Evidence from a labor market experiment in Uganda." *Econometrica* 88 (6): 2369–2414. (Cited on page 27).

Alfonsi, Livia, Mary Namubiru, and Sara Spaziani. 2022. *Meet your future: Experimental evidence on the labor market effects of mentors.* Working Paper. (Cited on page 27).

Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103 (484): 1481–1495. (Cited on page 24).

Anukriti, S, Catalina Herrera-Almanza, and Mahesh Karra. 2022. *Bring a friend: Strengthening women's social networks and reproductive autonomy in India.* The World Bank. (Cited on page 27).

Ashraf, Nava, and Oriana Bandiera. 2018. "Social Incentives in Organizations." *Annual Review of Economics,* (cited on page 4).

Balboni, Clare, Johannes Boehm, and Mazhar Waseem. 2023. *Firm Adaptation in Production Networks: Evidence from Extreme Weather Events in Pakistan.* Working Paper. (Cited on page 28).

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2009. "Social Connections and Incentives in the Workplace: Evidence From Personnel Data." *Econometrica* 77 (4): 1047–1094. (Cited on page 4).

———. 2010. "Social Incentives in the Workplace." *The Review of Economic Studies* 77 (2): 417–458. (Cited on page 4).

Bandiera, Oriana, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2022. *The Search for Good Jobs: Evidence from a Six-year Field Experiment in Uganda.* Working Paper. (Cited on page 27).

Beaman, Lori. 2016. "Social Networks and the Labor Market." In *The Oxford Handbook of the Economics of Networks,* edited by Yann Bramoulle, 649–672. (Cited on page 4).

Beaman, Lori, Niall Keleher, and Jeremy Magruder. 2018. "Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi." *Journal of Labor Economics* 36 (1): 121–157. (Cited on page 28).

Blattman, Christopher, and Stefan Dercon. 2018. "The Impacts of Industrial and Entrepreneurial Work on Income and Health: Experimental Evidence from Ethiopia." *American Economic Journal: Applied Economics* 10 (3): 1–38. (Cited on page 6).

Bloom, Nick, and John Van Reenen. 2011. "Human Resource Management and Productivity." In *Handbook of Labor Economics,* edited by Orly Ashenfelter and David Card, 4:1697–1767. Elsevier. (Cited on page 6).

Bolte, Lukas, Nicole Immorlica, and Matthew O Jackson. 2020. *The Role of Referrals in Immobility, Inequality, and Inefficiency in Labor Markets.* Working Paper. (Cited on page 28).

Borker, Girija. 2021. *Safety first: Perceived risk of street harassment and educational choices of women.* Working Paper, Policy Research Working Papers. World Bank Group. (Cited on page 5).

Calvó-Armengol, Antoni, and Matthew O Jackson. 2004. "The Effects of Social Networks on Employment and Inequality." *The American Economic Review* 94 (3): 426–454. (Cited on page 27).

———. 2007. "Networks in labor markets: Wage and employment dynamics and inequality." *Journal of Economic Theory* 132 (1): 27–46. (Cited on page 27).

Caria, Stefano, Simon Franklin, and Marc Josef Witte. 2022. "Searching With Friends." *Journal of Labor Economics,* (cited on pages 5, 27).

Carranza, Eliana, Aletheia Donald, Florian Grosset, and Supreet Kaur. 2022a. *The Social Tax: Redistributive Pressure and Labor Supply.* Working Paper 30438. National Bureau of Economic Research. (Cited on page 8).

Carranza, Eliana, Robert Garlick, Kate Orkin, and Neil Rankin. 2022b. "Job Search and Hiring with Limited Information about Workseekers' Skills." *The American Economic Review* 112 (11): 3547–3583. (Cited on page 27).

Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2017. "Peer Effects in the Workplace." *The American Economic Review* 107 (2): 425–456. (Cited on page 4).

Donovan, Kevin, Will Jianyu Lu, and Todd Schoellman. 2023. "Labor Market Dynamics and Development." *The Quarterly Journal of Economics* 138 (4): 2287–2325. (Cited on page 6).

Dube, Arindrajit, Suresh Naidu, and Adam D Reich. 2022. *Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers.* Working Paper 30441. National Bureau of Economic Research. (Cited on pages 5, 28).

Fernandez, Roberto M, and Isabel Fernandez-Mateo. 2006. "Networks, Race, and Hiring." *American Sociological Review* 71 (1): 42–71. (Cited on page 27).

Fernandez, Roberto M, and Lourdes Sosa. 2005. "Gendering the Job: Networks and Recruitment at a Call Center." *The American Journal of Sociology* 111 (3): 859–904. (Cited on page 27).

Field, Erica, Seema Jayachandran, Rohini Pande, and Natalia Rigol. 2016. "Friendship at Work: Can Peer Effects Catalyze Female Entrepreneurship?" *American Economic Journal: Economic Policy* 8 (2): 125–153. (Cited on page 27).

Field, Erica, and Kate Vyborny. 2022. *Women's Mobility and Labor Supply: Experimental Evidence from Pakistan.* Working Paper. (Cited on page 5).

Franklin, Simon. 2018. "Location, search costs and youth unemployment: Experimental evidence from transport subsidies." *Economic Journal* 128 (614): 2353–2379. (Cited on page 5).

Friebel, Guido, Matthias Heinz, Mitchell Hoffman, and Nick Zubanov. 2023. "What Do Employee Referral Programs Do? Measuring the Direct and Overall Effects of a Management Practice." *The Journal of Political Economy* 131 (3): 633–686. (Cited on page 4).

Galenianos, Manolis. 2014. "Hiring through referrals." *Journal of Economic Theory* 152:304–323. (Cited on page 27).

———. 2020. "Referral Networks and Inequality." *The Economic Journal* 131 (633): 271–301. (Cited on page 27).

Georges-Kot, Simon, Dominique Goux, and Eric Maurin. 2017. "Following the Crowd: Leisure Complementarities beyond the Household." *Journal of Labor Economics* 35 (4): 1061–1088. (Cited on page 4).

———. 2022. *The Value of Leisure Synchronization.* Working Paper. (Cited on page 4).

Glaeser, Edward L, Bruce Sacerdote, and José A Scheinkman. 1996. "Crime and Social Interactions." *The Quarterly Journal of Economics* 111 (2): 507–548. (Cited on page 5).

Glaeser, Edward L, Jose A Scheinkman, and Bruce I Sacerdote. 2003. "The Social Multiplier." *Journal of the European Economic Association* 1 (2/3): 345–353. (Cited on page 5).

Gollin, Douglas, Remi Jedwab, and Dietrich Vollrath. 2016. "Urbanization with and without industrialization." *Journal of Economic Growth* 21 (1): 35–70. (Cited on page 17).

Goux, Dominique, Eric Maurin, and Barbara Petrongolo. 2014. "Worktime Regulations and Spousal Labor Supply." *The American Economic Review* 104 (1): 252–276. (Cited on page 4).

Hallberg, Daniel. 2003. "Synchronous leisure, jointness and household labor supply." *Labour Economics* 10 (2): 185–203. (Cited on page 4).

Hamermesh, Daniel S. 2000. *Togetherness: Spouses' Synchronous Leisure, and the Impact of Children.* Working Paper 7455. National Bureau of Economic Research. (Cited on page 4).

Herbst, Daniel, and Alexandre Mas. 2015. "Peer effects on worker output in the laboratory generalize to the field." *Science* 350 (6260): 545–549. (Cited on page 4).

Hjort, Jonas. 2014. "Ethnic Divisions and Production in Firms." *The Quarterly Journal of Economics* 129 (4): 1899–1946. (Cited on page 24).

Hoff, Karla, and Arijit Sen. 2011. "The Kin System as a Poverty Trap?" In *Poverty Traps,* edited by Samuel Bowles, Steven N Durlauf, and Karla Hoff, 95–115. Princeton, NJ: Princeton University Press. (Cited on page 8).

Jackson, Matthew O. 2020. *The Human Network: How Your Social Position Determines Your Power, Beliefs, and Behaviors.* Reprint édition. Vintage. (Cited on page 28).

Kiss, Andrea, Robert Garlick, Kate Orkin, and Lukas Hensel. 2023. *Jobseekers' Beliefs about Comparative Advantage and (Mis)Directed Search.* Working Paper. (Cited on page 27).

Kreindler, Gabriel. 2023. *Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications.* Working Paper 30903. National Bureau of Economic Research. (Cited on page 5).

Kreindler, Gabriel, Arya Gaduh, Tilman Graff, Rema Hanna, and Benjamin A Olken. 2023. *Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta.* Working Paper 31369. National Bureau of Economic Research. (Cited on page 5).

Le Barbanchon, Thomas, Maddalena Ronchi, and Julien Sauvagnat. 2023. *Hiring Difficulties and Firm Growth.* Working Paper. (Cited on page 6).

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies* 60 (3): 531–542. (Cited on page 16).

Mas, Alexandre, and Enrico Moretti. 2009. "Peers at Work." *The American Economic Review* 99 (1): 112–145. (Cited on page 4).

Mas, Alexandre, and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements." *The American economic review* 107 (12): 3722–3759. (Cited on page 5).

Minster, Clotilde, Nabil Samir, and Jean-François Arvis. 2022. *When public transportation is unavailable, firms have trouble keeping their workers: Morocco case study.* https://blogs.worldbank.org/arabvoices/when-public-transportation-unavailable-firms-have-trouble-keeping-their-workers-morocco. Accessed: 2023-8-15. (Cited on page 6).

Moghadam, Hamed M, Patrick A Puhani, and Joanna Tyrowicz. 2023. *Pension Reforms and Couples' Labour Supply Decisions.* Working Paper, Discussion Paper Series 16039. IZA. (Cited on page 4).

Park, Sangyoon. 2019. "Socializing at Work: Evidence from a Field Experiment with Manufacturing Workers." *American Economic Journal: Applied Economics* 11 (3): 424–455. (Cited on page 4).

Rodrik, Dani. 2016. "Premature deindustrialization." *Journal of Economic Growth* 21 (1): 1–33. (Cited on page 17).

Rubineau, Brian, and Roberto M Fernandez. 2013. "Missing Links: Referrer Behavior and Job Segregation." *Management Science* 59 (11): 2470–2489. (Cited on page 27).

Topa, G. 2011. "Labor markets and referrals." *Handbook of social economics,* (cited on page 1).

———. 2019. "Social and spatial networks in labour markets." *Oxford Review of Economic Policy,* (cited on pages 1, 4).

Tsivanidis, John Nicholas. 2023. *The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogota's Transmilenio.* Working Paper. (Cited on page 5).

Velásquez, Daniel. 2023. *Transit infrastructure, couples' commuting choices, and gender inequality.* Working Paper. (Cited on page 5).

Wahba, Jackline, and Yves Zenou. 2005. "Density, social networks and job search methods: Theory and application to Egypt." *Journal of Development Economics* 78 (2): 443–473. (Cited on page 27).

Wheeler, Laurel, Robert Garlick, Eric Johnson, Patrick Shaw, and Marissa Gargano. 2022. "LinkedIn(to) Job Opportunities: Experimental Evidence from Job Readiness Training." *American Economic Journal: Applied Economics* 14 (2): 101–125. (Cited on page 27).

Witte, Marc. 2021. *Why do workers make job referrals? Experimental evidence from Ethiopia.* Working Paper. (Cited on page 4).

World Bank. 2018. *Cashew Value-Chain Competitiveness Project: Project Appraisal Document.* Technical report. Washington, DC: The World Bank Group. (Cited on page 6).

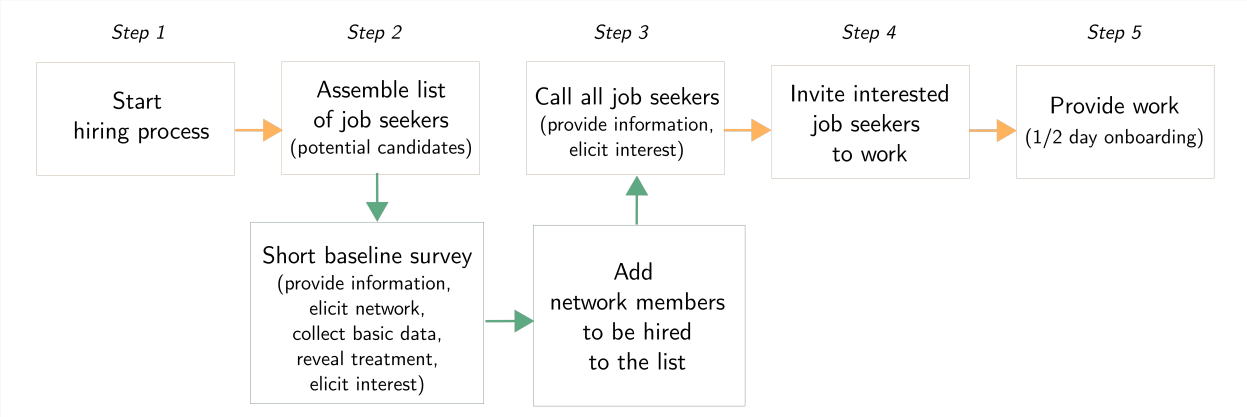
World Bank. 2020. *World Development Report 2020: Trading for Development in the Age of Global Value Chains.* Technical report. Washington, DC: The World Bank Group. (Cited on page 6).

World Bank Group. 2019. *Que la Route Soit Bonne : Améliorer la Mobilité Urbaine à Abidjan.* Report. Washington, D.C.: World Bank. http://documents.worldbank.org/curated/en/624341549322162402/Que-la-Route-Soit-Bonne-Am%7B%5C'e%7Dliorer-la-Mobilit%7B%5C'e%7D-Urbaine-%7B%5C%60a%7D-Abidjan. (Cited on page 6).

Zárate, Román D. 2022. *Spatial misallocation, informality, and transit improvements: Evidence from Mexico city.* Working Paper, Policy Research Working Papers. World Bank. (Cited on page 5).

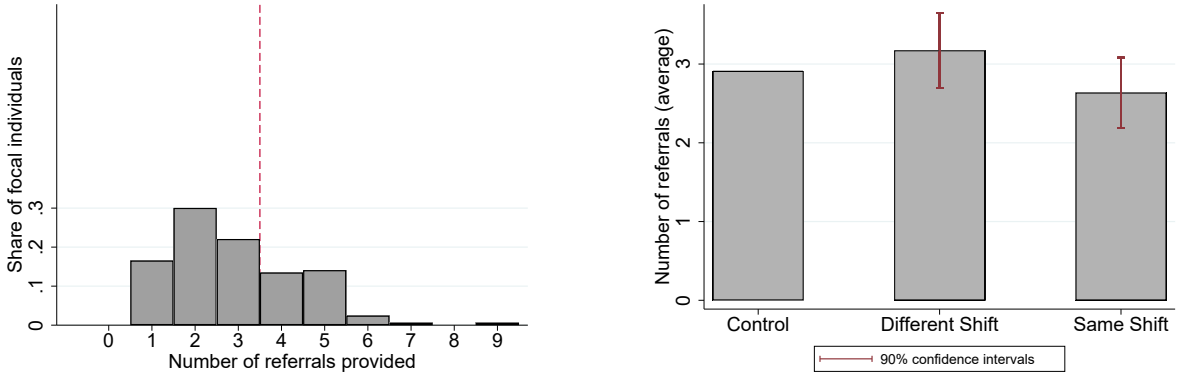
A.1 Appendix: Figures

Figure A1: Common Implementation Process



Notes: Visual description of the way in which the implementation of the experiment is embedded in the natural hiring process of the partner firms. The top row represents the natural hiring process, and the bottom row the additional activities required to implement the experiment. This process applies to both the factory and the sales jobs.

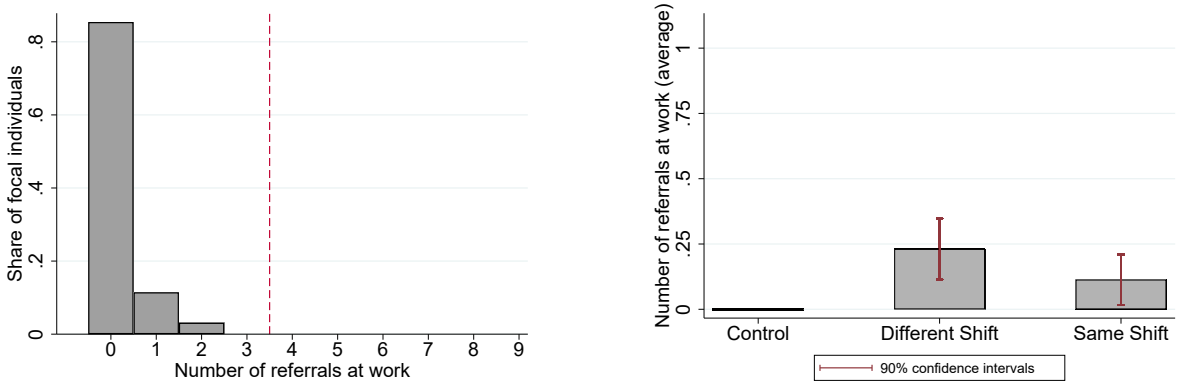
Figure A2: Size of Networks Elicited at Baseline, Factory Experiment



(A) Distribution **(B)** Balance

Notes: Distribution and average across each treatment arm of the number of referrals elicited at baseline from the focal job seekers in the factory experiment. Sample restricted to job seekers providing at least 1 referral. On the left panel, the dashed red line indicates that up to three network members are hired for each focal job seeker; any job seeker to the left of the line will have its reported network saturated with job offers. On the right panel, 90% confidence intervals for the difference with the Control group are reported. N=163 focal job seekers.

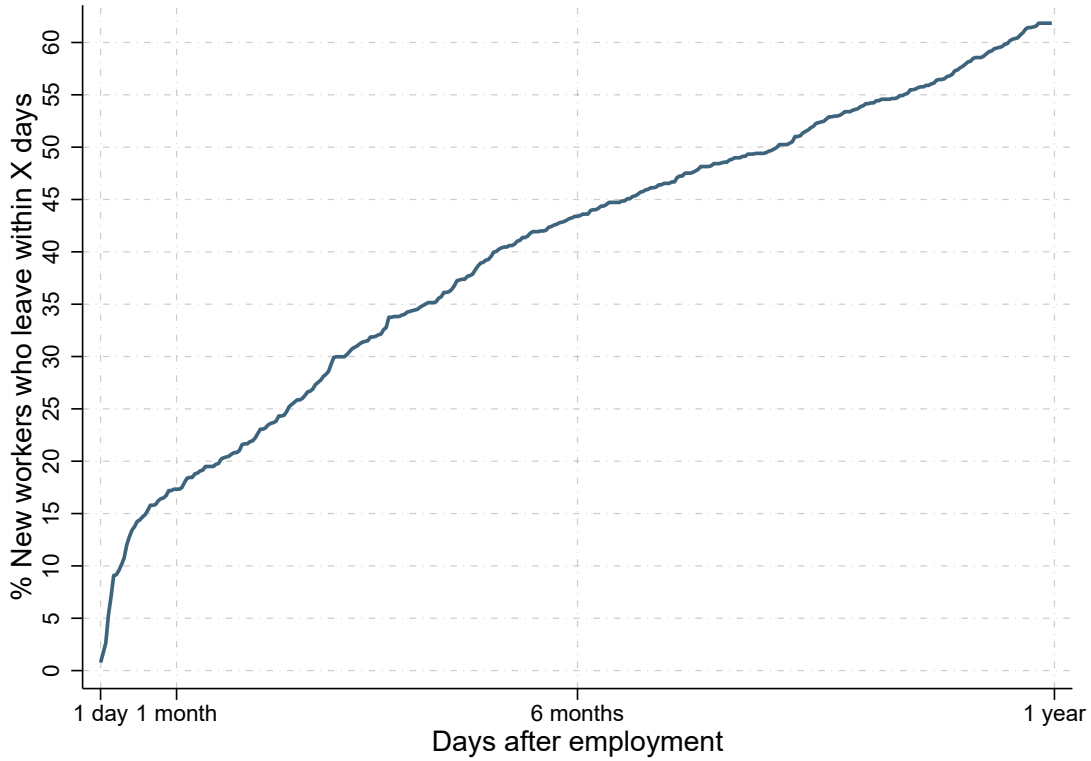
Figure A3: Number of Network Members at Work, Factory Experiment



(A) Distribution (treatment only) **(B)** Balance

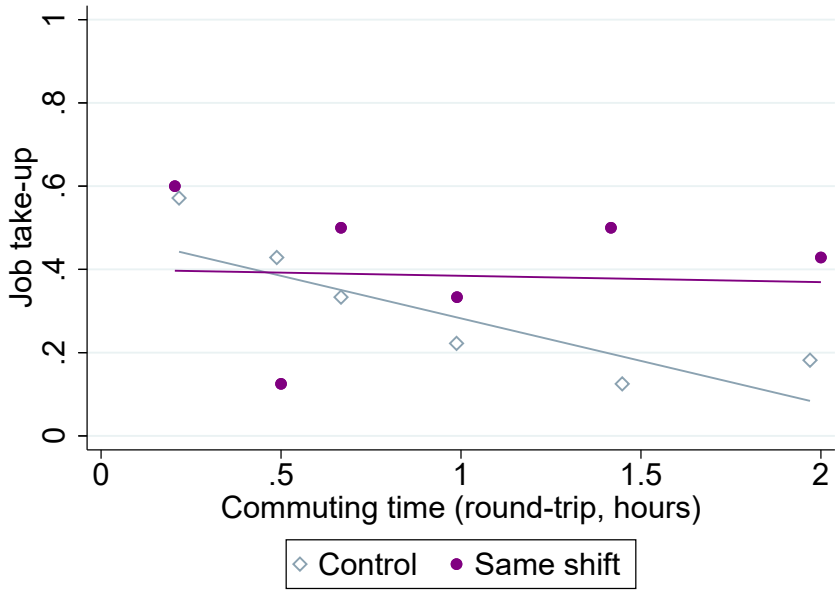
Notes: Left panel: distribution of the number of network members (referrals) who come to work at the factory at least one day, for each focal job seeker assigned to either the Same Shift or the Different Shift treatment arms. Right panel: average number of network members who come to work at the factory at least one day, separately by treatment arm. 90% confidence intervals for the difference with the Control group are reported. N=163 focal job seekers.

Figure A4: Turnover at the Factory



Notes: Share of workers who quit the factory within X days, from 1 day to 1 year. Sample of all factory workers who started work at least one year before the end of the available human resources data—including both workers from the factory experiment and non-experiment workers. N=1,431 workers.

Figure A5: Job take-up by Binned Commuting Time, Factory Experiment



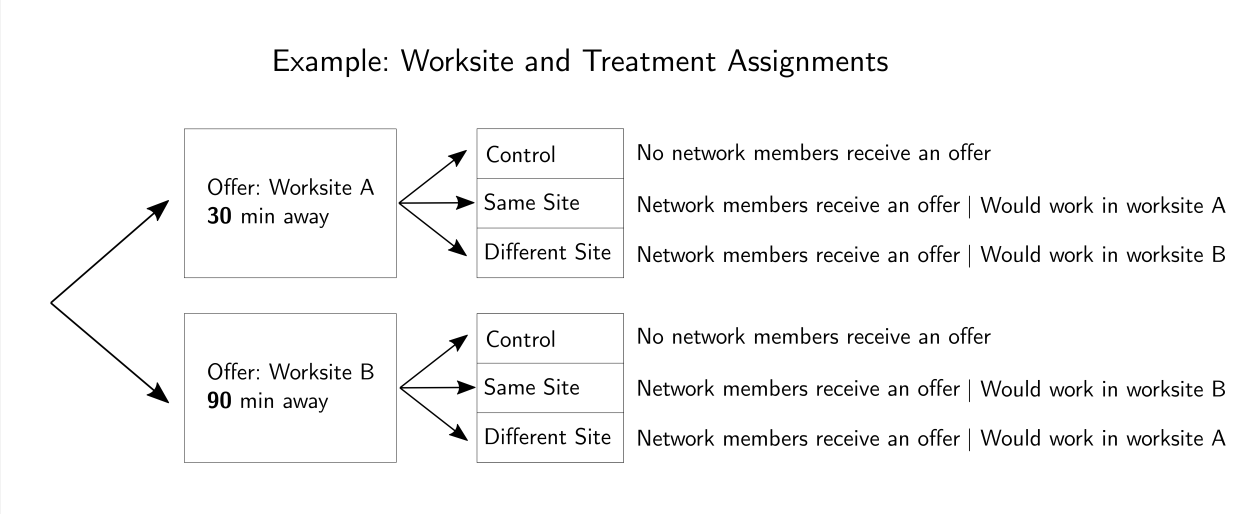
Notes: Commuting time to the factory elicited at baseline from the focal job seekers. Results are reported as a binscatter, together with a line of best fit computed over the underlying raw data. We remove three outliers with a reported commuting time of over 3 hours. N=81.

Figure A6: Absenteeism at the Factory



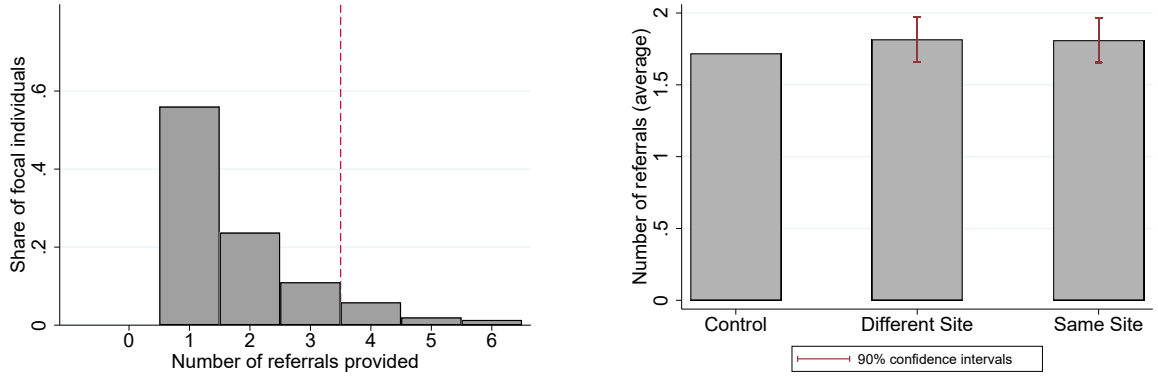
Notes: Sample of all factory workers, whether hired through the factory experiment or not. The blue dotted line represents the number of workers in the workforce, who are expected to come to work that day. (A worker is classified as being in the workforce between their first and last day at the factory.) The red line represents the number of workers actually coming to work that day. Days in the bottom 5% of the attendance distribution (where at most 64 workers are present) are excluded, to account for holidays.

Figure A7: Illustration of the Design, Sales Experiment



Notes: Graphical representation of the cross-randomization implemented in the sales experiment. Each focal job seeker is randomly assigned to either of two worksites (in this example, a focal job seeker can either be assigned to a worksite located 30 minutes or 90 minutes away from their home) and to either of three primary treatment groups (control, Same Shift, Different Shift).

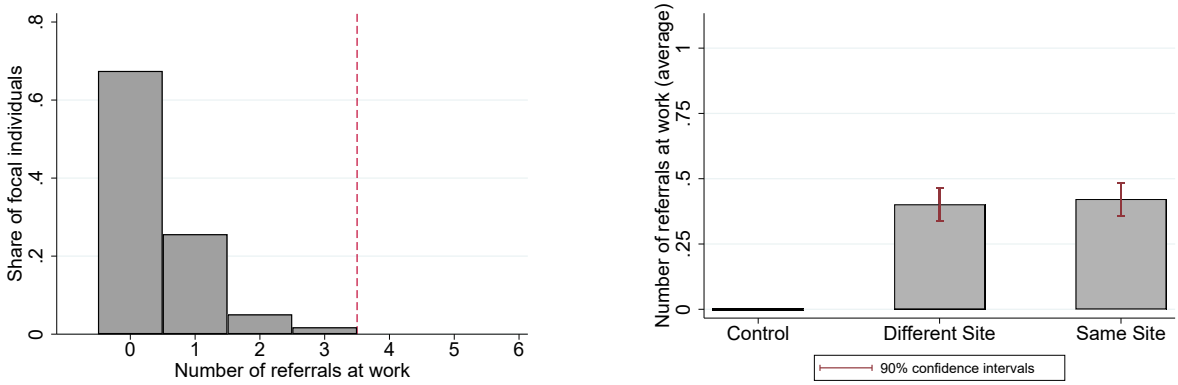
Figure A8: Size of Networks Elicited at Baseline, Sales Experiment



(A) Distribution **(B)** Balance

Notes: Distribution and average across each treatment arm of the number of referrals elicited at baseline from the focal job seekers in the sales experiment. Sample restricted to job seekers providing at least 1 referral. On the left panel, the dashed red line indicates that up to three network members are hired for each focal job seeker; any job seeker to the left of the line will have its reported network saturated with job offers. On the right panel, 90% confidence intervals for the difference with the Control group are reported. N=873 focal job seekers.

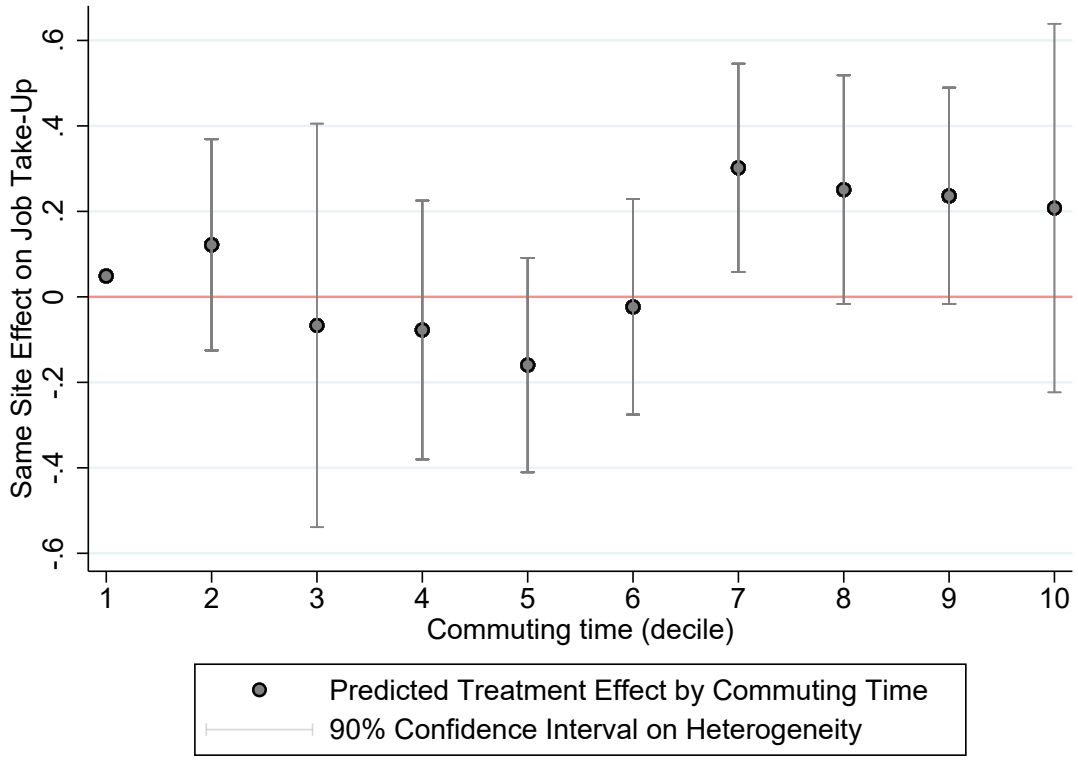
Figure A9: Number of Network Members at Work, Sales Experiment



(A) Distribution (treatment only) **(B)** Balance

Notes: Left panel: distribution of the number of network members (referrals) who come to work on the marketing campaign at least one day, for each focal job seeker assigned to either the Same Site or the Different Site treatment arms. Right panel: average number of network members who come to work on the marketing campaign at least one day, separately by treatment arm. 90% confidence intervals for the difference with the Control group are reported. N=873 focal job seekers.

Figure A10: Same Site treatment Effect by Commute Time, Sales Experiment



Notes: Effect of being assigned to the Same Site treatment arm relative to Control, separately for each decile of assigned commute time. The 90% confidence intervals represent the uncertainty associated with the heterogeneity in treatment effects in each decile relative to the first decile. The regression underlying this graph controls for the commute time to each of the two possible worksites and their interactions.

A.2 Appendix: Tables

Table A1: Baseline Balance, Factory Experiment

Variable	(1)	(2) Control	(3) Different Shift	(4)	(5) Same Shift	(6)
	N	Mean/SE	Mean/SE	P-value	Mean/SE	P-value
Panel A: Jobseeker Characteristics						
Is a woman	163	.72	.65	.472	.82	.210
		(.45)	(.48)		(.39)	
Age	163	29	28	.369	31	.090
		(6)	(6)		(5)	
Years of education	163	8	8	.588	7	.699
		(5)	(6)		(6)	
Has a partner	163	.49	.50	.936	.57	.438
		(.50)	(.50)		(.50)	
Has young children	163	.40	.21	.023	.48	.445
		(.49)	(.41)		(.51)	
Panel B: Jobseeker Baseline Income-Generating Activities						
Any IGA, past 3 months	163	.73	.73	.994	.61	.203
		(.45)	(.45)		(.49)	
Any earnings, past 30 days	163	.61	.71	.256	.52	.359
		(.49)	(.46)		(.51)	
Earnings, past 30 days (FCFA)	163	31522	39099	.349	30582	.908
		(37500)	(47893)		(44807)	
Any earnings, past 7 days	163	.46	.56	.307	.41	.581
		(.50)	(.50)		(.50)	

Notes: This table presents summary statistics for focal job seekers' baseline covariates by treatment group in the sales experiment and tests whether there are any statistically significant differences between experimental arms at baseline. Col. (1) presents the number of observations, and Cols. (2), (4), and (5) the mean and standard error of each covariate. Cols. (3) and (5) present the p-value from the t-test of equality of means between the Different Site and Same Site treatment arms, respectively, and the Control Group. Panel A includes job seekers' characteristics. Having young children is an indicator variable that equals 1 if the respondent has at least one child under their care at the time of the survey. Panel B includes job seekers' baseline income-generating activities: whether they had any in the past 3 months, as well as their earnings in the past 30 days and 7 days. Earnings in the past 30 days windsorized at the 99th percentile. Exchange rate in December 2022: 620 FCFA = 1 USD.

Table A2: Baseline Balance, Sales Experiment

Variable	(1)	(2) Control	(3) Different Site	(4)	(5) Same Site	(6)
	N	Mean/SE	Mean/SE	P-value	Mean/SE	P-value
Panel A: Jobseeker Characteristics						
Is a woman	873	.38	.40	.800	.32	.113
		(.49)	(.49)		(.47)	
Age	873	25	25	.856	25	.807
		(5)	(4)		(5)	
Years of education	873	15	15	.766	15	.383
		(2)	(2)		(2)	
Has a partner	873	.09	.11	.508	.10	.689
		(.29)	(.31)		(.30)	
Has young children	873	.20	.18	.595	.17	.370
		(.40)	(.38)		(.37)	
Panel B: Jobseeker Baseline Income-Generating Activities						
Any IGA, past 3 months	873	.62	.64	.692	.60	.527
		(.48)	(.48)		(.49)	
Any earnings, past 30 days	873	.40	.39	.776	.37	.379
		(.49)	(.49)		(.48)	
Earnings, past 30 days (FCFA)	873	42519	39790	.684	32827	.113
		(77818)	(74149)		(57974)	
Any earnings, past 7 days	873	.22	.20	.536	.19	.396
		(.42)	(.40)		(.39)	
Currently a student	873	.39	.43	.312	.42	.422
		(.49)	(.50)		(.49)	

Notes: This table presents summary statistics for focal job seekers' baseline covariates by treatment group in the sales experiment and tests whether there are any statistically significant differences between experimental arms at baseline. Col. (1) presents the number of observations, and Cols. (2), (4), and (5) the mean and standard error of each covariate. Cols. (3) and (5) present the p-value from the t-test of equality of means between the Different Site and Same Site treatment arms, respectively, and the Control Group. Panel A includes job seekers' characteristics. Having young children is an indicator variable that equals 1 if the respondent has at least one child aged below 5 at the time of the survey. Panel B includes job seekers' baseline income-generating activities: whether they had any in the past 3 months, as well as their earnings in the past 30 days and 7 days. Earnings in the past 30 days windsorized at the 99th percentile. It also includes an indicator that equals 1 if they are a student at the time of the survey. Exchange rate in December 2022: 620 FCFA = 1 USD.

A.3 Appendix: Amenity Value

We build a simple model of referral and job take-up decisions, which we primarily use to derive the estimation strategy for the amenity value of offering a job to network members with possible interactions.

A.3.1 Model set-up

The model considers two types of individuals: focal job seekers and referred individuals.

The focal job seekers make two decisions: they decide whether to refer a network member for the job, and whether to take up a job offer. The utility they can derive from these decisions is characterized by four parameters: the wage offered ω, their cost of effort e_f, their cost of making a referral c, and their value of working with the person they can refer π_f.

The referred individuals make only one decision if they have been referred and receive a job offer: whether to take up the offered job. If they do not receive a job offer, they do not make any decision. The utility they derive from their decision is characterized by three parameters: the wage offered ω, their cost of effort e_r, and their value of working with the person who referred them π_r. Consistent with our study setting, the wage offered is the same for the focal job seekers and their network members.

The model has four steps:

1. The focal job seeker learns e_f and π_f
2. The focal job seeker decides whether to refer a network member
3. The focal job seeker learns e_r and π_r
4. The focal job seeker and the referred individual decide whether to take up an offer

If the focal job seeker and the individual they refer both receive a job offer, a coordination game is played in Step 4. We allow for communication between the players and therefore study any Pareto-dominant pure strategy Nash equilibrium.

A.3.2 Equilibrium decisions

We can partition the set of possible values for the focal job seeker's cost of effort e_f into three segments, and use it to describe the equilibrium decisions of the focal job seekers.

A. If $\omega - e_f \in \left(-\infty, -\pi_f + \frac{c}{1-Pr(\omega - e_r + \pi_r < 0)}\right]$

$\rightarrow$ The focal job seeker does not make a referral (step 1), does not go to work (step 2)

B. If $\omega - e_f \in \left(-\pi_f + \frac{c}{1-Pr(\omega - e_r + \pi_r < 0)}, 0\right]$

→ The focal job seeker makes a referral (step 1), goes to work only if the referred individual goes to work (step 2)

C. If $\omega - e_f \in (0, +\infty)$

→ The focal job seeker makes a referral (step 1), goes to work (step 2)

A.3.3 Calibration

We do not need to fully calibrate the model in order to achieve the objective of estimating the amenity value of making an offer to network members who could interact with the focal job seeker.

We seek to estimate δ: the wage-equivalent of the amenity value. Specifically, it is the increase in wage that would be sufficient to achieve the same increase in job take-up among the focal job seekers in the control group as the Same Site treatment did in the sales experiment. Formally, we search for δ such that:

$$Pr\left(\omega + \Delta\omega + \delta - e_f > 0\right) - Pr\left(\omega + \Delta\omega + -e_f > 0\right) = \beta_{ITTrescaled} \tag{A1}$$

where $\beta_{ITTrescaled}$ is the treatment effect estimate from the sales experiment, re-scaled to be estimated on the full sample of initial job seekers (unconditional on referral) for comparability.

Our objective is to calibrate Equation A1 to recover δ. We know the exact values of ω (the offered wage during the pilot) and $\Delta\omega$ (the natural wage increase). We have estimated the treatment effect of being assigned to the Same Site arm relative to Control in section 4, and thereby have an estimate of $\beta_{ITTrescaled}$ (to re-scale the estimate, we simply need to multiply the ITT estimate from section 4 by the share of focal job seekers in the initial sample who provided at least one referral).

The only missing element to calibrate Equation A1 is the cumulative distribution function of the cost of effort e_f. We can calibrate any such CDF that has at most two parameters using the response to the natural wage variation. In the control group of the sales experiment, the referred individuals do not receive a job offer. As such, any focal job seeker who takes up the offered job has a cost of effort located on the segment C. If we measure the share of the focal job seekers (on the full sample, unconditional on providing a referral) in the control group who take up the offered job, we can therefore estimate $Pr\left(\omega - e_f > 0\right)$ (before the natural wage increase) and $Pr\left(\omega + \Delta\omega - e_f > 0\right)$ (after the natural wage increase). These two estimates sufficient for our purpose.

If we assume a uniform CDF on the cost of effort, Equation A1 becomes:

$$\delta = \beta_{ITTrescaled} \frac{\Delta\omega}{Pr\left(\omega + \Delta\omega - e_f > 0\right) - Pr\left(\omega - e_f > 0\right)} \tag{A2}$$

And we can estimate $\hat{\delta} = 0.19$

If we instead assume a normal CDF on the cost of effort, Equation A1 becomes:

$$\delta = \sigma \left(\Phi^{-1} \left(Pr \left(\omega + \Delta\omega - e_f > 0 \right) + \beta_{ITTrescaled} \right) - \Phi^{-1} \left(Pr \left(\omega + \Delta\omega - e_f > 0 \right) \right) \right) \quad \text{(A3)}$$

with

$$\sigma = \frac{\Delta\omega}{\Phi^{-1} \left(Pr \left(\omega + \Delta\omega - e_f > 0 \right) \right) - \Phi^{-1} \left(Pr \left(\omega - e_f > 0 \right) \right)}$$

And we can estimate $\hat{\delta} = 0.13$