Along for the Ride: The Labor Supply Effects of Commuting Together

Aletheia Donald and Florian Grosset-Touba*

November 4, 2026

Abstract

Commuting is notoriously challenging in lower-income country cities, fragmenting urban labor markets and eroding potential gains from agglomeration. We study whether jobseekers' ability to travel to work with people they know reduces this friction and increases labor supply. In two field experiments in urban Côte d'Ivoire, we randomize whether a jobseeker's network members also receive job offers and whether they are assigned to work at the jobseeker's own site or shift. The first, with a marketing firm, also randomizes job location and hence commute time. Working alongside network members raises take-up by 12 percentage points, while working at a different site has no effect. The effect rises with commute length, without detectable productivity losses. In the second, with a factory, take-up of long-term formal jobs rises by 21 percentage points and more workers remain employed nine months later. Factory workers who commute together also miss work on the same days. Commuting together is a valuable job amenity, which may help explain why firms hire through referrals and cluster within cities.

*Donald: World Bank (adonald@worldbank.org); Grosset-Touba: CREST, ENSAE, Institut Polytechnique de Paris & CEPR (Corresponding author: florian.grosset-touba@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, Emily Breza, Hannah Farkas, Rob Garlick, Ed Glaeser, Louise Guillouet, Gabriel Kreindler, Suresh Naidu, Anna Papp, Tommaso Porzio, Imran Rasul, Jeff Shrader, Jack Willis, 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 Poverty 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 and Peirong Shi for excellent research assistance. AI tools were used to assist with writing code and with searching for relevant external data and references, all of which were checked by the authors. 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

Cities foster growth by bringing firms and workers together. Commuting costs are one of the main forces limiting these benefits. As cities grow, trips to work become longer and more congested, offsetting part of the gains from agglomeration (e.g., Duranton and Puga 2020). Much of this cost is not monetary, with commuting ranking among the least enjoyable activities of the day (Kahneman and Krueger 2006). In the cities of lower-income countries, these costs are especially high. Most workers travel in crowded, informal minibuses and shared taxis, on trips that are slow, tedious and often unsafe. As a result, firms struggle to hire and retain workers who live further away, while many residents limit where they work, shop, or study to areas closer to home.[1] The experience of the trip, however, depends on who one travels with—commuting with people one knows, such as friends or neighbors, can make it safer and more pleasant. This suggests a link between social networks and labor supply decisions: people's desire to work at a given firm may depend on whether they can travel there with their network members.

Both costly commutes and the habit of commuting together are common in Abidjan, the economic capital of Côte d'Ivoire. Among workers we survey at a large factory, the round trip to work takes over an hour and commuting together is common: 39% of workers travel to work with at least one co-worker (40% of women and 33% of men). Asked why, both men and women most often cite safety and a more pleasant trip. Personnel records indicate that workers who commute together coordinate their days of presence and absence, and also tend to quit the factory at the same time. Together, these facts suggest that workers value commuting together and time their work around it. This social nature of commuting motivates the question we study: does the ability to commute with others from one's social network increase labor supply?

We develop a causal test to empirically examine this possibility.[2] We consider a set of jobseekers—whom we refer to as "focal" jobseekers—deciding whether to accept a job offer, and elicit their networks at baseline. We randomly vary whether members of these networks receive a job offer at the same firm and whether they are assigned to work at the same time and place as their focal jobseekers, thereby inducing random variation in their ability to commute together. In addition, we randomly vary the focal jobseekers' commuting time

1. For academic and policy reviews on transport in cities of lower-income countries, see Kumar and Barrett 2008; Lall et al. 2017; Bryan et al. 2020; Storeygard 2026. On how it influences where people shop or study, see Borker 2021; Vitali 2023.

2. We take home locations as given. Over the time horizons we study, relocating is costly: formal jobs are concentrated near the port and in high-rent areas, far from where most workers live (**WorldBank2019**), and moving closer to work would often mean leaving the neighbors one could commute with.

from home to work. Varying both dimensions lets us separate the role of commuting together from other reasons that job offers to network members might matter. If jobseekers value the ability to commute together, offers to their network members should raise take-up more the longer the focal jobseeker's assigned commuting time, but only if they can actually travel together. We test for effects on job take-up, and also consider retention, attendance, and productivity.

We implement this design across two field experiments in Abidjan. The first is with a marketing firm hiring 848 focal jobseekers for short-term street canvassing jobs. Work takes place at either of two worksites, in different boroughs of the city. At baseline, each focal jobseeker refers people for similar positions at the same firm; these referrals form their network members. Most refer friends, and most name people they could commute with. We randomly assign each focal jobseeker to Control (no network member receives an offer), Same Site (up to three network members receive an offer to work at the focal jobseeker's site) or Different Site (up to three network members receive an offer to work at the other site). Crucially, we also randomize each focal jobseeker's own worksite between the two options, which induces exogenous variation in their commuting time.

We find strong effects of the ability to commute together on labor supply to the firm. Among our main sample of 677 focal jobseekers, whose offered network members could all commute with them, those assigned to Same Site are 12.4 percentage points more likely to take up the job than those in Control, a 38% increase ($p = 0.008$), and work 32% more days ($p = 0.024$). In the full sample of 848, which includes jobseekers who referred people they could not commute with, the effect is smaller (7.1 percentage points, $p = 0.091$). Focal jobseekers assigned to Different Site are no more likely to take up the job than those in Control, and significantly less likely than those assigned to Same Site ($p = 0.033$). Offering jobs to network members matters mainly when they work at the focal jobseeker's site. Consistent with a role for commuting together, the Same Site effect rises with the randomly assigned commuting time, by 7.4 percentage points per round-trip hour ($p = 0.04$). Some focal jobseekers assigned to Different Site could still share part of the trip with their network members, because both worksites lie in the same direction from their home. Among them, the treatment effect also rises with commuting time. It does not among those whose routes diverge, though the difference between the two groups is not statistically significant.

Offering jobs to network members could also affect labor supply through channels other than commuting together, for instance by changing redistributive pressure, beliefs about the firm, or opportunities to interact during breaks. We design the experiment to directly assess the relevance of these channels. Comparing Different Site to Control isolates the

effect of offering jobs to network members when interactions during the commute and the workday are limited. Comparing Same Site to Different Site among focal jobseekers with short, randomly assigned commutes isolates the effect of interacting at the worksite. Neither comparison yields a significant difference in take-up, suggesting these channels play a limited role in our setting. The Same Site effect therefore largely reflects the value jobseekers place on commuting together. We express this value in wage terms, leveraging a wage increase between the pilot and the main experiment that lets us estimate the elasticity of take-up to the wage. To raise take-up as much as Same Site assignment does, the firm would need to increase the offered wage to individuals hired alone by 9 to 23%. The ability to commute together thus represents a sizable non-wage job amenity.

The value to the firm of offering this amenity depends on its productivity consequences, which are ambiguous ex ante. Workers hired alongside friends could shirk together or learn from each other (Mas and Moretti 2009; Bandiera et al. 2010; Park 2019). The marginal focal jobseekers drawn in by the treatment, and the network members hired with them, could be more or less productive than those who would have taken the job anyway (Beaman and Magruder 2012; Burks et al. 2015; Srivastava 2026). Using detailed productivity data from two independent sources, we find no statistically significant differences in productivity across treatment arms, or between focal jobseekers and network members. For a firm, hiring focal jobseekers with long commutes alongside network members they can commute with raises average productivity by 0.02 standard deviations, which is neither economically nor statistically significant. As such, the labor supply increases do not come at a productivity cost to the firm.

What makes commuting together valuable? Monetary cost-sharing is not a driver: the informal minibuses and shared taxis most workers use charge per passenger, and the Same Site effect varies with commuting time but not with commuting cost. Gendered attitudes toward commuting are not the reason either. Though they may still constrain labor supply in our setting, the Same Site effect is similar for men and women. Safety explains part of the effect. In Control, jobseekers who feel less safe coming home at night are less likely to take up the job, especially the 49% who were not already commuting to work at baseline. Among them, the Same Site effect is concentrated among those who do not feel very safe, for both men and women. Among the 51% who were already commuting, the effect does not vary with perceived safety, pointing to a more pleasant trip as the likely explanation. This is consistent with whom jobseekers refer: when they name someone they could commute with, it is someone they feel very close to in 71% of cases.

We conduct a second field experiment with a multinational company operating an agri-

processing factory, offering long-term jobs to 163 focal jobseekers in Abidjan. This allows us to test whether our main results hold for a different job and population, mostly women with limited schooling, and to assess the persistence of the effects. We implement the same three-arm design, built on the factory's two non-overlapping shifts in place of the two worksites. The results closely match those of the sales experiment. Focal jobseekers assigned to Same Shift are 20.6 percentage points more likely to take up the factory job than those in Control ($p = 0.030$). Conditional on take-up, attendance is similar, and retention is 96.5 days longer, though imprecisely estimated ($p = 0.10$). Together with higher take-up, this raises total days worked by 133% ($p = 0.012$), and focal jobseekers assigned to Same Shift remain 12.9 percentage points more likely to be employed at the factory nine months later ($p = 0.074$). As in the first experiment, assignment to Different Shift has no effect on take-up, attendance, or retention.

Our study advances our understanding of urban transport and its impacts on economic activities. A growing body of structural work shows that transport infrastructure projects have large effects on the spatial distribution of people, jobs and commuting flows in cities of lower-income countries, and studies how best to design such projects (e.g., Zárate 2022; Kreindler et al. 2026; Tsivanidis 2026; Velasquez-Cabrera 2026). In partial equilibrium, reduced-form studies in Ethiopia and South Africa find that temporary transport subsidies stimulate job search, with mixed downstream consequences for employment (Franklin 2018; Abebe et al. 2021; Banerjee and Sequeira 2023; Caria et al. 2023). Most of this work treats commuting as an individual decision and focuses on lowering its time or monetary cost. We show that commuting is also a social activity. Whether jobseekers can travel with their network members is key to understanding labor supply decisions—including job take-up, attendance, and retention.

Our results also complement experiments in settings with highly conservative gender attitudes, such as India or Pakistan. There, offering women work-from-home jobs or women-only or free transport raised their labor supply (Dasgupta and Datta 2023; Jalota and Ho 2024; Garlick et al. 2025; Ho et al. 2026), though not always in the short run (Chen et al. 2026). We study a less gender-restrictive environment, closer to the experience of women across urban sub-Saharan Africa. In this setting, safety concerns also play a role in shaping the value of commuting together with network members. But safety matters for men as well as women, and is not the only mechanism: making long commutes more pleasant is also a key benefit.

Our results can help explain firm hiring behavior, particularly the extensive use of referrals. A large literature has studied this topic through the lens of information asymmetries: hiring

through referrals can reduce search, screening, and monitoring costs (see Beaman 2016; Topa 2019; Afridi and Dhillon 2022, for reviews). On the worker side, recent research shows that workers 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 2025). We point to a further motive: jobseekers value being able to commute together to work with people they know, especially when facing long commutes. Hiring network members can therefore represent an important non-wage job amenity, enabling firms to increase the labor supply they face at prevailing wages.

More broadly, our results shed light on how social ties shape the functioning of labor markets in lower-income countries. Social ties can sustain norms that restrict labor supply or the sharing of opportunities, from wage floors in spot labor markets (Breza et al. 2025) to redistributive pressure (Carranza et al. 2025) and knowledge hoarding (Cefalà et al. 2025). We show that social ties can also raise labor supply: being able to commute with network members makes people more willing to take a job and to keep it.[3]

These roles of commute and social ties in shaping work decisions are not unique to lower-income countries. In the United States and Europe, commuting is one of the job attributes that workers weigh most when searching for jobs and supplying labor (Mas and Pallais 2017; Marinescu and Rathelot 2018; Le Barbanchon et al. 2021; Dube et al. 2022). Transport infrastructure upgrades and transport subsidies also have major economic effects there (Duranton and Turner 2012; Phillips 2014; Ahlfeldt et al. 2015; Heblich et al. 2020; Severen 2023; Chizeck and Mbonu 2025). In addition, spouses and relatives coordinate the timing of their work to be able to spend leisure time together (Hamermesh 2002; Hallberg 2003; Goux et al. 2014; Georges-Kot et al. 2017; 2024; Moghadam et al. 2024). Our results show that such coordination extends to the trip to work. Yet the value of commuting together may well be smaller in high-income cities. Travel there is generally faster (Couture et al. 2018; Akbar et al. 2023a; 2023b), and commuting alone by car is more pleasant than commuting alone in a crowded minibus. Social ties beyond spouses and close relatives may also carry less economic value. Whether being able to commute together would raise labor supply as much in a high-income city remains an open question.

In what follows, we present the setting (Section 2), then the sales experiment (Section 3) and the factory experiment (Section 4). Section 5 concludes.

3. In contrast, job search assistance programs can reduce the employment of the beneficiaries' network members by changing information sharing and job search assistance within networks (Caria et al. 2023) or re-allocating work among household members (Afridi et al. 2025).

2 Motivation

2.1 Commute and Work in Abidjan

We work in Abidjan, the economic capital of Côte d'Ivoire. 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 2019). These characteristics are common across many cities in lower-income countries.[4] Workers commute to these jobs from home each day, as in most labor markets. The export-oriented manufacturing of parts of Asia, where firms house migrant workers on site, tend to be an exception (Pun and Smith 2007; Goodburn and Mishra 2024). This way to reduce commuting costs requires both land availability and workers willing to live in dorms with co-workers.

Workers reach these jobs on foot and, for longer distances, on the informal minibusses (*gbakas*) and shared taxis (*woro-woros*) that are used for most public-transport trips in the city. Gbakas seat 18 to 32 passengers and woro-woros four or five, and between them they carry close to two-thirds of public-transport trips in Abidjan (République de Côte d'Ivoire 2015; World Bank 2018a). These trips are slow, crowded, and require navigating transfers, waiting, and walking portions that are tedious and, at off-peak hours, unsafe. Traveling with a friend or acquaintance is one of the few ways to make the trip more pleasant. Fares are charged per passenger and do not fall when riders travel together, so any value of commuting with others does not stem from sharing the monetary cost of the trip. Both modes operate on a fill-and-go basis, and the routes served and the number of vehicles on each are set by the operators' associations rather than by individual drivers or riders—features common to minibus systems across the continent (Conwell 2026; Mbonu and Eaglin 2026).

In this setting, urban unemployment is high and labor force participation limited at 68%, comparable to the sub-Saharan African norm (Christiaensen and Premand 2017).[5] Nonetheless, firms often report difficulties in hiring and retaining blue-collar workers at prevailing

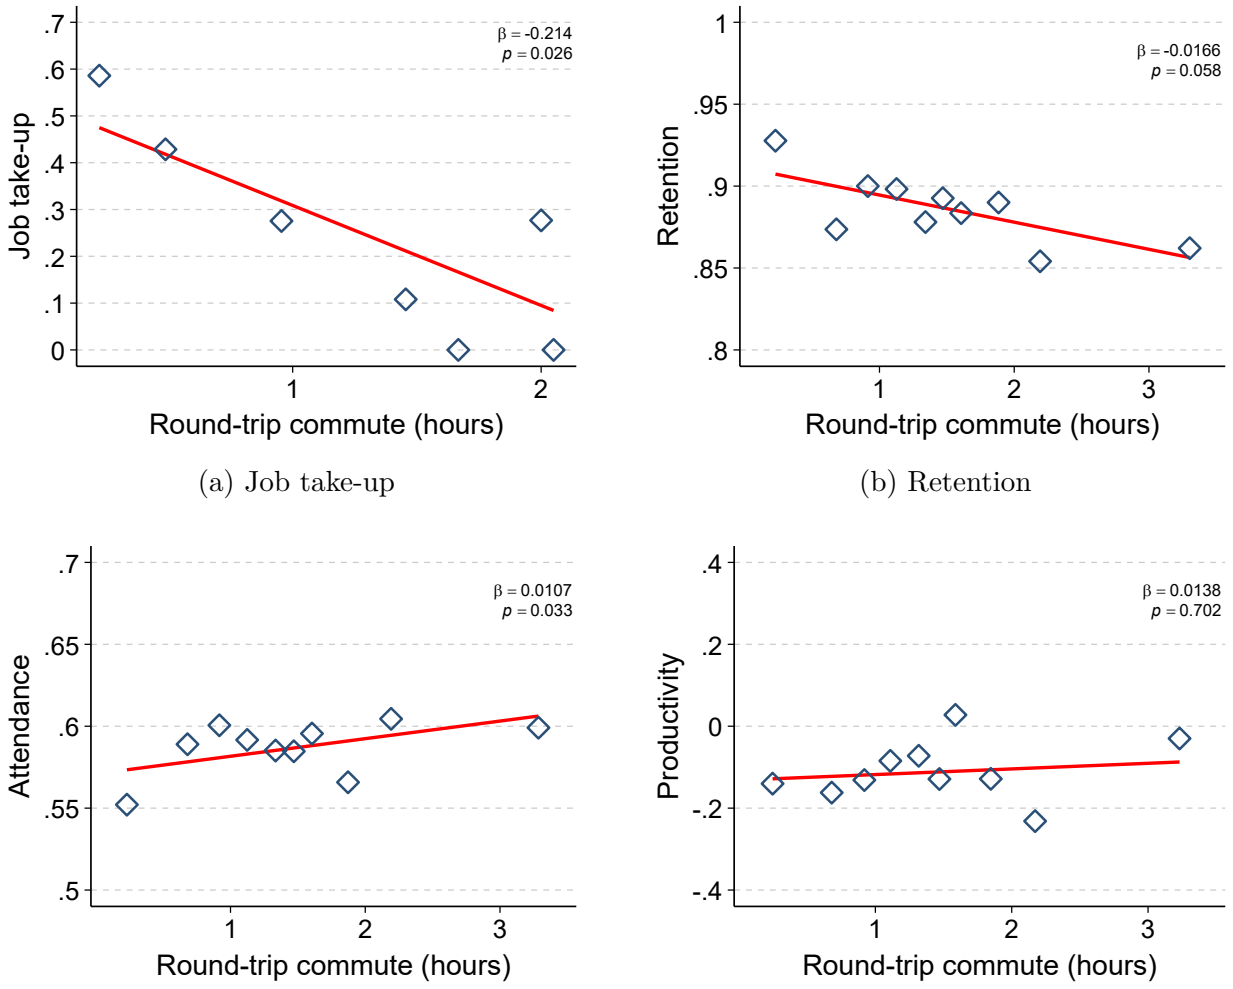
4. For instance, in Morocco, industry executives report unreliable transport services as a source of high employee turnover, and transportation to be a major concern for employees and labor union representatives (Minster et al. 2022).

5. Male participation is also in line with South Asia: the one notable difference is a far smaller gender gap, as women participate at 80% of the male rate against 43% across South Asia (World Bank 2025). While gender gaps in earnings exist, barriers to women working outside the household are limited: three-quarters of Ivorians report that women are rarely or never prevented by their spouses or other family members from taking paid employment (Afrobarometer 2024).

wages (World Bank 2018b; 2020). This is also the case for our partner companies, whose jobs pay well relative to workers' available alternatives—roughly two times typical outside earnings—yet prove difficult to fill. For instance, at our partner factory, only 28% of the individuals who reported being interested in a production job and received a job offer ever came to work, and 42% 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. 2021; 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).

These hiring and retention difficulties can be directly related to commuting costs. From our experiment with a large factory, individuals offered a job (alone, without network members) are significantly less likely to take up the offered job when they live farther away from work (Figure 1, Panel (a)). In addition, when merging factory personnel records with a separate survey of workers, enabling us to know where workers live and how long they commute to go to work, we find that individuals are significantly less likely to retain their job when they live farther away (Figure 1, Panel (b)). By contrast, attendance is, if anything, slightly higher among workers with longer commutes, and productivity is unrelated to commuting time (Figure 1, Panels (c) and (d)). This could credibly reflect selection: the workers who decide to take up the job and retain it despite long commutes are likely to be more motivated.

Figure 1: Labor Supply and Commuting Time



(a) Job take-up

(b) Retention

(c) Attendance

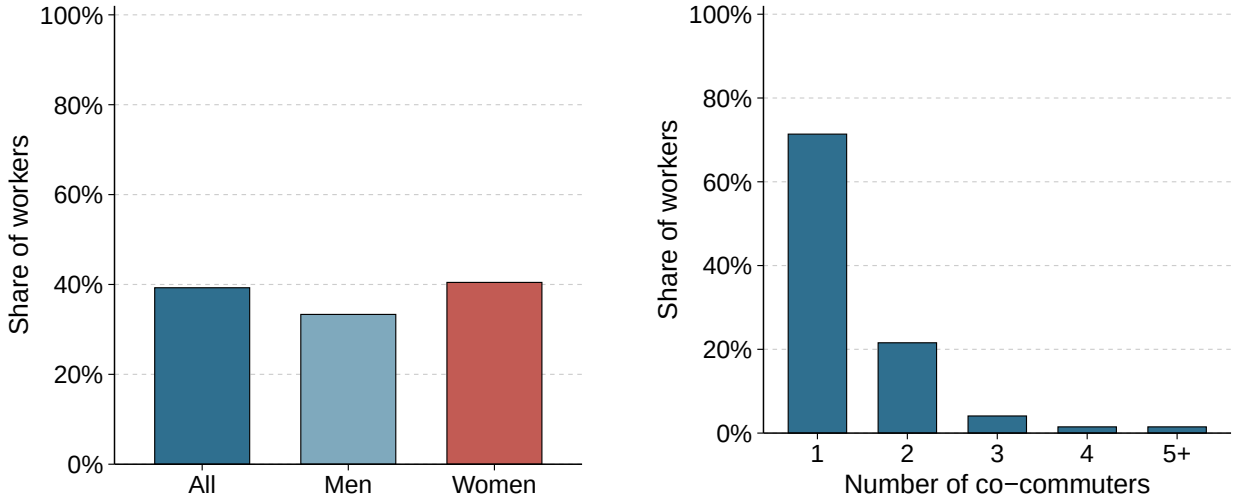
(d) Productivity

Notes: Correlation between labor supply decisions and commuting time. Binscatter of the data, together with OLS estimate, p-value and fitted line from regressing each labor supply dimension on commuting time. Appendix Table A1 reports the corresponding regression results. Evidence from workers at a large factory in Abidjan, Côte d'Ivoire. Panel (a): job take-up decisions from the control group of the factory experiment (presented in Section 4). Panels (b)–(d): retention, attendance, and productivity from non-experiment factory workers (sample and data presented in Appendix B.1). Retention is the share of weeks since hiring in which the worker remained employed. Attendance is the share of working days the worker came to work, between their first and last day on the job. Productivity is the worker's output per day worked, standardized and averaged over their working days.

2.2 Commuting Together and Labor Supply

Commuting together is a common practice in this setting. In a network survey of 685 workers at our partner factory, 39% report commuting with a co-worker. Those who do usually commute with just one other person (Panel B). These links are persistent: in 37% of these cases, the co-commuter is someone the worker already knew before starting at the factory rather than someone met on the job. In addition, the joint commute is often made with people the workers enjoy the company of: about 40% are with someone the worker also sees outside work. Consistently, when asked to rank their reasons for commuting together, over a third of the respondents place commuting becoming more pleasant first. Safety is also a major concern, with over half of the respondents placing it first. Almost none rank co-commuting making it cheaper as a benefit, which is consistent with the fixed price per head paid by workers (Figure 3).

Figure 2: Prevalence of Co-Commuting



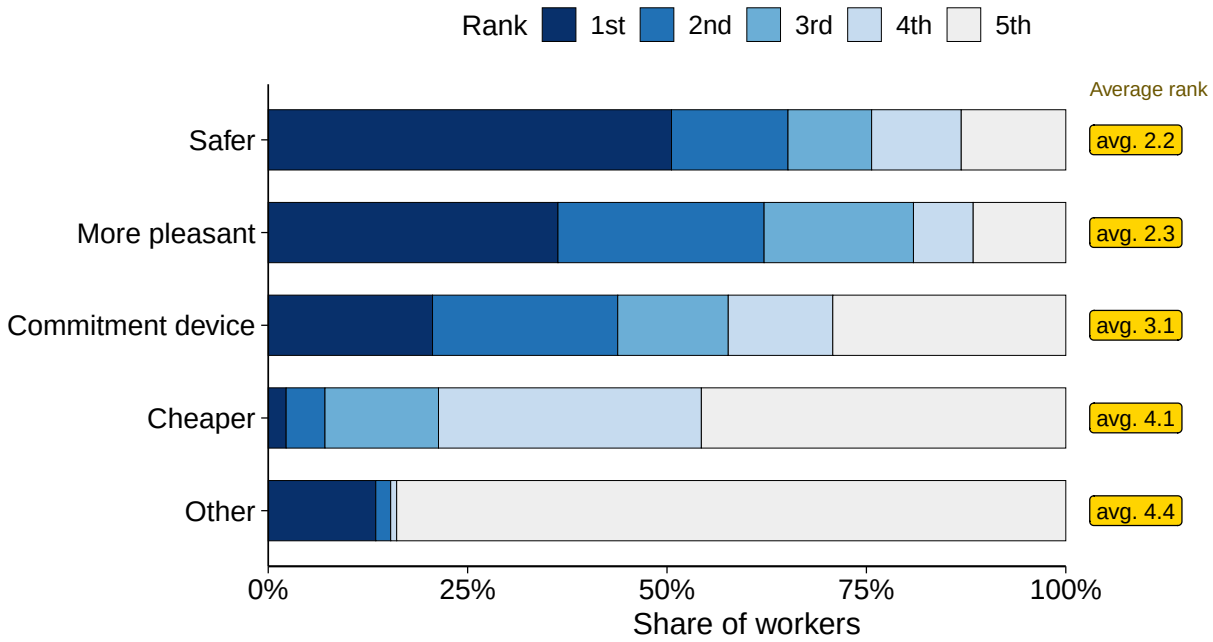
(a) Share with at least one co-commuter

(b) Number of co-commuters

Notes: Data from a survey of 685 workers at the partner factory. Panel (a): share of respondents who report at least one co-commuter—a co-worker they commuted to work with on at least three different days—overall and by gender. Panel (b): distribution of the number of co-commuters among the 269 respondents who report at least one.

Both men and women often commute with co-workers, and provide broadly similar motives for doing so. The share commuting with a co-worker is somewhat higher for women than for men (40% versus 33%, Panel A of Figure 2). For both women and men, *safer* and *more pleasant* are the two reasons most often ranked first. Women rank *safer* first more often than men (55% versus 24%), while men rank *more pleasant* first slightly more often (39% versus 36%; Appendix Figure A7). The survey alone cannot tell whether these gendered differences reflect how women and men report their reasons or genuine differences in what they value,

Figure 3: Reasons for Commuting Together



Notes: Data from the respondents of Figure 2 who reported at least one co-commuter. They were asked to rank five potential benefits from commuting together ("other" being an option), from most (1st) to least (5th) important. Each bar shows the share of respondents giving the reason each rank; the box gives the reason's average rank. Respondents could give the same rank to more than one reason.

nor whether they translate into different complementarities in labor supply. We test both directly in the sales experiment (Section 3.6).

These patterns of joint commute at the factory have direct implications for labor supply. We find that individuals choose the timing of their labor supply decisions—when they are present or absent to work, and when they quit—to coordinate with their co-commuters.

To study the role of co-commuting in shaping labor supply decisions, we match the network survey to the factory's daily attendance records. We regress a worker's absence on a given day on the number of their co-commuters absent that day, following Equation (1).

$$\text{Absent}_{it} = \beta\, A_{it} + \lambda_{\ell(i,t)} + \mu_t + \varepsilon_{it}, \tag{1}$$

where Absent_{it} equals one if worker i is absent on day t, and $A_{it} = N_{it} - P_{it}$ is the number of i's co-commuters absent that day: N_{it} is the number of i's co-commuters who work at the factory on day t (between their first and last day there) and P_{it} the number of them present. The sample is the days within i's own time at the factory on which $N_{it} > 0$, and standard errors are clustered by worker.

The link fixed effect $\lambda_{\ell(i,t)}$ is specific to worker i and to which of i's co-commuters work at

the factory on day t: a new link starts whenever one of them joins or leaves the factory. It ensures that the estimates on attendance are not driven by co-commuters joining or leaving the factory. Because each link belongs to a single worker, it also absorbs differences in absence rates across workers, so that the estimate compares days of the same worker with the same co-commuters. The day fixed effects μ_t absorb seasonal variation in attendance. To absorb shocks that could jointly move the labor supply of workers who travel together, Columns (2)–(4) of Table 1 progressively add two sets of fixed effects. Day-by-neighborhood fixed effects $\eta_{n(i)t}$, where $n(i)$ is the neighborhood the worker reports living in, absorb shocks such as weather or a local transport disruption. Day-by-production-unit fixed effects $\delta_{u(i,t)t}$, where $u(i,t)$ is the worker's main production unit in the month of day t, absorb shocks such as a managerial decision to run a unit below capacity. For the association to reflect something other than coordination, a shock would have to move co-commuters' attendance together without moving that of others in the same neighborhood or production unit on the same day.

Note that 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 read the association as the magnitude of a causal effect. We test only for its presence, which would be zero absent any coordination.

We find that a worker is 31 percentage points more likely to be absent on a day when a co-commuter is absent than on a day when none is, against a baseline absence rate of 13% (Table 1, Col. 1). Adding day-by-neighborhood fixed effects barely changes the estimate. It falls to about 12 percentage points once day-by-production-unit fixed effects are also included, but remains large and highly significant (Col. 4). This indicates that common shocks explain part of the association: the estimate falls by more than 60% once workers in the same production unit on the same day are compared. But they do not explain all of it. Even then, a co-commuter's absence raises the worker's probability of being absent by nearly as much as the baseline absence rate, which points to co-commuters coordinating their attendance rather than only sharing the same shocks.

Coordination is also contemporaneous: a co-commuter's absence predicts the worker's absence on the same day, with no detectable lead or lag effects (Appendix Figure A6), consistent with day-to-day coordination rather than slow-moving correlated trends.

Coordination in our setting is concentrated among co-commuters who are also friends: a friend's absence raises the worker's own-day absence probability by 16 percentage points, against 9 for a co-commuter the worker does not otherwise see outside work ($p < 0.05$,

Table 1: Coordination of Absences and Quits Among Co-Commuters at the Factory

	(1)	(2)	(3)	(4)
Panel A. Absences (daily)				
Nbr Co-Commuters absent	0.307***	0.288***	0.145***	0.120***
	(0.026)	(0.034)	(0.020)	(0.025)
	[0.000]	[0.000]	[0.000]	[0.000]
N: Worker-days	66,936	47,378	65,235	45,544
N: Workers	239	165	238	162
Panel B. Quits (weekly)				
Nbr Co-Commuters quitting	0.153***	0.153***	0.122***	0.103*
	(0.039)	(0.052)	(0.039)	(0.054)
	[0.000]	[0.004]	[0.002]	[0.057]
N: Worker-weeks	8,423	6,007	8,243	5,818
N: Workers	204	144	203	142
Day (week) × neighborhood FE	No	Yes	No	Yes
Day (week) × prod. unit FE	No	No	Yes	Yes

Notes: Sample of factory workers, administered a survey to elicit the workers they commute with and matched to the firm's personnel records. Panel A regresses whether the worker is absent on a given day on the number of her co-commuters absent that day, with day and link fixed effects (one per worker and set of co-commuters working at the factory; a new link starts whenever a co-commuter joins or leaves). Panel B regresses whether the worker's quit in a given week on the number of her co-commuters quitting that week, with week and worker fixed effects. Columns (2)–(4) add day (week) × neighborhood and day (week) × production-unit fixed effects to absorb common shocks. Standard errors clustered at the respondent level in parentheses; p-values in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table A3, Panel A, Col. 4), pointing to the importance of the social relationship rather than mere co-presence on the same route. Panel B reinforces that the relevant margin is friendship, not length of acquaintance: co-commuters who knew each other before the factory coordinate no differently from those who met on the job.

These effects extend beyond attendance only: co-commuters also leave the factory together. Quits are harder to study than absences. They are rare: 63 of the 204 workers in Panel B of Table 1 leave the factory during the period we study. They are also observed with a lag. We consider that a worker has left the factory once they have not appeared in the records for 90 days, so a last day close to the end of the data may simply be that of a worker still employed; for 481 of the 685 respondents, the last day in the records falls in the final 90 days. We therefore end the quit sample on 24 April 2022, 90 days before the end of the data, so that every departure in it is observed. We also study quits by week rather than by day: a quit happens once, and co-commuters who leave together need not leave on the same day. Only 8 quits fall on the same day as a co-commuter's quit, against 16 in the same week. The

outcome is whether the worker's last day falls in the week, and the regressor is the number of co-commuters whose last day falls in that week. Because a co-commuter's quit changes the set of active co-commuters, and hence the link, we replace link fixed effects with worker fixed effects.

The effects on quits mirror those on absences, pointing to coordination in labor supply decisions. A worker is 15 percentage points more likely to quit in a week in which a co-commuter quits, against a baseline weekly quit rate of 0.6% (Table 1, Panel B, Col. 1). The estimate falls to 10 percentage points but remains significant once both sets of fixed effects are included (Col. 4).

Building on these patterns, we design and implement two experiments to provide causal evidence on how the ability to commute together affects whether and how much people work.

3 Sales Experiment

Our first experiment tests whether the ability to commute with others from one's social network changes labor supply. We implement it in street canvassing, a sales job in the service sector—one absorbing a large and growing share of employment across sub-Saharan Africa (Rodrik 2016; Gollin et al. 2016). We present the setting (Section 3.1) and design of this experiment (Section 3.2), as well as the main results on job take-up decisions (Section 3.3). We estimate the productivity consequences for the firm of hiring network members who can commute together (Section 3.4), and the monetary value of this non-wage job amenity for workers (Section 3.5). Finally, we provide evidence on the reasons why workers value commuting together (Section 3.6).

3.1 Setting

Our partner is Coin Afrique, the largest digital classified-ads company in francophone Africa, which ran a street-canvassing campaign in Abidjan to grow its app's user base: workers approached people in public, explained the app, and helped them install it and sign up. The work requires a smartphone and some technical abilities, so the applicant pool skews toward university graduates and is gender-mixed, letting us study men and women in the same setting. Pay is 9,500 FCFA (∼15 USD) per day and the jobs run four to five days.[6] Each week, workers are dispatched to work at the same time in two different worksites, in

6. The campaign ran in two identical waves. The first offered five days of work, the second four, for budget reasons.

different boroughs of the city and far enough apart that workers assigned to different sites cannot interact during the day and their ability to commute together is reduced.

The sample is 848 eligible focal jobseekers, balanced across arms (Appendix Table A6).[7] They are young (mean age 25), largely university-educated, and weakly attached to work—about 40% had any earnings in the past 30 days—so the offered job is desirable relative to their outside option.

The network members are referrals. At the hiring stage, before any treatment is assigned, the baseline survey asks each focal jobseeker whether they know other people who would be suitable and available for a similar position at the same firm. To maximize the chance of eliciting individuals with whom the focal jobseeker could commute, the survey asks first about people living in their own neighborhood, then about people living elsewhere. To limit strategic referrals, focal jobseekers are told that the people they refer enter a separate applicant pool, so they will not compete with them. Their own offer and pay do not depend on whether the referred individuals take up, or on their productivity. Appendix C.1 describes the hiring process, the baseline survey that records the referrals and announces the treatment, how the referrals who receive an offer are chosen, and the timeline.

Focal jobseekers refer 1.8 network members on average. Nearly all (91%) refer at most three, so in the treatment arms we saturate their referral networks with job offers. Most of the people they refer could commute together with them if assigned to the same worksite (88% on average). Doing so could make the commute more pleasant: most are friends (70% on average) and people the focal jobseekers feel very close to. As expected due to common homophily in networks, network members resemble the focal jobseeker in gender, age, and education (Appendix Table A7).[8]

3.2 Design

Our objective is to vary individuals' ability to commute with network members, and test whether labor supply changes. In this first experiment, we leverage the existence of two worksites of the same firm and randomly vary two dimensions.

7. Testing equality of means across the three arms, we find one difference, between the Same Site and Control focal jobseekers: those in the Same Site arm earned less in the 30 days before baseline (33,100 vs. 47,200 FCFA, $p = 0.067$), though they were no less likely to have had any earnings. All dimensions are balanced for the Different Site focal jobseekers relative to Control. Our results are similar whether or not we control for baseline earnings (Appendix Table A18).

8. Among male focal jobseekers, 82% of network members are men; among female focal jobseekers, 73% are women. Moreover, 61% of focal jobseekers are within 2.5 years of their network members' average age, and 60% are within one year of their average years of schooling.

First, we randomly assign each focal jobseeker to one of three arms. In *Control*, their network members do not receive a job offer. In *Same Site*, their network members receive an offer and are assigned to the focal jobseeker's worksite. In *Different Site*, their network members receive an offer but are assigned to the other worksite. This induces exogenous variation in the ability of the focal jobseeker to commute to work with their network members. When assigned to the Different Site arm, people can at most commute some of the way together, if they broadly live in the same direction to the two potential worksites. When assigned to the Same Site arm, people can commute together all of the way. Of course, this variation only applies to the focal jobseekers who referred at baseline individuals with whom they could commute with (rather than people who live far away and are not on their commuting path). Second, we randomly assign each focal jobseeker to work in one of the two potential worksites. This generates exogenous variation in their commuting times. Appendix C.1 details the announcement conditions.

Individuals assigned to work in the Same Site as their network members can commute together, but could also interact during the work day. Preventing these additional interactions is practically infeasible, if only because one could not prevent workers to freely interact during breaks. As such, comparing the labor supply of the focal jobseekers assigned to Same Site, relative to Different Site or Control, does not identify the role of joint commute. This is why we introduce the second source of variation: if the Same Site effects were fully driven by interactions on the job, these effects should not vary with commuting time. Heterogeneity in the Same Site treatment effects by commuting time therefore indicates a role for joint commute.

By contrast, we expect limited heterogeneity of the Different Site effects by commuting time, since the focal jobseekers have limited ability to commute to work together. If anything, any heterogeneity we find should be driven by the focal jobseekers who have a large co-commuting potential to the different worksites, because the two sites are broadly in the same direction from where they live. The details of this additional test are provided in Appendix C.3

The validity of this test would be threatened if the value of on-the-job interactions were correlated with commuting time. This might happen, for instance, if people who live in generally more remote places benefit differently from working alongside their friends than people living in a more central location in the city. Our second random variation—the focal jobseeker's worksite assignment—is designed to directly address that potential concern. Specifically, we isolate the exogenous variation in commuting time that we induce, and identify the causal heterogeneity in the Same Site treatment effects by commuting time. The estimation follows Borusyak and Hull (2023), as detailed below.

Offering jobs to network members could also affect labor supply through channels other than commuting together, for instance by changing redistributive pressure, beliefs about the firm, or opportunities to interact during breaks. We design the experiment to directly assess the relevance of these channels. These effects are indeed interesting in their own right, and can be used to benchmark the importance of joint commute. Comparing Different Site to Control isolates the effect of offering jobs to network members when interactions during the commute and the workday are limited. Comparing Same Site to Different Site among focal jobseekers with short, randomly assigned commutes isolates the effect of interacting at the worksite.

We estimate intent-to-treat effects from a comparison of means at the jobseeker level,

$$y_i = \alpha^s + \beta_1^s \text{DifferentSite}_i + \beta_2^s \text{SameSite}_i + \varepsilon_i^s, \tag{2}$$

where y_i is the outcome for focal jobseeker i and the two indicators mark assignment to Different Site or Same Site. Focal jobseekers were randomized in two waves, and their probability of being assigned to each treatment arm differs between the two waves. Regressions weight each focal jobseeker by the inverse of their probability of being assigned to their arm, so that each arm represents the whole sample, with robust standard errors. The main outcome is take-up—whether the focal jobseeker ever came to work.

We estimate heterogeneity by commute by interacting the treatment indicators with C_i, focal jobseeker i's round-trip commuting time to the assigned worksite (in hours, or an indicator for an above-median commute):

$$y_i = \alpha^s + \beta_1^s \text{DifferentSite}_i + \beta_2^s \text{SameSite}_i + \gamma^s C_i + \theta_1^s \,\text{DifferentSite}_i \times C_i + \theta_2^s \,\text{SameSite}_i \times C_i + \varepsilon_i^s. \tag{3}$$

Commuting time depends both on where the focal jobseeker lives and on which worksite is drawn, and only the draw is random. Following Borusyak and Hull (2023), we compute each jobseeker's expected commute under the actual assignment procedure, $\bar{C}_i$: the round-trip time to each candidate worksite, weighted by the probability that the procedure sends the jobseeker there. We use it in two ways.

First, we add controls for $\bar{C}_i$ to Equation (3).[9] Second, we estimate Equation (3) by two-stage least squares, instrumenting C_i and its interactions with the arms by the recentered

9. When interacting the arms with the above-median indicator, we control for indicators for each value of its expectation; with the continuous commute time, for indicators for the deciles of its expectation. Results are similar if we instead control for the continuous expected values.

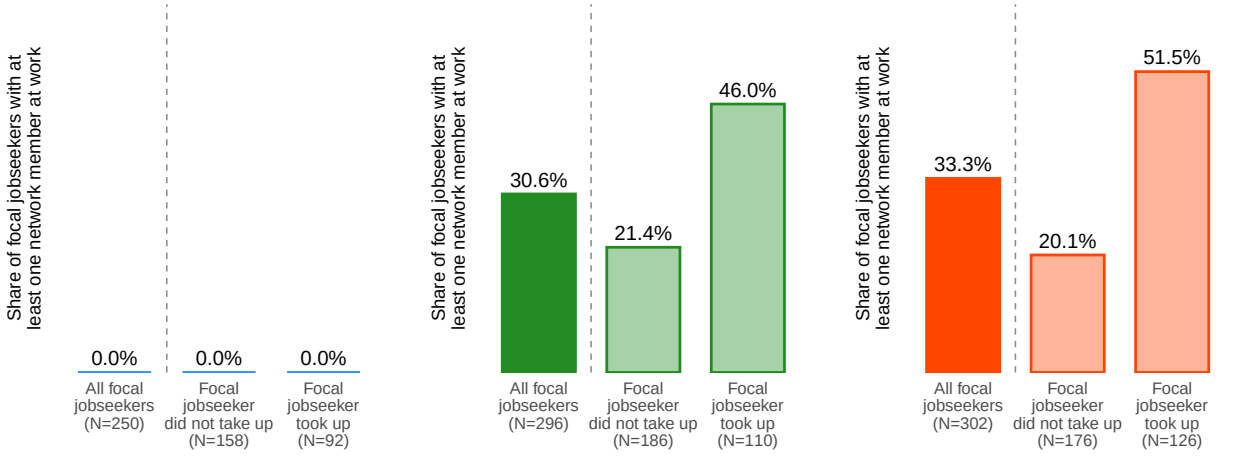
commute, $\tilde{C}_i = C_i - \bar{C}_i$, and its interactions with the arms. The recentered commute reflects only the random worksite draw: by construction, its expected value over the possible draws is zero for every jobseeker. The instrumented interactions therefore remain valid even if the treatment effects vary with where jobseekers live.

This robustness costs statistical power, because most of the variation in commuting time comes from where jobseekers live, not from the draw. The recentered commute has a little more than half the standard deviation of the realized commute (0.69 against 1.21 hours). The instrumented estimates can therefore be expected to be much less precise than the OLS ones.

3.3 Main Effects

The jobs offered to network members reach them: among the network members selected to receive a job offer, 73% are reached by the firm's hiring call. Most of those not reached had either invalid phone numbers or did not answer despite several attempts from the firm to reach them. Among those reached, 32% take up the job, on par with the take-up rate of Control focal jobseekers (33%, Table 2). Ultimately, a third of the focal jobseekers in the Same Site and Different Site arms have at least one network member working, against none in Control. Consistent with a coordination of labor supply decisions within networks, we find that network members are more likely to take up when their focal jobseekers also take up (Figure 4).

Figure 4: Network Members at Work



(a) Control (b) Different Site (c) Same Site

Notes: Share of focal jobseekers, in each treatment arm, with at least one network member who took up the job offer. "All focal jobseekers" is the unconditional within-arm share. The remaining bars are split by whether the focal jobseeker took up the offer. In Control, no network member was offered the job, so the share is zero by construction. Inverse-propensity weights applied where relevant.

Table 2: Treatment Effects, Sales Experiment

	Job take-up (worked $\geq$ 1 day)			Days worked
	(1)	(2)	(3)	(4)
Different Site	0.0294	0.0244	0.0249	0.0789
	(0.0463)	(0.0436)	(0.0416)	(0.213)
	[0.526]	[0.575]	[0.549]	[0.711]
Same Site	0.124	0.0884	0.0707	0.490
	(0.0469)	(0.0438)	(0.0418)	(0.216)
	[0.008]	[0.044]	[0.091]	[0.024]
p(Same Site = Different Site)	0.033	0.127	0.254	0.044
NMs who could commute together	All	≥ 1	—	All
Control mean	0.329	0.345	0.347	1.52
Observations	677	776	848	677

Notes: Treatment effects relative to Control. The dependent variable is job take-up (worked at least one day) in columns (1)–(3) and days worked in the first work week in column (4) (at most five in the first wave, four in the second). Inverse-propensity weights applied where relevant. Robust standard errors in parentheses, p-values in brackets. p(Same Site = Different Site): Wald test of equal effects. Columns (1) and (4) keep focal jobseekers whose offered network members could *all* commute together with them ("If you had to go to the same place in Abidjan at the same time, could you make the trip together?"), column (2) those with *at least one*; column (3) has no restriction. *Control mean* is the mean of the dependent variable in the Control arm.

Same Site raises take-up by 12.4 percentage points against a Control mean of 0.33, a 38% increase. It also raises days worked by 0.49 (Table 2, Col. 4). The Different Site effect is close to zero and insignificant across outcomes. Offering the job to network members does little unless they are assigned to the same worksite as the focal jobseeker. As such, the labor supply of the focal jobseekers in the Same Site treatment arm is also significantly higher than in the Different Site treatment arm: complementarities in labor supply driven by the ability to interact exist and are strong.

These effects are estimated on the focal jobseekers whose offered network members could all commute with them. This characteristic is elicited at baseline, so assignment remains randomized within the subsample. The restriction sharpens the estimand to focal jobseekers whose networks can actually commute together and raises power. The take-up effect is smaller but of the same sign in the full sample of all jobseekers, which also includes those who could not commute together with the network members they referred (Col. 3).

The Same Site treatment effects increase with the length of the commute. The interaction of the Same Site arm with round-trip commute time to the assigned worksite is 7.4 points

Table 3: Treatment Effects by Commuting Time, Sales Experiment

	Above-median commute			Commute time (hours)		
	(1)	(2)	(3)	(4)	(5)	(6)
Covariate	-0.0572	-0.0644	-0.0684	-0.0112	-0.0115	-0.0276
	(0.0744)	(0.0874)	(0.118)	(0.0283)	(0.0334)	(0.0527)
	[0.443]	[0.462]	[0.563]	[0.691]	[0.731]	[0.601]
Different Site	-0.0170	-0.0178	-0.0311	-0.0419	-0.0479	-0.0650
	(0.0563)	(0.0566)	(0.0688)	(0.0915)	(0.0927)	(0.142)
	[0.763]	[0.753]	[0.651]	[0.647]	[0.606]	[0.647]
Different Site × Covariate	0.137	0.137	0.178	0.0335	0.0384	0.0450
	(0.0990)	(0.0994)	(0.155)	(0.0377)	(0.0380)	(0.0642)
	[0.166]	[0.170]	[0.251]	[0.374]	[0.312]	[0.483]
Same Site	0.0372	0.0379	0.0501	-0.0270	-0.0273	-0.0151
	(0.0574)	(0.0576)	(0.0731)	(0.0910)	(0.0914)	(0.154)
	[0.517]	[0.511]	[0.493]	[0.767]	[0.765]	[0.922]
Same Site × Covariate	0.242	0.242	0.210	0.0712	0.0745	0.0663
	(0.0989)	(0.0993)	(0.162)	(0.0370)	(0.0369)	(0.0696)
	[0.014]	[0.015]	[0.194]	[0.055]	[0.044]	[0.341]
Expected-commute controls	No	Yes	No	No	Yes	No
Estimator	OLS	OLS	IV	OLS	OLS	IV
First-stage F			368.3			11.3
Control take-up	0.33	0.33	0.33	0.33	0.33	0.33
Observations	677	677	677	677	677	677

Notes: Regressions of job take-up (an indicator equal to one if the focal jobseeker worked at least one day) on the Same Site and Different Site indicators, a commute measure (the "Covariate"), and its interaction with each arm. The estimation sample is the focal jobseekers for whom *all* of their offered network members could commute together with them, elicited at baseline ("If you had to go to the same place in Abidjan at the same time, could you make the trip together?"). In columns (1)–(3) the covariate is an indicator for an above-median round-trip commute to the randomly assigned worksite; in columns (4)–(6) it is that commute in hours. The expected commute averages the commute to the two candidate worksites, using the probability that the assignment procedure sends the focal jobseeker to each. Columns (2) and (5) control for it: indicators for each value of the expected above-median indicator in column (2), indicators for its deciles in column (5). Columns (3) and (6) instrument the commute measure and its interactions with the arms by the commute minus its expected value and that residual's interactions with the arms; the residual reflects only the random worksite draw. *First-stage F* is the heteroskedasticity-robust Kleibergen–Paap rk Wald F statistic. Inverse-propensity weights applied where relevant. Each cell reports the coefficient, the robust standard error (in parentheses), and the p-value [in brackets]. *Control take-up* is the take-up rate in the Control arm.

per hour ($p = 0.04$, Table 3, Col. 5). The same pattern emerges when splitting commute time at its median: the effect is 24 percentage points larger above the median and near zero below it (Col. 1, plotted in Appendix Figure A15). This allows us to isolate commuting from jobseekers interacting on-the-job. If working alongside network members were driving the effect, it would not vary with the length of the trip.

Using only the variation from the random worksite draw gives almost the same interactions: 6.6 points per hour and 21 points above the median (Cols. 6 and 3), about 90% of the estimates that control for the expected commute. They are not statistically significant on their own ($p = 0.34$ and 0.19) because they are less precise, not because they are smaller: their standard errors are 1.9 and 1.6 times larger, close to what the smaller variation of the recentered commute predicts (Section 3.2). The design-based estimates are therefore consistent with the main estimates: they differ in precision, not in magnitude.

The treatment effect heterogeneity is not driven by a few commute lengths. Appendix Figure A16 reports the Same Site effect within 30-minute bins of the commute. It is small and not significant up to two hours, and positive in every bin beyond, from 28 to 70 percentage points, significantly so at three, three and a half and four hours. Where it starts to rise is harder to read, because answers bunch on round values: 26% of focal jobseekers report a round trip of exactly two hours. These reported times are coarse. An answer of two hours broadly means that the commute is long, not that it lasts exactly two hours.

The Different Site effect increases with the commute, but only where focal jobseekers can still share part of the trip with their network members. The Different Site arm is not a clean "no co-commuting" condition: tracing each focal jobseeker's and their network members' public-transport trips from home to work, we find that focal jobseekers assigned to different sites still share part of the commute with a network member, whereas almost all Same Site focal jobseekers could commute together (Figure A10). However, splitting the Different Site arm by this residual overlap, the heterogeneity by commuting time appears only among focal jobseekers with a high degree of overlap: 26.9 percentage points vs. 1.9 where the overlap is low—though with only about 215 focal jobseekers per group, the two are not statistically distinguishable ($p = 0.21$; Table A10). The small heterogeneity of the Different Site effect by commuting time thus reflects the ability to commute together, not the site assignment itself.

Two comparisons speak to channels other than commuting together. Comparing Different Site to Control isolates the effect of offering jobs to network members when interactions during the commute and the workday are limited: the difference in take-up is 2.9 percentage

points ($p = 0.53$, Table 2, Col. 1). Comparing Same Site to Different Site among focal jobseekers with below-median commutes isolates the effect of interacting at the worksite: the difference is 5.4 percentage points ($p = 0.33$, Table 3, Col. 1). Neither comparison yields a significant difference in take-up, suggesting these channels play a limited role in our setting.

3.4 Limited Productivity Consequences for Firms

Whether hiring network members together carries a productivity cost is of first-order importance for firms, and the sign is ambiguous ex ante. Workers who know each other may learn faster and motivate each other, but they may also shirk and chat together. There is also a selection effect: the marginal jobseekers, who take up only when their network members are offered a job, may differ in productivity from the inframarginal ones. Likewise, the referred individuals may differ in productivity from the referrers.

We measure productivity from two independent sources: the evening debriefs conducted by the workers' field supervisors, and the firm's app logs. At the end of each workday, workers are gathered by their supervisors who record four dimensions of sales: the number of people the worker approached, the number they helped to access the app or the associated website, the number they helped to create an account, and the number they helped to turn on notifications.[10] The firm also collects data through its app. Each workers carries an individual QR code, which links every download to the worker who induced it. For 20 days after each download we observe the user's interactions with the app along multiple dimensions. More details on these data are provided in Appendix C.4.

Being a digital classified ads company, the firm mostly cares about the number of contacts between potential buyers and sellers that are induced by the campaign. This is proxied in the data by the number of times a user clicked to access a seller's contact details, which gives us a measure of quality-adjusted productivity. In addition, since we collect multiple dimensions of productivity, we construct three standardized indices (Anderson 2008; Kling et al. 2007): one from the field supervisor's evening debrief, one from the firm's app logs measuring users' activity on the day each user was recruited, and one from the firm's app logs measuring users' activity for 20 days after they download the app. The latter is our preferred measure, since it comes closest to the objective of the firm.

Overall, the increase in take-up induced by our Same Site hiring intervention does not come at the expense of productivity (Figure 5, Panel A). On average, focal jobseekers who took up

10. The supervisors conduct regular spotchecks throughout the day to ensure that workers are not shirking and to verify the accuracy of the evening reports. Pay is based on daily attendance only, so neither workers nor supervisors have a financial incentive in misreporting.

the job when hired together with their network members to work in the same worksite have similar productivity levels as those hired alone. While the difference is never statistically significant, these workers are on average slightly more productive if considering the app logs data, or slightly less productive if considering the evening debriefs.

Importantly, the value from the Same Site hiring is concentrated among focal jobseekers with above-median commuting time. For them, whatever the data source we consider, the Same Site workers are always more productive than the Control workers; although this relationship remains non-statistically-significant (Figure 5, Panel C). This holds whether we consider the productivity indices, or the key outcome of seller contacts (Appendix Table A11, Panel B).

While firms are interested in overall effects (i.e., what is the average productivity of the workers who come to work), we can be interested in disentangling the interaction and selection effects of hiring workers to work in the Same Site. This is enabled by the fact that we find no treatment effects on job take-up for individuals with below-median commuting time. There, the productivity effects are fully driven by the interaction effects of being hired with network members, while the selection channel is shut down. The effects are also statistically insignificant. The point estimates are negative for both the Different Site and Same Site treatment arms when considering the evening debrief's productivity index and the app logs for the seller contacts (Appendix Table A11, Columns (3), (4) and (5)).

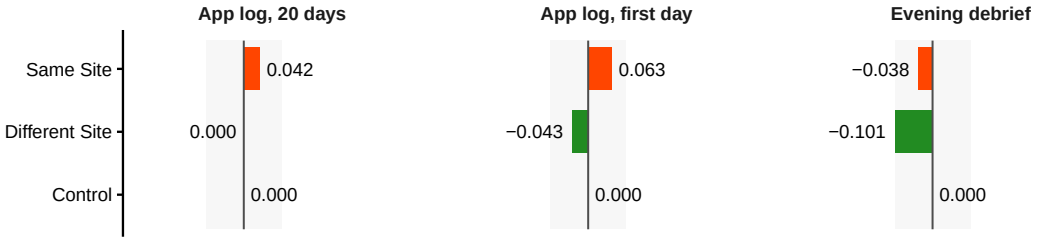
Taken together, these point estimates suggest a potential drop in effort for jobseekers hired with network members (but not driven by the fact of working together), offset by the selection effect for jobseekers with long commutes. We emphasize, though, that the estimates remain too noisy for a conclusive answer: we can at best rule out average productivity losses larger than about a quarter of a standard deviation, whether we consider the indices or each of their components separately (Appendix Figure A11). Hiring through networks raises labor supply without significantly reducing the average productivity of the workers hired.

The firm also cares about the quality of the network members being hired. In our setting, the individuals being referred at the hiring stage, who receive a job offer and decide to take up the job, have a slightly lower productivity than the focal jobseekers, to whom the firm originally wanted to offer a job. However, this average productivity difference is minimal: depending on the productivity index considered, the network members are on average 0.03 to 0.06 standard deviations less productive than the focal jobseekers (Figure 5, Panel B).

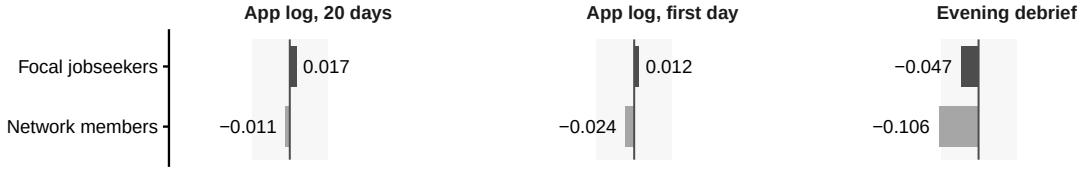
Overall, we find no statistically significant differences in productivity across treatment arms, or between focal jobseekers and network members. For a firm, hiring focal jobseekers with long commutes alongside network members they can commute with raises average produc-

Figure 5: Productivity of Focal Jobseekers and Network Members, Sales Experiment

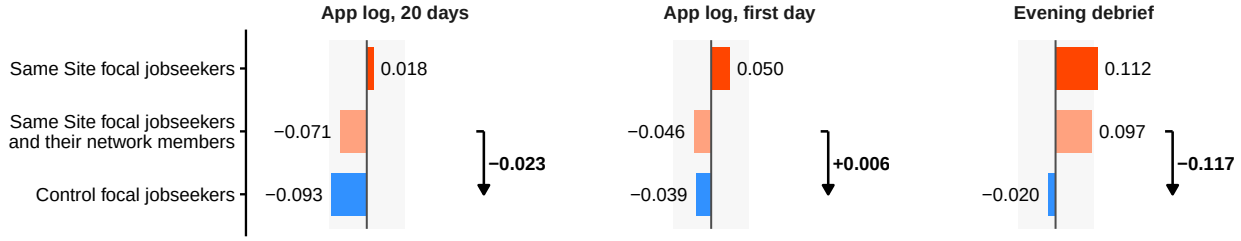
(a) Focal jobseekers, by treatment arm



(b) Focal jobseekers and network members



(c) Focal jobseekers with above-median commuting time, and their network members



Notes: Average productivity of workers in the sales experiment. Productivity is an index of sales activity, built separately from three data sources: the field supervisors' evening debriefs, and the firm's app log on the day each user was recruited and over the 20 days after each download. Each component of the index is standardized against the Control group, averaged within a data source, and standardized again. Panel (a) includes all focal jobseekers from our study sample, who took up the job. Panel (b) compares all focal jobseekers who took up the job with the network members who took up the jobs. Panel (c) keeps focal jobseekers with above-median commuting time, retaining the normalization from Panel (a): Same Site focal jobseekers who took up the job, the same pooled with the network members hired through them, and Control focal jobseekers, with arrows reporting the difference in average productivity of the Control focal jobseekers relative to the pooled group. The light grey band marks ±0.1 standard deviations. The figure reports point estimates only. Appendix Table A12 reports the standard error and p-value of each comparison made in the text; none is significant at the 10% level. Inverse-propensity weights applied where relevant.

tivity by 0.02 standard deviations, which is neither economically nor statistically significant: the labor supply gains do not come at a productivity cost.

3.5 The Amenity Value of Commuting Together

We express the value of commuting together as the wage increase that would raise take-up by as much as Same Site assignment. Take-up was low during the operational pilot of the sales experiment, and the daily wage was raised from 7,000 to 9,500 FCFA (∼USD 11 to 15), a 36% increase. Among focal jobseekers in the Control and Different Site arms, where offering jobs to network members does not raise take-up, take-up rose from 9.3% in the two pilot rounds to 23.4% in the two rounds that followed (Appendix C.5). Raising take-up by one percentage point thus cost 177 FCFA per worker-day.

Comparing this response with the Same Site effect recovers, under assumptions on the distribution of jobseekers' reservation wages, the wage increase equivalent to commuting together. If take-up rises linearly with the wage, matching the 12.4 percentage point Same Site effect requires raising the wage by 23%; a normal distribution of reservation wages, or the smaller effect in the full sample, gives lower values. Depending on these assumptions, the ability to commute together is equivalent to a wage increase of 9 to 23%.

3.6 Why Workers Value Commuting Together

Why would jobseekers value commuting together? Five types of mechanisms could be plausible: sharing the monetary cost of the trip, gendered attitudes toward commuting, motivating each other to avoid absenteeism, greater safety, and commute becoming more pleasant. We can rule out the first three, and find empirical support for the last two: commuting together being safer and more pleasant.

Cost-sharing. Cost-sharing can be ruled out as a mechanism. In the informal public transport used by most workers, fares are charged per passenger and do not fall when riders travel together. So, commuting together or alone leaves the monetary cost of the trip unchanged. Nor does it lower the time cost: both common modes of transport—*gbakas* and *woro-woros*—fill up from whoever is waiting at the stop, and the routes and number of vehicles are set by the operators' associations. One might still ask whether traveling together makes a vehicle fill up, and so leave, sooner. This is unlikely to matter: commuting together adds at most one or two riders, the network members hired alongside the focal jobseeker, to stops used by many other commuters.

We can also test directly whether commuting cost drives our results. Commuting time and cost are positively correlated, but with enough variation that we can separately identify their interaction with the treatment effects (Appendix Figure A12). While the Same Site effects increase significantly with commuting time (Table 3), they do not vary with commuting costs (Appendix Table A14). When both interactions enter the same regression, the interaction with commuting cost is small and not significant, while the interaction with commuting time remains significant and close to its value in Table 3 (Appendix Table A14).

Gendered attitudes. Gendered attitudes toward commuting are unlikely to account for the effects. Such attitudes are relatively permissive in this setting to begin with (Section 2).

Moreover, if commuting together were valuable mainly because it makes commuting socially acceptable for women, the Same Site effect would be concentrated among them. Instead it is 11.8 percentage points for men and 12.6 for women, statistically indistinguishable ($p = 0.94$, Appendix Table A17), despite men and women differing at baseline in earnings and childcare (Appendix Table A16). The Same Site heterogeneity by commuting time is larger for women (Appendix Table A17), though this difference is significant only in the continuous-commute specification and rests on 233 women.

Note that because both men and women applied to these jobs, our sample is selected on willingness to commute to work overall. Gendered attitudes could still be deterring labor supply earlier, at the application stage.

Motivation. Workers can suffer from temporary lack of motivation. Having someone expecting them to go and commute together to work can help overcome this lack of motivation (for instance to avoid the social cost of letting the network member go to work alone, or by being called and motivated by the network member). A sophisticated present-bias worker, anticipating these future motivational benefits, may then prefer to take up a job when they can expect to commute and be motivated by their network members.

While theoretically plausible, this idea is unlikely to apply in our setting. Due to the short-term nature of the offered jobs, absenteeism is indeed almost nonexistent. Conditional on taking up the job, 85% of the workers in the control are present for all the days they could work, and 82% in the same site group. As such, there is limited scope for increased motivation as a mechanism for this sales experiment.

Safety. Safety is a mechanism, driving part but not all the treatment effects. First, safety perceptions influence labor supply decisions. In our sample, 54% of the jobseekers report

feeling very safe when commuting back home at night (52% among women, and 55% among men). For individuals offered the job alone (control group), those reporting feeling less safe coming home at night are less likely to take up (Appendix Figure A13). This holds true for both genders, but is especially strong for women: in the control group, 31% of those reporting feeling very safe take-up, against 26% for those feeling somewhat safe, and none of the four women in the control group who report not feeling safe took up the job (Appendix Figure A14). (For men, take-up decreases from 42% to 27% for those reporting feeling very safe to not safe.)[11]

We report heterogeneity in treatment effects by baseline safety perception in Appendix Table A15. The Same Site treatment effects are large for the focal jobseekers who do not feel very safe commuting back home at night (+16.2 percentage points, $p = 0.017$). These effects are twice larger than for those reporting feeling very safe (+8.5 percentage points, $p = 0.193$), even though the difference is not statistically significant.

Our sample includes both individuals who have been commuting to work at baseline (despite, for some of them, not feeling safe commuting back at night) and others who were not commuting, either because they didn't have a baseline income-generating activity or because they had one close to their home. By revealed preference, those commuting to work despite not feeling safe can be expected to be less elastic to safety when responding to our job offers. We indeed find that in the control group, the relationship between safety and take-up is especially strong for individuals who weren't commuting to work at baseline, twice larger than for those who were already commuting to an income-generating activity.

Consistently, safety appears to be the main mechanism driving the treatment effects for the 49% of the focal jobseekers with no baseline experience commuting to work, but not for the 51% others. Results disaggregated by this baseline characterstic are presented in Appendix Table A15. For individuals who weren't already commuting at baseline, the Same Site treatment effects are fully concentrated among individuals who report not feeling very safe (−1.9pp for individuals reporting feeling very safe at baseline, against 21.6pp for those reporting not feeling very safe, $p = 0.026$). For the others, the Same Site effects are also large but remain similar whether they feel safe or not (+17.9pp and +11.3pp, respectively). This is not a gendered pattern: both men and women exhibit the same behavior of treatment effects concentrated among those not feeling very safe (Cols. 4-5).

11. We also ask whether the jobseekers have ever been bothered or harassed on their way home from work. However, the answers are not correlated with take-up in the control group, suggesting that this elicitation of safety perception may not be relevant for labor supply decisions. The treatment effects are, if anything, larger for individuals who never have been harassed (Appendix Figure A13).

Taken together, these results indicate that safety concerns are likely to explain part of the results–especially for the sample of jobseekers who were not already commuting to work at baseline, and for whom safety constraints may play an especially important role–but not all of them.

Value of social interactions. We find a strong social component attached to commuting together. This is directly reflected in the type of co-commuters being referred, and consistent with the factory evidence where workers report that commute becoming more pleasant is among the two main benefits from being able to commute together (Section 2.2).

Specifically, when the focal jobseekers refer individuals with whom they could commute with, they are referring individuals to whom they feel especially close and with whom they are used to spend time regularly. They feel very close to the individuals they refer in 71% of the cases, against not very close or not close at all in only 4% of the cases. The remaining 25% are individuals they feel somewhat close to.[12] At baseline, the focal jobseekers routinely interact with the people they refer: 88.5% of them report spending at least 2 hours with them in a usual week, with a median weekly time together of 9.4h.

By contrast, when the focal jobseekers refer individuals with whom they could not commute (i.e., those included in the 848 individuals in the full sample of Column 3, Table 2), these are also people they are less close to. They report feeling very close to only 43% of them, and spend twice less time together at median.[13] This indicates that commuting together bears a social component: the focal jobseekers decide that, when they refer someone they could commute with, that person will also be someone they value spending social time with.

4 Factory Experiment

We conduct a second field experiment, with a multinational company operating an agri-processing factory in Abidjan. This allows us to test whether our main results hold for a different job and population: women with limited schooling, rather than the young university graduates who canvass for our sales partner. It also allows us to assess the persistence of the effects. Given the short-term nature of the jobs, the sales experiment can speak to job

12. They feel especially close to the family members they refer, with very close relationships reported for over 90% of the referred cousins, siblings, and parents. They also feel very close to 71% of the friends they refer and, by definition, to a lower 40% of the acquaintances they refer.

13. This is not just driven by the composition of who they refer: even when referring friends, they only report being very close to 50% of them, against 73% when they could commute together. This is also not driven by living close to each other: among non-neighbors, the focal jobseekers report being much closer to the non-neighbors they could commute with than those they cannot commute with.

take-up and productivity but not to retention. This second experiment lets us follow focal jobseekers for up to a year and estimate effects on the attendance and retention of the job takers.

4.1 Setting

Our partner is Olam, a large multinational and one of the largest employers in Côte d'Ivoire. Olam planned to expand its cashew-processing activities and open a new factory in Anyama, on the outskirts of Abidjan. The factory was hiring for production roles: the manual peeling and grading of nuts, an individual and repetitive task, and the packaging of processed nuts, done by a group of workers along a production line. Around three-quarters of these roles are traditionally filled by women. Skills are difficult to screen ex ante, so the eligibility criteria are essentially the legal requirements for formal work: being at least 18 years old and producing an identity document. Contracts are open-ended, with the expectation of a long-run employment relationship. Workers are paid 4,000 FCFA (∼6 USD) per day of work, slightly above the national minimum wage. And of relevance to our experiment, a large share of Olam's workers find their jobs through their networks. The factory has only one worksite, but two non-overlapping shifts: a morning and an afternoon shift. The shifts are large, with over fifty workers in each, so that workers assigned the same shift are unlikely to be able to interact during work. Workers keep the same co-workers throughout, but whether they work mornings or afternoons rotates weekly. Network members on different shifts work at different hours, so they cannot commute together.

The sample is 163 eligible focal jobseekers, balanced across arms (Appendix Table A19).[14] Most have some work experience but limited schooling, averaging 29 years of age and a middle-school education, and over half report having a partner. Consistent with the composition of the factory's workforce, 72% are women. The job is desirable relative to their outside option: nearly a third had no income-generating activity in the three months before applying, and for the rest it pays over twice their baseline earnings for similar hours.

Two network members receive an offer per focal jobseeker. As for the sales experiment, the network members are elicited at the hiring stage, and treatment is announced afterwards. Appendix D.1 describes the hiring process, the timeline, and the randomization. Focal

14. Testing equality of means across three arms, we find a small number of differences, all between the Different Shift and Control focal jobseekers. All dimensions are balanced for the Same Shift focal jobseekers relative to Control. Specifically, focal jobseekers in the Different Shift arm are less likely to have spent time on childcare in the past week, and more likely to have had any earnings in the past 30 days, than those in Control. Our results are similar whether or not we control for these two baseline variables, with the retention effect slightly more precise (Appendix Table A21).

jobseekers refer 2.9 network members on average. On average, 45% of them are friends, 25% siblings and 11% cousins (Appendix Table A20).

4.2 Design

The design mirrors the sales experiment's (Section 3.2). We induce random variation in the focal jobseeker's ability to commute with their network members, by randomly assigning them to *Control*, in which their network members receive no offer, *Same Shift*, in which their network members receive an offer and would work in the same shift, and *Different Shift*, in which their network members receive an offer but would work in different shifts. The shifts are large, and so workers assigned the same shift are unlikely to be able to interact besides commuting together and possibly spending breaks together. There is only one worksite, so no possibility to induce random variation in commuting time in this second experiment.

Factory records compile each worker's attendance every day. They are used to calculate pay and are therefore accurate.[15] From them we measure take-up, the days each focal jobseeker works, how long they are retained, and their attendance while retained. Retention and attendance are observed only for those who take up, so those estimates combine direct and selection effects.

As in the sales experiment, we estimate intent-to-treat effects from a comparison of means at the jobseeker level,

$$y_i = \alpha^f + \beta_1^f \text{DifferentShift}_i + \beta_2^f \text{SameShift}_i + \varepsilon_i^f, \tag{4}$$

where y_i is the outcome for focal jobseeker i, and the two indicators mark assignment to Different Shift or Same Shift. Focal jobseekers were randomized in two hiring waves, and their probability of being assigned to each arm differs across waves. As in the sales experiment, regressions weight each focal jobseeker by the inverse of their probability of being assigned to their arm. We report robust standard errors and randomization-inference p-values, which are valid in this small sample.

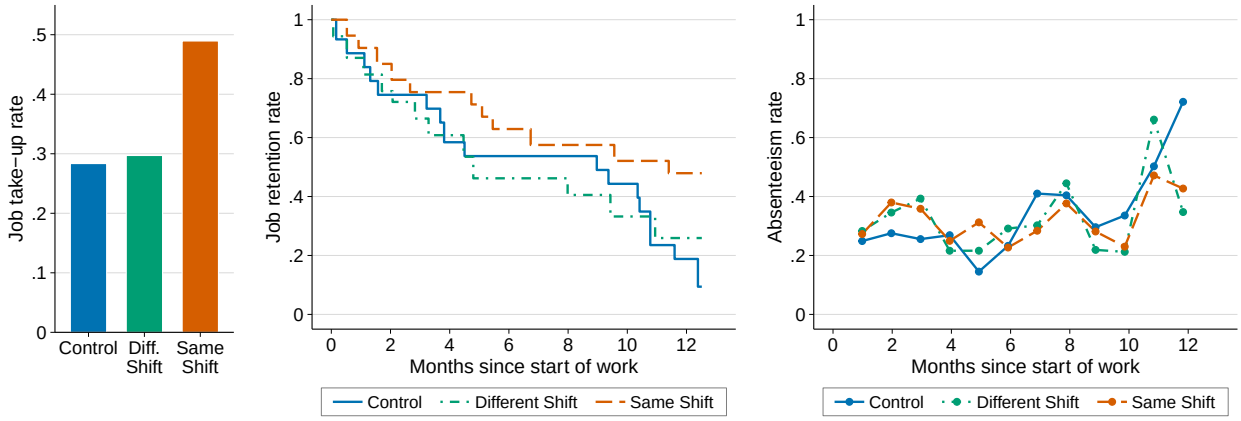
4.3 Results

Take-up in the Control arm illustrates the hiring difficulty our partner faces. These are jobseekers pre-selected for their potential interest , offered a job paying over twice their baseline earnings, and only 28% ever come to work. Of those who do, absenteeism is high

15. The factory was not collecting worker-level production data at the time of the experiment, so we have no productivity measures here.

and a large share leave within months (Figure 6). This is the pattern of hiring and retention difficulty firms across sub-Saharan Africa report at prevailing wages.

Figure 6: Patterns of Labor Supply, Factory Experiment



(a) Take-up

(b) Retention, if took up

(c) Monthly absenteeism, while job retained

Notes: Take-up, retention, and absenteeism from raw data, for the 163 focal jobseekers of the factory experiment. Take-up equals 1 if the focal jobseeker worked at least one day. Retention is the share of focal jobseekers who took up the job still working at the factory in month d. Absenteeism is the share of the factory's open days on which retained workers are absent, computed for each month of tenure over the workers still employed in that month. Inverse-propensity weights applied where relevant.

Being assigned to the Same Shift arm, relative to Control, raises the total number of days worked at the factory by 52.5, against a Control mean of 39.5: a 33% increase (Table 4, Col. 1). This sharp increase in labor supply mostly comes from the extensive margins: whether a focal jobseeker takes up the job, and how long they stay. Same Shift indeed raises take-up by 20.6 percentage points against a Control mean of 0.28, a 73% increase (Col. 2). Among those who take up the job, retention is 96.5 days higher, though imprecisely estimated ($p = 0.101$, (Col. 3), while attendance is unchanged (Col. 4). By contrast, and consistent with the sales experiment, the Different Shift effect is close to zero on every outcome. Offering network members a job does little unless they could commute with the focal jobseeker.

Table 5 shows retention results over the course of the year, recording whether each is still employed at the factory at increasing horizons. These outcomes are defined for everyone offered the job, so they combine the decision to start with the decision to stay and require no conditioning on take-up. Being assigned to the Same Shift as your network members raises employment at the factory at every time cut-off, from 20.6 percentage points at one day or longer, to 12.9 points at nine months or longer, with randomization-inference p-values between 0.010 and 0.074. Moreover, the Different Shift effect stays near-zero throughout, and the two coefficients are significantly different at every threshold (p-value between 0.036 and 0.055).

Table 4: Main Treatment Effects, Factory Experiment

	Days worked (1)	Take-up (2)	Retention (3)	Attendance (4)
Different Shift	2.89	0.0136	13.1	-0.00940
	(17.2)	(0.0847)	(65.4)	(0.0621)
	[0.725]	[0.860]	[0.732]	[0.962]
Same Shift	52.5	0.206	96.5	-0.0357
	(23.9)	(0.0950)	(61.6)	(0.0576)
	[0.012]	[0.030]	[0.101]	[0.610]
Control mean	39.5	0.284	224	0.710
Observations	163	163	57	57

Notes: Each column regresses the outcome on indicators for the Different Shift and Same Shift arms; Control is the omitted group. Columns (1)–(2) include all 163 focal jobseekers; columns (3)–(4) the 57 who took up the job. Retention counts the days from the first to the last day at work; attendance is the share of the factory's open days in between on which the worker came. Inverse-propensity weights applied where relevant. Robust standard errors in parentheses; randomization-inference p-values in brackets.

Suggestive evidence indicates that in this setting leaving the factory often means losing wage employment, rather than leaving for a better job opportunity. Our sample for the network survey conducted among workers from the same factory, several months after this experiment, includes 61 individuals who had left the factory at the time of the survey.[16] Despite having left on average 45 days before the survey, only 13% of them reported any earnings since leaving. Even among those who earned something, median earnings since leaving were less than a quarter of their previous pay at Olam. As such, the treatment effects on job take-up and turnover are likely to translate into increases in overall employment and earnings.

The Same Shift effect is concentrated among focal jobseekers with longer commutes. It is small and not significant for commutes of up to one hour, and positive and significant beyond (Appendix Figure A17). This evidence is suggestive rather than causal. The factory sample is small and underpowered for heterogeneity analysis, and commuting time is not randomized here, so it may proxy for other characteristics that shape treatment effects. Still, it is fully consistent with the results from the sales experiment that has a larger sample size and randomly varies commuting time.

16. We do not have survey evidence from the workers enrolled in the factory experiment besides their baseline survey, and thus cannot study their potential activities outside the factory directly. However, the sample of workers enrolled in the experiment has broadly similar characteristics as the sample of workers from the network survey, with all workers working at the same factory (Appendix Table A2).

Table 5: Factory Employment at Increasing Horizons, Factory Experiment

Still employed:	$\geq$ 1 day	$\geq$ 8 days	$\geq$ 2 mo.	$\geq$ 4 mo.	$\geq$ 6 mo.	$\geq$ 9 mo.
	(1)	(2)	(3)	(4)	(5)	(6)
Different Shift	0.0136	0.0157	0.0138	0.0150	-0.0151	-0.0319
	(0.0847)	(0.0829)	(0.0770)	(0.0711)	(0.0654)	(0.0637)
	[0.860]	[0.807]	[0.812]	[0.717]	[0.948]	[0.743]
Same Shift	0.206	0.225	0.205	0.204	0.156	0.129
	(0.0950)	(0.0940)	(0.0910)	(0.0870)	(0.0839)	(0.0822)
	[0.030]	[0.016]	[0.017]	[0.010]	[0.033]	[0.074]
p(Same Shift = Different Shift)	0.055	0.036	0.047	0.042	0.050	0.055
Control mean	0.284	0.265	0.211	0.166	0.152	0.152
Observations	163	163	163	163	163	163

Notes: Each column regresses an indicator for still being employed at the factory at the horizon in its header on the Same Shift and Different Shift indicators, for all 163 focal jobseekers offered the job. A focal jobseeker counts as employed at a horizon if the span from their first to their last day at the factory, both included, is at least that long. Those who never came count as zero, so column (1) is job take-up (Table 4). Inverse-propensity weights applied where relevant. Robust standard errors in parentheses; randomization-inference p-values in brackets. p(Same Shift = Different Shift) tests that the two arms have equal effects; *Control mean* is the share employed in the Control arm.

5 Conclusion

In cities of lower-income countries, where transport is difficult and costly, being able to commute together with network members is a major non-wage job amenity. We find this consistently across different types of jobseekers, each offered a different type of job in Abidjan: short-term sales work and long-term factory employment. Offering this amenity increases job take-up without reducing retention, attendance or productiivty, leading to increases in overall labor supply to the firm.

This has immediate implications for firms' personnel policies. Firms can increase the labor supply they face by eliciting and hiring referrals at the hiring stage, providing they enable them to commute together with the person who referred them. However, this can come at the cost of increased volatility: workers who commute together tend to be absent and quit at the same time. The level of interdependencies in the production process, as well as the cost of hiring and training new workers, will likely influence which effect dominates and whether firms prefer hiring through referrals or the open labor market.

We find that commute being safer and more pleasant are two key benefits from being able to commute together. This suggests additional avenues through which firms could increase the labor supply they face, including arranging secure shuttles or encouraging the creation of new social ties among workers—to help them learn about the set of potential co-commuters,

and make them more pleasant to commute with.

The implications of our work extend beyond firms' decisions, to any setting where commuting costs can limit the take-up and long-term efficacy of policies. Targeting social groups who could commute together, rather than individuals in isolation, can be a potent policy lever. This naturally applies to active labor market interventions, which have been widely implemented across lower-income countries with often limited success in raising long-term employment for participants. It could also be relevant for domains such as health and education, where the sustained participation of both participants and practitioners are long-standing concerns.[17]

Our results also shed light on the spatial distribution of economic activity within cities. They suggest a potentially novel microfoundation for the agglomeration preferences of firms within cities: by locating where large numbers of workers are already going to work, they increase the likelihood that the jobseeker they hire would be able to commute with some of them. Their location decisions thereby allows them to increase the labor supply they face at prevailing wages. Likewise, urban planners involved in the design of zoning or transport infrastructure should consider how it would influence the ability of workers to commute together. For instance, having jobs scattered across the city could reduce road congestion but make co-commuting with neighbors more difficult.

More broadly, this work contributes to the discussion on the existence and effects of multiple equilibria in labor supply. A body of theory shows that complementarities across individuals can give rise to multiple equilibria (Cooper and John 1988; Sah 1991; Murphy et al. 1993). Glaeser et al. (1996, 2003) develop indices of the strength of such interactions, and show that they generate far more variation in aggregate outcomes than independent decisions would. Their applications include criminal activity and unemployment. Alesina et al. (2005) build on this idea to discuss the possibility that complementarities in labor supply could help explain why hours worked drastically reduced in Europe starting in the early 1970s but not in the United States. Our study provides direct causal evidence on the existence of strong complementarities in labor supply, supporting 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

17. In the areas of entrepreneurial training and clinic visits, take-up and effectiveness have been shown to increase when potential participants are invited with network members, especially for women (Field et al. 2016; Anukriti et al. 2026). We bring this idea to the domain of labor supply and introduce the role of joint commute as a key driver. We thereby add network-based hiring to a growing literature designing and evaluating labor market interventions in lower-income settings. Previously-considered interventions include mentoring (Alfonsi et al. 2024), training (Alfonsi et al. 2020), matching (Bandiera et al. 2025), skills assessments (Carranza et al. 2022; Kiss et al. 2023) and job search training (Wheeler et al. 2022).

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 have persistently different patterns of labor supply, including formal wage employment. This also implies the presence of a social multiplier amplifying the labor supply consequences of shocks. Localized disruptions to the transport network, or shocks leading some workers to be absent or quit, will propage to their co-commuters and lead to potentially large aggregate labor supply consequences.

References

Abebe, Girum, Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn. 2021. "Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City." *Review of Economic Studies* 88 (3): 1279–1310.

Afridi, Farzana, and Amrita Dhillon. 2022. "Social Networks and the Labor Market." In *Handbook of Labor, Human Resources and Population Economics,* edited by Klaus F. Zimmermann, 1–18. Springer International Publishing.

Afridi, Farzana, Amrita Dhillon, Sanchari Roy, and Nikita Sangwan. 2025. *Social Networks, Gender Norms and Labor Supply: Experimental Evidence Using a Job Search Platform.* QBS Working Paper 2025/02. Queen's University Belfast, Queen's Business School.

Afrobarometer. 2024. *Afrobarometer Round 10 Survey in Côte d'Ivoire.* Online Data Analysis facility, https://www.afrobarometer.org/online-data-analysis/.

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2015. "The Economics of Density: Evidence from the Berlin Wall." *Econometrica* 83 (6): 2127–2189.

Akbar, Prottoy A., Victor Couture, Gilles Duranton, and Adam Storeygard. 2023a. "Mobility and Congestion in Urban India." *American Economic Review* 113 (4): 1083–1111.

———. 2023b. *The Fast, the Slow, and the Congested: Urban Transportation in Rich and Poor Countries.* NBER Working Paper 31642. Revise and resubmit, *Quarterly Journal of Economics*. National Bureau of Economic Research.

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.

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.

Alfonsi, Livia, Mary Namubiru, and Sara Spaziani. 2024. *Meet Your Future: Experimental Evidence on the Labor Market Effects of Mentors.* G^2LM|LIC Working Paper 87. IZA – Institute of Labor Economics.

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.

Anukriti, S, Catalina Herrera-Almanza, and Mahesh Karra. 2026. "Bring a Friend: Leveraging Financial and Peer Support to Improve Women's Reproductive Agency in India." *Journal of Development Economics* 180:103706.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2010. "Social Incentives in the Workplace." *Review of Economic Studies* 77 (2): 417–458.

Bandiera, Oriana, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2025. "The Search for Good Jobs: Evidence from a Six-Year Field Experiment in Uganda." *Journal of Labor Economics* 43 (3): 885–935.

Banerjee, Abhijit, and Sandra Sequeira. 2023. "Learning by Searching: Spatial Mismatches and Imperfect Information in Southern Labor Markets." *Journal of Development Economics* 164:103111.

Beaman, Lori. 2016. "Social Networks and the Labor Market." In *The Oxford Handbook of the Economics of Networks,* edited by Yann Bramoullé, Andrea Galeotti, and Brian W. Rogers, 648–672. Oxford University Press.

Beaman, Lori, and Jeremy Magruder. 2012. "Who Gets the Job Referral? Evidence from a Social Networks Experiment." *American Economic Review* 102 (7): 3574–3593.

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.

Bloom, Nick, and John Van Reenen. 2011. "Human Resource Management and Productivity." In *Handbook of Labor Economics,* edited by O. Ashenfelter and D. Card, 4:1697–1767. Elsevier.

Borker, Girija. 2021. *Safety First: Perceived Risk of Street Harassment and Educational Choices of Women.* Policy Research Working Paper 9731. World Bank.

Borusyak, Kirill, and Peter Hull. 2023. "Nonrandom Exposure to Exogenous Shocks." *Econometrica* 91 (6): 2155–2185.

Breza, Emily, Supreet Kaur, and Nandita Krishnaswamy. 2025. "Social Norms as a Determinant of Aggregate Labor Supply: Evidence from Decentralized Spot Markets." *Journal of Political Economy.*

Bryan, Gharad, Edward L. Glaeser, and Nick Tsivanidis. 2020. "Cities in the Developing World." *Annual Review of Economics* 12:273–297.

Burks, Stephen V, Bo Cowgill, Mitchell Hoffman, and Michael Housman. 2015. "The Value of Hiring through Employee Referrals." *Quarterly Journal of Economics* 130 (2): 805–839.

Caria, Stefano, Simon Franklin, and Marc Josef Witte. 2023. "Searching with Friends." *Journal of Labor Economics* 41 (4): 887–922.

Carranza, Eliana, Aletheia Donald, Florian Grosset-Touba, and Supreet Kaur. 2025. "The Social Tax: Redistributive Pressure and Labor Supply." *Econometrica* 93 (6): 2273–2308.

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

Cefalà, Luisa, Franck Irakoze, Pedro Naso, and Nicholas Swanson. 2025. *The Economic Consequences of Knowledge Hoarding.* Working Paper.

Chen, Yutong, Kerem Coşar, Devaki Ghose, Shirish Mahendru, and Sheetal Sekhri. 2026. "Gender-Specific Transportation Costs and Female Time Use: Evidence from India's Pink Slip Program." *Economic Journal.*

Chizeck, Seth, and Oluchi Mbonu. 2025. *The Role of the 'Fare' in Welfare: Public Transportation Subsidies and Their Effects on Low-Income Households.* Working Paper.

Christiaensen, Luc, and Patrick Premand, eds. 2017. *Côte d'Ivoire Jobs Diagnostic: Employment, Productivity, and Inclusion for Poverty Reduction.* Jobs Series 2. Report No. AUS13233. Washington, DC: World Bank.

Conwell, Lucas. 2026. *Privatized Provision of Public Transit.* Working Paper. Available at SSRN: https://ssrn.com/abstract=4489658. University College London.

Cooper, Russell, and Andrew John. 1988. "Coordinating Coordination Failures in Keynesian Models." *Quarterly Journal of Economics* 103 (3): 441–463.

Couture, Victor, Gilles Duranton, and Matthew A. Turner. 2018. "Speed." *Review of Economics and Statistics* 100 (4): 725–739.

Dasgupta, Aparajita, and Ashokankur Datta. 2023. *Gendered Transport Subsidy and Its Short Run Effect on Female Employment: Evidence from Delhi's Pink Pass Scheme.* Economics Discussion Paper 105. Ashoka University.

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

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.

Duranton, Gilles, and Diego Puga. 2020. "The Economics of Urban Density." *Journal of Economic Perspectives* 34 (3): 3–26.

Duranton, Gilles, and Matthew A. Turner. 2012. "Urban Growth and Transportation." *Review of Economic Studies* 79 (4): 1407–1440.

Ewarren Financial Services and DigitalTransport4Africa. 2023. *Abidjan Paratransit GTFS.* GTFS dataset, version of 26 July 2023, published by Data Transport. https://gitlab.com/digitaltransport/data/africa/abidjan_2/-/tree/main/Abidjan_2023.

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.

Franklin, Simon. 2018. "Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies." *Economic Journal* 128 (614): 2353–2379.

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." *Journal of Political Economy* 131 (3): 633–686.

Garlick, Robert, Erica Field, and Kate Vyborny. 2025. *Women's Mobility and Labor Supply: Experimental Evidence from Pakistan.* IZA Discussion Paper 17883. IZA.

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.

———. 2024. "The Value of Leisure Synchronization." *American Economic Journal: Applied Economics* 16 (1): 351–376.

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

———. 2003. "The Social Multiplier." *Journal of the European Economic Association* 1 (2-3): 345–353.

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

Goodburn, Charlotte, and Soumya Mishra. 2024. "Beyond the Dormitory Labour Regime: Comparing Chinese and Indian Workplace–Residence Systems as Strategies of Migrant Labour Control." *Work, Employment and Society* 38 (2): 505–526.

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

Hallberg, Daniel. 2003. "Synchronous Leisure, Jointness and Household Labor Supply." *Labour Economics* 10 (2): 185–203.

Hamermesh, Daniel S. 2002. "Timing, Togetherness and Time Windfalls." *Journal of Population Economics* 15 (4): 601–623.

Heblich, Stephan, Stephen J. Redding, and Daniel M. Sturm. 2020. "The Making of the Modern Metropolis: Evidence from London." *Quarterly Journal of Economics* 135 (4): 2059–2133.

Ho, Lisa, Suhani Jalota, and Anahita Karandikar. 2026. "Bringing Work Home: Flexible Work Arrangements as Gateway Jobs for Women in West Bengal." *Quarterly Journal of Economics* 141 (4): 3037–3094.

Jalota, Suhani, and Lisa Ho. 2024. *What Works for Her? How Work-from-Home Jobs Affect Female Labor Force Participation in Urban India.* SSRN Working Paper 4739387. SSRN.

Jungle Bus, Systra, and OpenStreetMap Côte d'Ivoire. 2021. *AbidjanTransport2021.* GTFS dataset exported from OpenStreetMap, version 202112291340. https://jungle-bus.github.io/AbidjanTransport2021/.

Kahneman, Daniel, and Alan B. Krueger. 2006. "Developments in the Measurement of Subjective Well-Being." *Journal of Economic Perspectives* 20 (1): 3–24.

Kiss, Andrea, Robert Garlick, Kate Orkin, and Lukas Hensel. 2023. *Jobseekers' Beliefs about Comparative Advantage and (Mis)Directed Search.* IZA Discussion Paper 16522. IZA.

Kling, Jeffrey R, Jeffrey B Liebman, and Lawrence F Katz. 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica* 75 (1): 83–119.

Kreindler, Gabriel, Arya Gaduh, Tilman Graff, Rema Hanna, and Benjamin A Olken. 2026. "Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta." *American Economic Review* 116 (9): 3330–3379.

Kumar, Ajay, and Fanny Barrett. 2008. *Stuck in Traffic: Urban Transport in Africa.* Technical report AICD Research Notes. World Bank.

Lall, Somik Vinay, J. Vernon Henderson, and Anthony J. Venables. 2017. *Africa's Cities: Opening Doors to the World.* Washington, DC: World Bank.

Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet. 2021. "Gender Differences in Job Search: Trading off Commute against Wage." *Quarterly Journal of Economics* 136 (1): 381–426.

Le Barbanchon, Thomas, Maddalena Ronchi, and Julien Sauvagnat. 2023. *Hiring Difficulties and Firm Growth.* CEPR Discussion Paper 17891. CEPR.

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

Marinescu, Ioana, and Roland Rathelot. 2018. "Mismatch Unemployment and the Geography of Job Search." *American Economic Journal: Macroeconomics* 10 (3): 42–70.

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

Mas, Alexandre, and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements." *American Economic Review* 107 (12): 3722–3759.

Mbonu, Oluchi, and Christopher Eaglin. 2026. *Market Segmentation and Coordination Costs: Evidence from Johannesburg's Minibus Networks.* Working Paper. Harvard University.

Minster, Clotilde V., Nabil Samir, and Jean-François Arvis. 2022. *When Public Transportation Is Unavailable, Firms Have Trouble Keeping Their Workers: Morocco Case Study.* World Bank Blog. Accessed August 15, 2023. https://blogs.worldbank.org/arabvoices/when-public-transportation-unavailable-firms-have-trouble-keeping-their-workers-morocco.

Moghadam, Hamed M, Patrick A Puhani, and Joanna Tyrowicz. 2024. "Pension Reforms and Couples' Labour Supply Decisions." *Labour Economics* 91:102627.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny. 1993. "Why Is Rent-Seeking So Costly to Growth?" *American Economic Review* 83 (2): 409–414.

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

Phillips, David C. 2014. "Getting to Work: Experimental Evidence on Job Search and Transportation Costs." *Labour Economics* 29:72–82.

Pun, Ngai, and Chris Smith. 2007. "Putting Transnational Labour Process in Its Place: The Dormitory Labour Regime in Post-Socialist China." *Work, Employment and Society* 21 (1): 27–45.

République de Côte d'Ivoire. 2015. *Schéma directeur des transports urbains du Grand Abidjan.* Rapport final, vol. 3, projet de développement du schéma directeur d'urbanisme du Grand Abidjan (SDUGA). JICA.

Rodrik, Dani. 2016. "Premature Deindustrialization." *Journal of Economic Growth* 21 (1): 1–33.

Sah, Raaj K. 1991. "Social Osmosis and Patterns of Crime." *Journal of Political Economy* 99 (6): 1272–1295.

Severen, Christopher. 2023. "Commuting, Labor, and Housing Market Effects of Mass Transportation: Welfare and Identification." *Review of Economics and Statistics* 105 (5): 1073–1091.

Shi, Peirong, Florian Grosset-Touba, and Émilien Schultz. 2025. *CanterburyCommuto: A Python Package to Compute and Visualize the Overlap of Route Pairs.* Python package, version 0.2.5. Also available on PyPI. https://github.com/PeirongShi/CanterburyCommuto.

Srivastava, Kartik. 2026. *Familiar Strangers: Evidence from Referral-Based Hiring Experiments in India.* Working Paper. Harvard University.

Storeygard, Adam. 2026. "Transport in Low- and Middle-Income Countries." *Regional Science and Urban Economics,* 104231.

Topa, Giorgio. 2019. "Social and Spatial Networks in Labour Markets." *Oxford Review of Economic Policy* 35 (4): 722–745.

Tsivanidis, Nick. 2026. "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio." *American Economic Review* 116 (2): 418–463.

Velasquez-Cabrera, Daniel. 2026. *Transit Infrastructure, Couples' Commuting Choices, and Gender Earnings Inequality.* Working Paper. Claremont McKenna College.

Vitali, Anna. 2023. *Consumer Search and Firm Location: Theory and Evidence from the Garment Sector in Uganda.* Working Paper. University College London.

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.

Witte, Marc J. 2025. *Why Do Workers Make Job Referrals? Experimental Evidence from Ethiopia.* IZA Discussion Paper 18258. IZA.

World Bank. 2018a. *Abidjan Urban Mobility Project (P167401): Concept Project Information Document and Integrated Safeguards Data Sheet.* Project Information Document PIDISDSC24943. Washington, DC: World Bank.

———. 2018b. *Cashew Value-Chain Competitiveness Project: Project Appraisal Document.* Project Appraisal Document PAD2121. Washington, DC: World Bank.

———. 2019. *Que la route soit bonne: Améliorer la mobilité urbaine à Abidjan.* Situation économique en Côte d'Ivoire, huitième édition. Washington, DC: World Bank.

———. 2020. *World Development Report 2020: Trading for Development in the Age of Global Value Chains.* Washington, DC: World Bank.

———. 2025. *Labor Force Participation Rate (Modeled ILO Estimates).* World Development Indicators, indicators SL.TLF.CACT.ZS (total), SL.TLF.CACT.MA.ZS (male), SL.TLF.CACT.FE.ZS (female), and SL.TLF.CACT.FM.ZS (female-to-male ratio).

Zárate, Román D. 2022. *Spatial Misallocation, Informality, and Transit Improvements: Evidence from Mexico City.* Policy Research Working Paper 9990. World Bank.

Appendix

This appendix contains the following supplementary materials:

- Appendix A: Labor Supply Patterns
- Appendix B: Commuting Together: Supplemental Material
  - Appendix B.1: Survey Protocols
  - Appendix B.2: Additional Exhibits
- Appendix C: Sales Experiment: Supplemental Material
  - Appendix C.1: Implementation Protocols
  - Appendix C.2: Labor Supply to Commute Elasticity
  - Appendix C.3: Co-Commuting Potential
  - Appendix C.4: Productivity
  - Appendix C.5: Amenity Value Calibration
  - Appendix C.6: Mechanisms
  - Appendix C.7: Gender
  - Appendix C.8: Additional Exhibits
- Appendix D: Factory Experiment: Supplemental Material
  - Appendix D.1: Implementation Protocols
  - Appendix D.2: Additional Exhibits

A Labor Supply Patterns

This section complements Figure 1, which relates labor supply at the factory to commuting time. Table A1 reports the corresponding regressions.

Figure 1 draws on two samples. Panel (a) uses the 67 focal jobseekers in the Control arm of the factory experiment, who were offered a job without network members. Their commuting time is the round trip to the factory they expected at baseline. The four who expected more than two hours are set to 2.05 hours. Panels (b)–(d) use the 685 surveyed workers of Appendix B.1. In these panels, commuting time and the outcome are both adjusted for hire-date fixed effects, with the sample mean added back, so that each panel compares workers hired on the same day. Panel (a) is not adjusted, since jobseekers have no hire date.

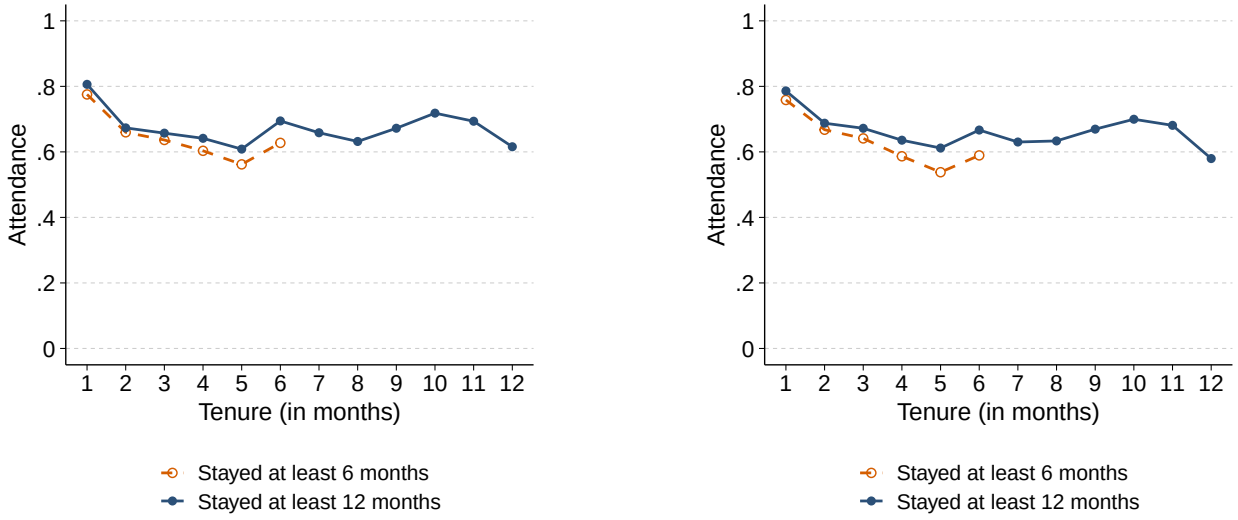
The remaining figures use the factory's personnel and production records to relate the margins of labor supply to one another. Figure A1 shows how attendance evolves with time on the job, Figure A2 relates productivity to attendance, Figure A3 relates retention to attendance in the first two weeks, and Figure A4 relates productivity to attendance in the first two weeks.

Table A1: Labor Supply and Commuting Time: Regressions

	Job take-up (1)	Retention (2)	Attendance (3)	Productivity (4)
Commuting time (round-trip hours)	-0.214	-0.0166	0.0107	0.0138
	(0.0940)	(0.00871)	(0.00501)	(0.0359)
	[0.026]	[0.058]	[0.033]	[0.702]
Mean of dependent variable	0.284	0.886	0.587	-0.114
Observations	67	542	628	264

Notes: Each column regresses the outcome of the corresponding panel of Figure 1 on round-trip commuting time in hours. Column (1): take-up among the jobseekers in the Control arm of the factory experiment. Column (2): retention is the share of weeks since hiring in which the worker remains employed, measured until 90 days before the end of the data, as for quits in Table 1. Column (3): all surveyed workers. Column (4): surveyed workers who appear in the factory's daily production records; productivity is in standard deviations of output per day worked. Columns (2)–(4) include hire-date fixed effects to account for seasonality and differences in potential retention. Inverse-propensity weights applied where relevant. Each cell reports the coefficient, the robust standard error, and the p-value in brackets.

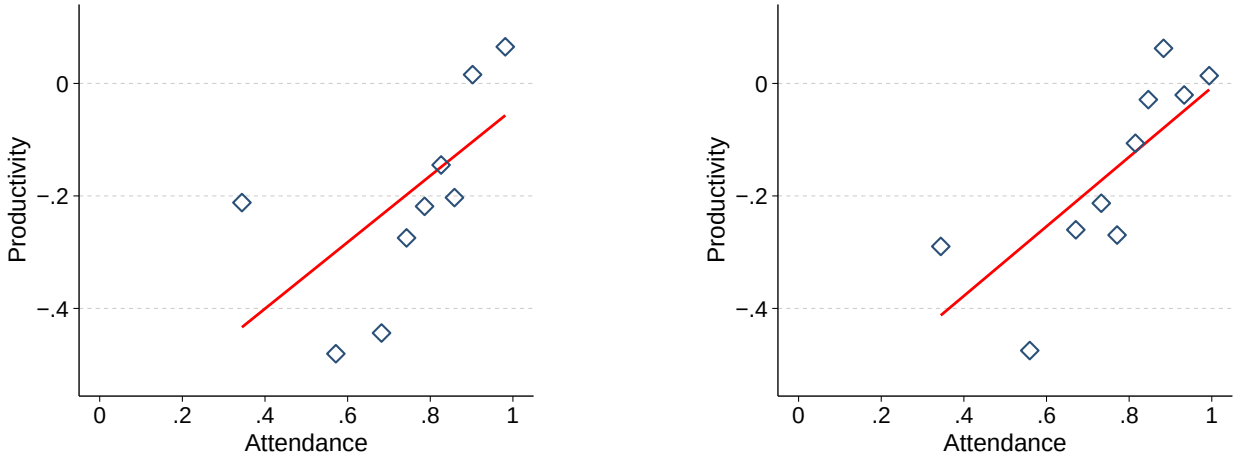
Figure A1: Attendance over time on the job



(a) All factory workers (b) Surveyed workers

Notes: Evolution of factory workers' attendance with time spent on the job. Panel A: sample of 828 workers in the personnel records who stayed at least six months on the job (orange dashed line, showing how their attendance evolved over these six months), 433 of whom stayed at least twelve (blue solid line, showing how their attendance evolved over these twelve months). Panel B: same as Panel A, for the subset of 491 of these workers (287 who stayed at least twelve months) who were surveyed as part of the network survey and were not part of the hiring experiment.

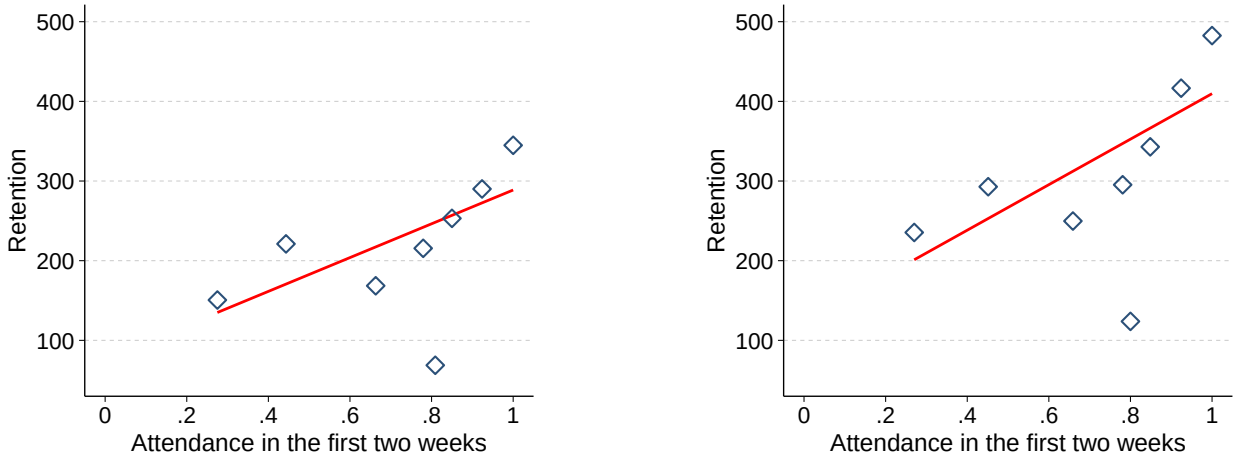
Figure A2: Attendance and Productivity



(a) All factory workers (b) Surveyed workers

Notes: Correlation between factory workers' average productivity and attendance. Productivity is daily output (quantity of cashew nuts peeled or graded each day), standardized at the production unit level and averaged at the worker level across all days of presence on the job. Attendance is the share of days in which the worker was present, from their first to last day at work. Each panel plots mean productivity within quantile bins of attendance (diamonds) and the fitted line of a linear regression of productivity on attendance. Panel A: sample of 409 workers with at least five days of presence at the factory during which production was recorded. Panel B: the subset of 266 of these workers who were surveyed as part of the network survey and were not part of the hiring experiment. The positive relationship between attendance and productivity is statistically significant in both panels ($p = 0.003$ in Panel A and $p = 0.004$ in Panel B).

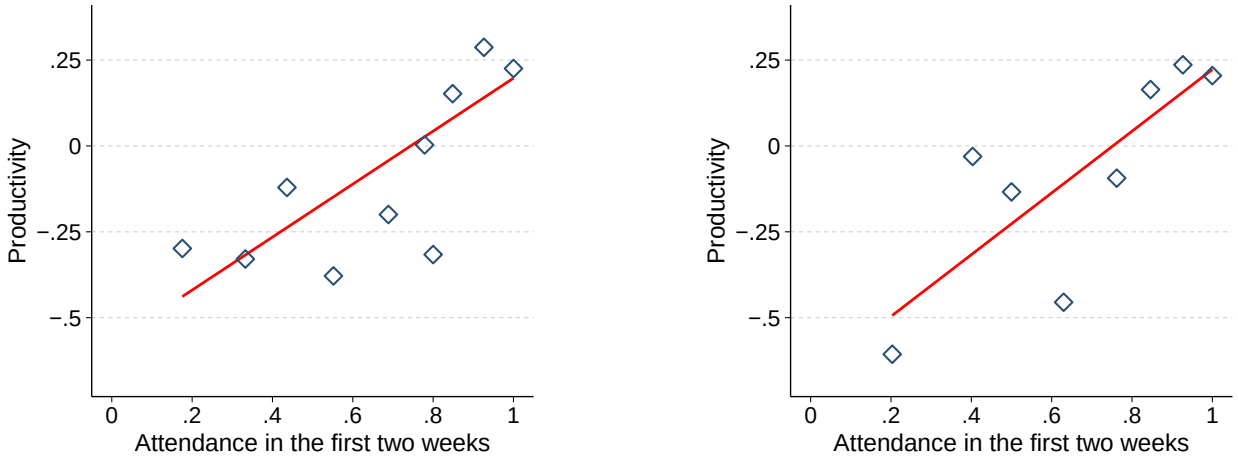
Figure A3: Early Attendance and Time on the Job



(a) All factory workers (b) Surveyed workers

Notes: Correlation between factory workers' retention (total number of days, between their first and last day on the job) and their early attendance (share of days present at work, in their first two weeks). Each panel plots mean retention within quantile bins of early attendance (diamonds) and the fitted line of a linear regression of retention on early attendance. Panel A: sample of all 1,713 workers in the personnel records who stayed on the job at least 14 days. Panel B: the subset of 651 of these workers who were surveyed as part of the network survey and were not part of the hiring experiment. The positive relationship between retention and early attendance is statistically significant, across both panels ($p < 0.001$).

Figure A4: Early Attendance and Productivity



(a) All factory workers (b) Surveyed workers

Notes: Correlation between factory workers' average productivity, as in Figure A2, and their early attendance, as in Figure A3. Each panel plots mean productivity within quantile bins of early attendance (diamonds) and the fitted line of a linear regression of productivity on early attendance. Panel A: sample of 357 workers with at least five days of presence at the factory during which production was recorded, and at least 14 days on the job. Panel B: the subset of 243 of these workers who were surveyed as part of the network survey and were not part of the hiring experiment. The positive relationship between early attendance and productivity is statistically significant, across both panels ($p < 0.001$).

B Commuting Together: Supplemental Material

B.1 Survey Protocols

A network survey was administered in two rounds to the factory workers employed at the time of each round, the first round from late February to April 2022 and the second round from late August to mid-September 2022. Each round targeted all the workers who appear in the factory's daily personnel records for more than three days over the four months before it. Because turnover at the factory is high, many workers were new by the second round: it surveyed the eligible workers who had not been interviewed in the first round, most of whom had joined the factory since November 2021. Workers interviewed in the first round were not surveyed again. Eligible workers were called in a random order to set up an interview. Since the personnel records include the workers hired through the factory experiment who took up the job, 43 of the respondents had been hired through the experiment, as focal jobseekers or as network members. We exclude them from the analysis for simplicity; the results are similar when they are included. The resulting analysis sample has 685 workers, 498 from the first round and 187 from the second.

The survey elicits the roster of co-workers with whom the respondent commutes, the reasons for commuting together in ranked order, whether the respondent knew each co-commuter before starting at the factory, and whether they interact at work or outside work. Respondents also report the borough and neighborhood in which they live, and how long it usually takes them to get from home to the factory, one way, which we double to obtain the round-trip commute. The 685 respondents live in 201 different neighborhoods of the city. Figure A5 maps where they live and the route they could take to the factory.

We also collect detailed personnel records from the factory. These records run from 3 August 2020 to 23 July 2022, so every focal jobseeker hired through the experiment is observed for at least a year after their first day at work. They include daily attendance data for the universe of workers, which can also be used to measure quits, and the production unit in which each worker works each day; we use the worker's main unit in each month, which takes 12 values in the sample of Table 1. Additionally, we obtain daily individual production records. These records start on 21 February 2022, run until July 2022, and cover two production units, peeling and grading. This comes too late to be used for the factory experiment, but covers 296 of the 685 surveyed workers, all of them women. The two units record output in different units, so we standardize each day's output by its mean and standard deviation over the worker-days of the same unit, then average it over the worker's days at work. Production is individual and directly depends on effort, so can be taken as a measure of individual daily

productivity. We then match the network links reported in the survey to these personnel records.

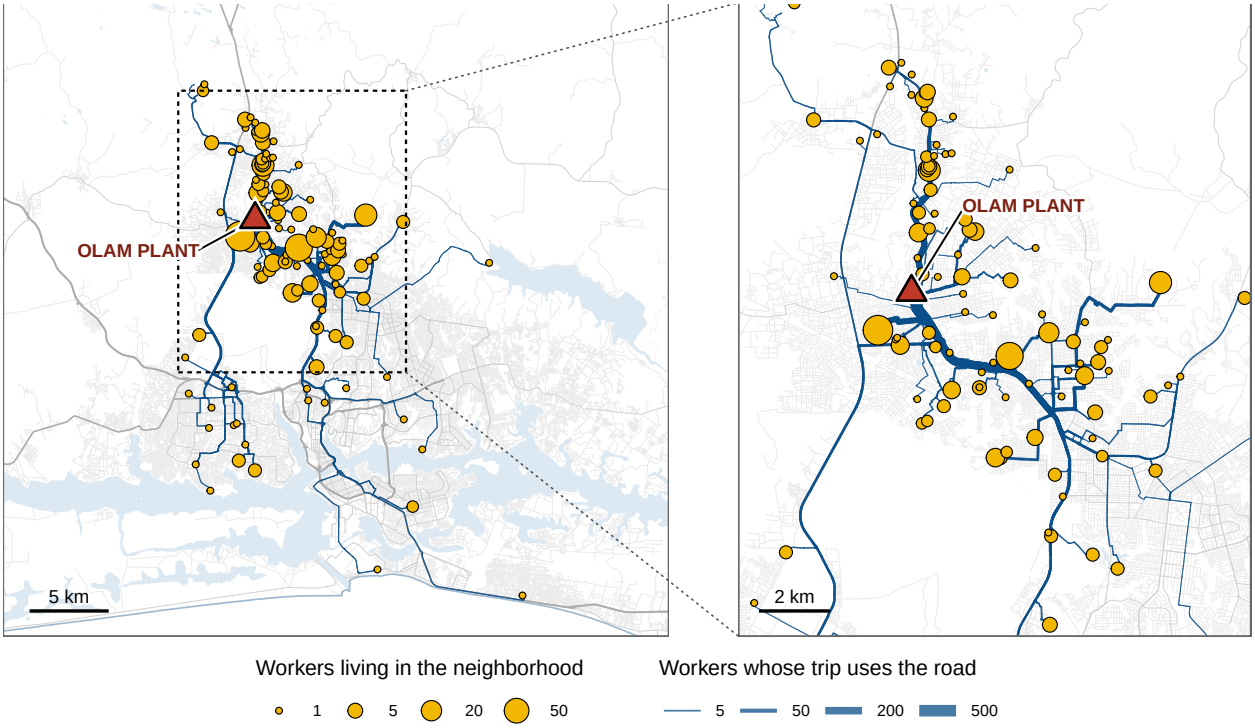
Table A2 compares the network-survey respondents who were not part of the hiring experiment with the focal jobseekers of the factory experiment at baseline.

Table A2: Network Survey and Factory Experiment Samples

	Network survey		Factory experiment	
Variable	Mean/SD	N	Mean/SD	N
	(1)	(2)	(3)	(4)
Is a woman	0.83	685	0.72	163
	(0.37)		(0.45)	
Age	33	685	29	163
	(7)		(6)	
Years of education	8	685	8	163
	(5)		(6)	
Has a partner	0.47	685	0.52	163
	(0.50)		(0.50)	
Any childcare	0.72	685	0.35	163
	(0.45)		(0.48)	

Notes: Cols. (1)–(2): the 685 factory workers in the network survey. Cols. (3)–(4): the 163 focal jobseekers of the factory experiment, at baseline. Cols. (1) and (3) report the mean and, in parentheses, the standard deviation. *Any childcare* equals 1 if the respondent has at least one child under their care (network survey) or spent any time on childcare in the past seven days (factory experiment). Inverse-propensity weights applied where relevant.

Figure A5: Home Locations and Routes to the Factory, Network Survey



Notes: Where the factory workers of Table A2 live, and the public-transport trips they would take to the factory. Each circle is a neighborhood, drawn at its center and sized by the number of workers living there. Trips are the fastest public-transport trip from each neighborhood to the factory leaving at 7:00 on a weekday, traced with GraphHopper on the transit lines of Appendix C.3; a road is drawn wider the more workers' trips use it. The right map enlarges the dashed area of the left map. Map data: © OpenStreetMap contributors (2022); transit lines: Jungle Bus (2021), Ewarren Financial Services and DigitalTransport4Africa (2023).

B.2 Additional Exhibits

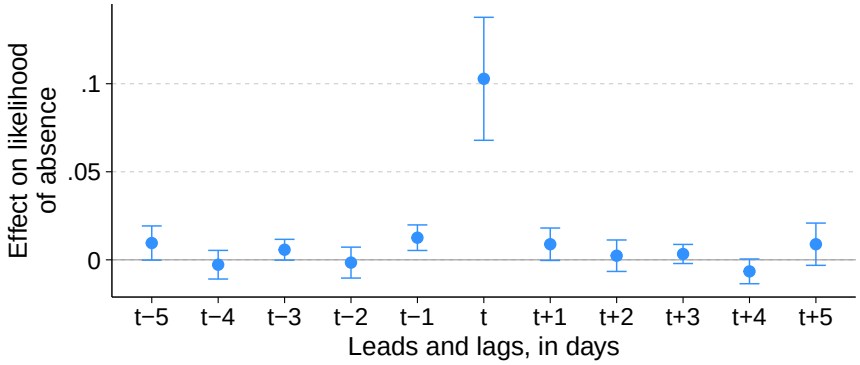
Figure A6 asks whether the coordination of absences in Table 1 happens on the same day or reflects absences that move together over several days. On worker-days, we estimate

$$\text{Absent}_{it} = \sum_{k=-5}^{5} \beta_k \, A_{i,t+k} + \lambda_{\ell(i,t)} + \delta_{u(i,t)t} + \eta_{n(i)t} + \varepsilon_{it}, \tag{A1}$$

where Absent_{it} and N_{it} are as in Equation (1), and $A_{i,t+k} = N_{it} - P_{i,t+k}$, with $P_{i,t+k}$ the number of i's co-commuters present at the factory k calendar days later ($|k|$ days earlier when $k < 0$). For $k = 0$, A_{it} is the regressor of Equation (1): the number of co-commuters absent on day t. The fixed effects are those of Column (4) of Table 1: link fixed effects $\lambda_{\ell(i,t)}$, as in Equation (1), and day fixed effects interacted with the worker's production unit, $\delta_{u(i,t)t}$, and with the worker's neighborhood, $\eta_{n(i)t}$. The sample is the worker-days of Table 1 on which all ten leads and lags are observed. Standard errors are clustered at the worker level, and the figure plots the β_k with their 90% confidence intervals.

The coefficient β_0 measures same-day coordination, holding fixed the co-commuters' absences on the surrounding days. The lags ($k < 0$) ask whether a co-commuter's absence in the previous days still predicts the worker's absence today, as it would if the two missed work in shared spells. The leads ($k > 0$) ask whether the worker's absence today predicts the co-commuters' absences in the coming days. A slow-moving shock common to the pair would show up in both the leads and the lags; coordination that happens day by day, through the joint trip, shows up in β_0 alone.

Figure A6: Dynamics of Co-Commuter Coordination on Absences



Notes: Event-study estimates of the effect of co-commuters' absences on the focal worker's likelihood of being absent, by day relative to the co-commuter's absence (t), from five days before to five days after. The regression uses the fixed-effect structure of Column (4) of Table 1. The sample is non-experiment factory workers with at least one co-commuter observed on the same day. Whiskers are 90% confidence intervals. Standard errors are clustered at the focal-worker level.

Table A3 decomposes the coordination of absences of Table 1 by friendship and by whether

the two workers knew each other before the factory. Figure A7 shows the ranking of reasons for commuting together of Figure 3 separately for women and men. Table A4 estimates Column (1) of Table 1, for absences and quits, separately for women and men.

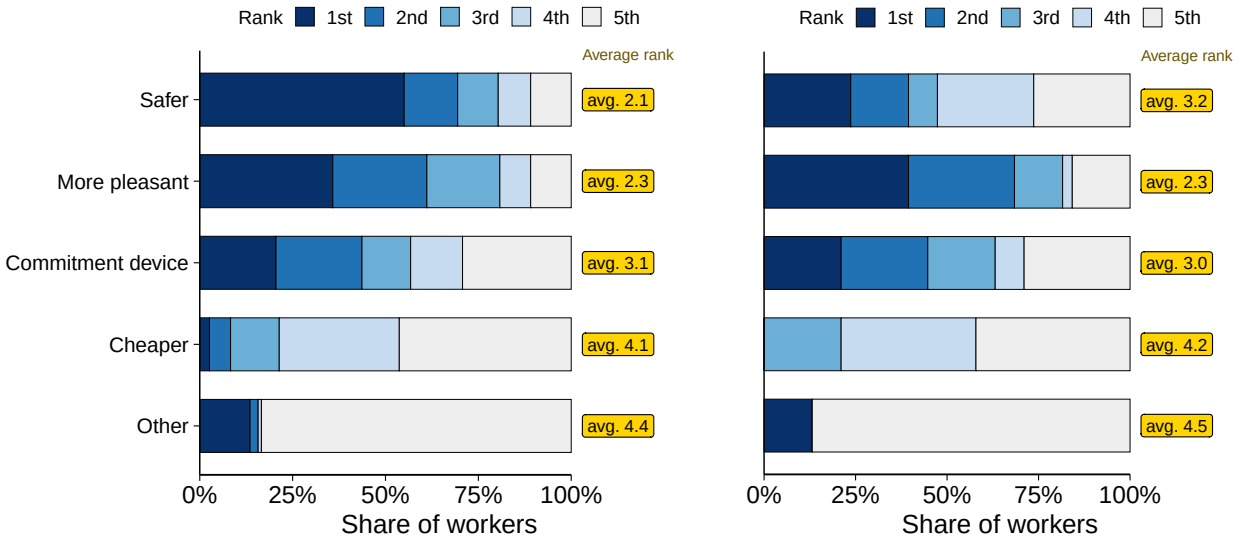
Table A3: Coordination of Absences Among Co-Commuters, by Friendship and by Pre-Factory Acquaintance

	(1)	(2)	(3)	(4)
Panel A. By friendship				
Nbr co-commuters absent who are:				
friends	0.365***	0.344***	0.186***	0.160***
	(0.030)	(0.043)	(0.026)	(0.038)
	[0.000]	[0.000]	[0.000]	[0.000]
non-friends	0.271***	0.253***	0.119***	0.092***
	(0.030)	(0.037)	(0.021)	(0.018)
	[0.000]	[0.000]	[0.000]	[0.000]
$p(\beta_F = \beta_N)$	0.007	0.051	0.004	0.024
Panel B. By pre-factory acquaintance				
Nbr co-commuters absent who are:				
known before the factory	0.277***	0.252***	0.138***	0.108***
	(0.036)	(0.045)	(0.023)	(0.027)
	[0.000]	[0.000]	[0.000]	[0.000]
met at the factory	0.317***	0.304***	0.147***	0.124***
	(0.030)	(0.035)	(0.023)	(0.028)
	[0.000]	[0.000]	[0.000]	[0.000]
$p(\beta_K = \beta_M)$	0.314	0.219	0.724	0.590
Day-neighborhood FE	No	Yes	No	Yes
Day-prod. unit FE	No	No	Yes	Yes

Notes: The sample, outcome, and fixed-effect specifications are as in Table 1. The single regressor of that table is decomposed into two, both entered in the same regression, which together partition the active co-commuters on each worker-day. Panel A splits co-commuters by friendship, where a co-commuter is classified as a friend if the focal worker also reports them in the out-of-work socialization module of the network survey. Panel B splits them by whether the focal worker knew them before working at the factory. The equality rows report cluster-robust Wald p-values. The absence rate on worker-days when neither regressor is positive is 0.13. Standard errors are clustered at the focal-worker level, with p-values in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A7: Reasons for Commuting Together, by Gender

(a) Women (b) Men



Notes: Stated main benefits from commuting to the factory with co-workers. As in Figure 3, separately by gender.

Table A4: Coordination of Absences and Quits Among Co-Commuters, by Gender

	Absences (daily)			Quits (weekly)		
	All	Women	Men	All	Women	Men
	(1)	(2)	(3)	(4)	(5)	(6)
Nbr Co-Commuters absent (quitting)	0.307***	0.311***	0.210***	0.153***	0.183***	-0.026
	(0.026)	(0.030)	(0.049)	(0.039)	(0.045)	(0.016)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.127]
Mean: Covar = 0	0.13	0.14	0.11	0.006	0.006	0.003
N: Worker-days (weeks)	66,936	55,721	11,174	8,423	7,078	1,339
N: Workers	239	206	33	204	171	33

Notes: Estimated coordination of absences and quits among co-commuters, separately by the worker's gender. Column (1) reproduces Column (1) of Table 1, Panel A, while column (4) reproduces Column (1) of Table 1, Panel B. Columns (2)-(3) and (5)-(6) show the effects separately by whether the survey respondent is a woman or a man. Standard errors clustered at the worker level in parentheses; p-values in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Sales Experiment: Supplemental Material

C.1 Implementation Protocols

Hiring process. The canvassing campaign was subcontracted to Atom-BTL Marketing, a local marketing agency with experience working with large formal firms in Côte d'Ivoire, which hired the workers and managed the work.

The experiment is implemented within the firm's usual hiring process. For its usual canvassing campaigns, it assembles a list of candidates, calls them, describes the specific openings, elicits interest, and invites those who are interested to a half-day onboarding session before their first day of work.

We insert two steps into this process, shown in Figure A8. First, we obtain the candidate list from the firm before it calls, and Innovations for Poverty Action administers a short baseline survey by phone. That survey provides information about the job, elicits the focal jobseeker's network, collects basic socio-economic characteristics, announces the treatment assignment, and elicits interest in the job. More information about these steps is provided below. Second, we return to the firm a list containing the eligible and interested focal jobseekers together with the network members of focal jobseekers assigned to a treated arm. The firm proceeds from there as it normally would.

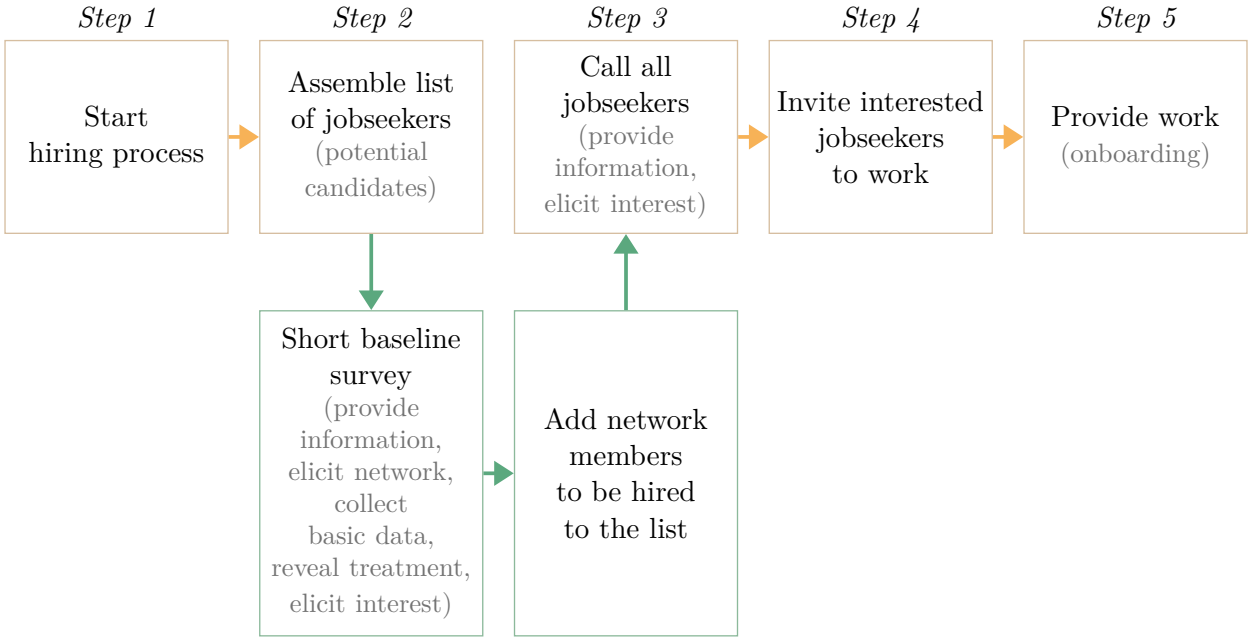
To maintain the implementation as natural as possible, the network members are not administered a survey. Those selected to receive an offer are directly called by the firm, to receive details about their offer and elicit their interest. When the firm receives the roster of individuals to hire (the eligible focal jobseekers and the selected network members), they are not informed of who is a focal jobseeker or a network member, nor of who referred whom. This is designed to minimize the scope for differential treatment from the managers.

The firms attempted at least three calls in a single day to reach each candidate. Candidates who could not be reached were not guaranteed an offer.

Before work begins there is a compulsory half-day of training at Atom-BTL's offices in Yopougon. Work then runs in one of the two weeks following the training, and focal jobseekers are told at baseline that they must be available for both weeks. The focal jobseekers were always scheduled in the same week as their network members. Consistently, when a focal jobseeker and a network member both come to work, 91% work in the same week.

Table A5 reports the treatment effects at each step of this process. In the first steps, almost all focal jobseekers say they are interested, even in the Control arm: 96% say so at baseline,

Figure A8: Implementation embedded in the natural hiring 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.

86% are reached by the firm's call, and 81% confirm their interest on that call. Since almost everyone says yes, neither treatment arm changes these steps (all effects below 3 percentage points, $p > 0.5$). The first step that requires a real decision is showing up for the training, and only 41% of Control focal jobseekers do. This is where the Same Site effect appears: 10.6 percentage points ($p = 0.030$). The effect carries over to job take-up (12.4 points, $p = 0.008$) and to the number of days worked (0.49 days, $p = 0.024$). The Different Site arm has no effect at any step.

Baseline survey. The survey is kept short and only collects minimal information to maintain the process as natural as possible, and avoid prompting the surveyed jobseekers to overly consider a specific aspect of the job being offered.

Network members are elicited from each of the surveyed jobseekers before the treatment assignment is revealed, so the elicited referrals are expected to be similar across treatment arms. The enumerators explain that the firm is looking to hire a large number of workers, and ask whether the respondent knows others who would be suitable and available for a similar position at the same firm. They first ask how many such people live in the respondent's own neighborhood and record their names, then do the same for people living elsewhere.

Respondents can give as many names as they wish; the enumerators suggest three to five. Named individuals must be at least 18 years old to be eligible for an offer. The individuals being referred, at the hiring stage, form the network members of the surveyed focal jobseekers. Focal jobseekers who name no referral are not eligible to participate in the study, since there is then no network member to whom a job could be offered.

Up to three network members can receive an offer, per focal jobseeker. Where a focal jobseeker refers three people or fewer, all of them receive an offer. Where more people are referred, which ones are selected to receive an offer is randomized. This was designed to provide exogenous variation in the type of network members receiving an offer, to test for heterogeneity in treatment effects along this dimension. In practice however, 91% of the focal jobseekers provide at most three names, so all available identifying variation comes from who was referred rather than this random election.

Treatment is announced to the focal jobseekers later in the same survey, after the referrals have been recorded. In Control, no network member receives an offer. In the Different Site treatment arm, the network members are assigned to the other worksite. In the Same Site treatment arm, they are assigned to the focal jobseeker's own worksite, hours, and team, so the two would see each other during the work day.

Focal jobseekers assigned to a treated arm learn during the survey which of their network members will receive an offer, whether those members would work at the same site or shift as them, and their own work hours and location. Focal jobseekers assigned to Control learn their own work hours and location.

Timeline. Hiring and work ran in weekly cohorts: each week, a new group of focal jobseekers was surveyed and offered the job, trained on a Friday, and worked the following week. Two pilot rounds came first, with training on 8 and 15 July 2022; their focal jobseekers are not part of our sample. The first wave then ran over 8 weeks, with training dates from 22 July to 16 September 2022, and accounts for 585 of the 848 focal jobseekers in our sample. The second wave ran over 3 weeks, with training from 11 to 23 November 2022, and accounts for the other 263. The job lasted five days of work in the first wave and four in the second.

Randomization. We randomized focal jobseekers within strata. Comparing similar people across treatment arms increases statistical power; we kept each stratum large enough to leave enough variation in treatment within it to conduct randomization inference. Focal jobseekers were randomized in batches, each time the firm added new candidates to its list: 11 batches in the first wave, from July 2022, and 5 in the second, from November 2022. Within a batch,

Table A5: Treatment Effects at Each Step of the Hiring Process, Sales Experiment

	Interested at baseline (1)	Reached by call (2)	Accepted on call (3)	Attended training (4)	Job take-up (5)	Days worked (6)
Different Site	0.00759	0.0133	-0.0250	0.0322	0.0294	0.0789
	(0.0191)	(0.0336)	(0.0405)	(0.0486)	(0.0463)	(0.213)
	[0.692]	[0.692]	[0.537]	[0.508]	[0.526]	[0.711]
Same Site	0.00317	0.00460	-0.0198	0.106	0.124	0.490
	(0.0196)	(0.0340)	(0.0402)	(0.0485)	(0.0469)	(0.216)
	[0.872]	[0.892]	[0.623]	[0.030]	[0.008]	[0.024]
p(Same Site = Different Site)	0.802	0.774	0.889	0.107	0.033	0.044
Control mean	0.956	0.864	0.807	0.409	0.329	1.52
Observations	677	677	677	677	677	677

Notes: Same specification and sample as column (1) of Table 2, for each step of the hiring process. Columns (2)–(6) equal zero for focal jobseekers who did not reach the step. The firm called every focal jobseeker interested at baseline; column (2) is whether the call reached them. Inverse-propensity weights applied where relevant. Robust standard errors in parentheses; p-values in brackets.

the strata are defined by gender, municipality of residence, and age (above or below the batch median), merging strata with fewer than 40 people. Within each stratum, each focal jobseeker was randomly assigned a treatment arm and a supervisor, and each supervisor worked at one of the two worksites of the week. Three of the five supervisors worked at one worksite and two at the other, so the worksite was assigned with probabilities 0.6 and 0.4, independently of the treatment arm.

The probability of each treatment arm differed between the two waves. In the first wave, the three arms had the same probability; in the second wave, Control had probability 1/5 and each treated arm 2/5. A focal jobseeker in the second wave was therefore more likely to be treated than one in the first, and the two waves may differ for other reasons, such as the time of year. Comparing arms without accounting for this would mix the effect of the treatments with differences between the waves. We therefore weight each focal jobseeker by the inverse of the probability of the arm they were assigned to, in their wave. With these weights, each wave contributes to every arm in proportion to its size, so the arms are compared on the same mix of waves.

Each week, the firm moved to a new pair of worksites. Focal jobseekers who had been randomized but had not yet answered the baseline call kept their treatment arm and supervisor, and their worksite was renamed to a site of the new pair. Which renamings applied to a person depended only on when they answered the baseline call, not on the supervisor they had been assigned. Enumerators did not know a focal jobseeker's treatment arm before

the call: the survey shows it only at the end, with the offer, so whether and when a focal jobseeker answered the call cannot depend on their arm. Following each person's renamings, we compute their probability of each worksite: 0.6 or 0.4 for 841 of the 848 focal jobseekers, and 1 for the other 7. These probabilities give the expected commute used as a control in Table 3. Finally, the first batch of the second wave re-randomized the first-wave candidates who had not answered the call or had said they were not yet available.

Sample. Table A6 reports baseline characteristics by treatment arm, and Table A7 describes the network members elicited at baseline.

Table A6: Baseline Balance, Sales Experiment

Variable	Control	Different Site		Same Site		
	Mean/SD	Mean/SD	P-value	Mean/SD	P-value	N
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Focal Jobseeker Characteristics						
Is a woman	0.39	0.40	0.840	0.33	0.185	848
	(0.49)	(0.49)		(0.47)		
Age	25	25	0.814	25	0.991	848
	(5)	(4)		(5)		
Years of education	14	14	0.913	14	0.641	848
	(2)	(2)		(2)		
Has a partner	0.09	0.10	0.527	0.09	0.892	848
	(0.28)	(0.30)		(0.28)		
Has young children	0.19	0.17	0.615	0.16	0.426	848
	(0.39)	(0.38)		(0.37)		
Panel B: Focal Jobseeker Baseline Income-Generating Activities						
Any IGA, past 3 months	0.63	0.65	0.792	0.59	0.320	848
	(0.48)	(0.48)		(0.49)		
Any earnings, past 30 days	0.42	0.41	0.833	0.37	0.269	848
	(0.49)	(0.49)		(0.48)		
Earnings, past 30 days (000s FCFA)	47.2	45.5	0.858	33.1	0.067	848
	(101.6)	(107.9)		(58.0)		
Any earnings, past 7 days	0.24	0.20	0.299	0.20	0.230	848
	(0.43)	(0.40)		(0.40)		
Panel C: Focal Jobseeker Baseline Commute						
Any commute	0.53	0.52	0.890	0.53	0.906	848
	(0.50)	(0.50)		(0.50)		
Commuting time (hours)	1.39	1.47	0.537	1.50	0.457	448
	(1.15)	(1.13)		(1.47)		
Commute with people they know	0.52	0.46	0.308	0.53	0.827	448
	(0.50)	(0.50)		(0.50)		

Notes: Baseline covariates of focal jobseekers by treatment arm. Cols. (1), (2), and (4) report the mean and, in parentheses, the standard deviation. Cols. (3) and (5) report the p-value on the corresponding treatment indicator from an OLS regression of the covariate on both treatment indicators, with robust standard errors. *Has young children* equals 1 if the respondent answers one or more to "How many children under 5 do you have?". Inverse-propensity weights applied where relevant. Exchange rate in December 2022: 620 FCFA = 1 USD.

Table A7: Elicited Network Members, Sales Experiment

Variable	Mean (1)	SD (2)	N (3)
Panel A: Number of Network Members			
Number of network members	1.83	1.19	848
Named more than 3 network members	0.09	0.29	848
Panel B: Relationship to the Focal Jobseeker			
Share: Friends	0.70	0.41	848
Share: Acquaintances	0.14	0.31	848
Share: Household members	0.09	0.26	848
Share: Siblings	0.07	0.23	848
Share: Cousins	0.07	0.22	848
Share: Parents	0.00	0.05	848
Share: Other relationship	0.02	0.12	848
Share: From same neighborhood	0.69	0.44	848
Share: Could commute together	0.88	0.30	848
Panel C: Similarity to the Focal Jobseeker			
Share: Men (among focal men)	0.82	0.33	534
Share: Women (among focal women)	0.73	0.40	314
Similar age (within 2.5 years)	0.61	0.49	848
Similar education (within 1 year)	0.60	0.49	830

Notes: Network members' names and characteristics elicited from focal jobseekers at baseline. One observation per focal jobseeker; shares are computed over all the network members referred by the focal jobseeker. *Could commute together*: the focal jobseeker says the two of them could make the trip together if they had to be in the same place at the same time. Similar education is missing when the focal jobseeker did not know any member's education. Inverse-propensity weights applied where relevant.

C.2 Labor Supply to Commute Elasticity

In the factory experiment, take-up in the Control arm falls with the expected commute (Figure 1). In the sales experiment it does not: Control take-up is flat in the commute to the assigned worksite (Table 3). The two jobs differ in where training takes place. At the factory, training and work happen at the same site. In the sales job, every new worker first attends an unpaid day of training at the firm's offices in Yopougon, and only then works at the assigned worksite.

The baseline survey does not ask about the trip to the training. We proxy it with the round trip to the firm's Yopougon worksite, 2.3 km from the training site, reported at baseline by the other focal jobseekers living at the same home location, or at the nearest home location with such answers. Table A8 relates training attendance (columns 1–3) and job take-up (columns 4–6) to this training trip, to the commute to the assigned worksite, and to both. In the Control arm, a one-hour longer training trip lowers training attendance by 6.1 percentage points ($p = 0.055$), while the commute to the worksite has no significant effect. The two treatment arms offset different trips. In the Different Site arm, the training-trip penalty disappears (interaction 10.8 points, $p = 0.009$), while the interaction with the commute to work, which network members cannot share, is small and not significant. In the Same Site arm, the reverse holds: attendance rises with the commute to work (interaction 9.7 points, $p = 0.011$), and the training trip matters little. Job take-up shows the same pattern, and both interactions keep their size when the two trips enter together (columns 3 and 6).

Two points help read these patterns. First, the Control focal jobseekers who still attend the training despite a long trip are likely those least sensitive to commuting, which would explain why, once trained, their take-up responds little to the commute to the assigned worksite. Second, network members can travel together to the single training site in both treatment arms, yet the training trip matters less in the Same Site arm than in the Different Site arm. One explanation is that the training trip is made once, on the training day, while the commute to work is made on each day of work: for Same Site focal jobseekers, who can share that daily commute, it weighs more in the decision than the trip to the training.

Table A9 repeats these regressions with the cost of the trips instead of their duration. We build the cost of the training trip in the same way, from the round-trip costs that the same respondents report for the Yopougon worksite (median 1,000 FCFA). In the Control arm, attendance falls by 17.6 percentage points for each additional 1,000 FCFA ($p = 0.001$), and the Different Site arm offsets this decline (19.7 points, $p = 0.009$). The cost of the trip to

the assigned worksite does not affect attendance.

Table A8: Training Attendance and Job Take-up by Training Trip and Commute Time, Sales Experiment

	Attended training			Job take-up		
	(1)	(2)	(3)	(4)	(5)	(6)
Training trip (hours)	-0.0612		-0.0585	-0.0311		-0.0304
	(0.0319)		(0.0319)	(0.0308)		(0.0308)
	[0.055]		[0.067]	[0.313]		[0.324]
Commute to work (hours)		-0.0423	-0.0381		-0.0112	-0.00909
		(0.0299)	(0.0302)		(0.0283)	(0.0281)
		[0.158]	[0.207]		[0.691]	[0.746]
Different Site	-0.180	-0.0469	-0.233	-0.128	-0.0419	-0.175
	(0.0957)	(0.0966)	(0.124)	(0.0930)	(0.0915)	(0.118)
	[0.060]	[0.628]	[0.059]	[0.169]	[0.647]	[0.139]
Different Site × Training trip	0.108		0.107	0.0813		0.0783
	(0.0411)		(0.0414)	(0.0403)		(0.0406)
	[0.009]		[0.010]	[0.044]		[0.055]
Different Site × Commute to work		0.0387	0.0276		0.0335	0.0247
		(0.0392)	(0.0392)		(0.0377)	(0.0377)
		[0.324]	[0.482]		[0.374]	[0.512]
Same Site	0.0487	-0.0982	-0.138	0.0960	-0.0270	-0.0416
	(0.0973)	(0.0949)	(0.125)	(0.0955)	(0.0910)	(0.121)
	[0.617]	[0.301]	[0.273]	[0.315]	[0.767]	[0.732]
Same Site × Training trip	0.0258		0.0163	0.0129		0.00506
	(0.0417)		(0.0417)	(0.0409)		(0.0409)
	[0.536]		[0.696]	[0.753]		[0.902]
Same Site × Commute to work		0.0969	0.0974		0.0712	0.0718
		(0.0380)	(0.0380)		(0.0370)	(0.0367)
		[0.011]	[0.011]		[0.055]	[0.051]
Control mean	0.409	0.409	0.409	0.329	0.329	0.329
Observations	677	677	677	677	677	677

Notes: Training trip: mean round trip to the firm's Yopougon worksite, 2.3 km from the training site, reported at baseline by the other focal jobseekers living at the same home location (or the nearest one with such answers). Commute to work: the focal jobseeker's own reported round trip to the assigned worksite. Sample of Table 3. For the 3 focal jobseekers without this measure, we set the training trip to the sample mean and control for an indicator. Inverse-propensity weights applied where relevant. Robust standard errors in parentheses; p-values in brackets.

Table A9: Training Attendance and Job Take-up by Training Trip and Commute Cost, Sales Experiment

	Attended training			Job take-up		
	(1)	(2)	(3)	(4)	(5)	(6)
Training trip (1,000 FCFA)	-0.176		-0.176	-0.0939		-0.0938
	(0.0513)		(0.0513)	(0.0502)		(0.0503)
	[0.001]		[0.001]	[0.062]		[0.062]
Commute to work (1,000 FCFA)		-0.0197	-0.0191		-0.00822	-0.00790
		(0.0500)	(0.0482)		(0.0469)	(0.0465)
		[0.695]	[0.693]		[0.861]	[0.865]
Different Site	-0.201	0.109	-0.125	-0.102	0.0906	-0.0415
	(0.101)	(0.0973)	(0.128)	(0.0978)	(0.0930)	(0.125)
	[0.047]	[0.262]	[0.330]	[0.299]	[0.331]	[0.741]
Different Site × Training trip	0.197		0.201	0.111		0.114
	(0.0749)		(0.0758)	(0.0726)		(0.0731)
	[0.009]		[0.008]	[0.126]		[0.119]
Different Site × Commute to work		-0.0555	-0.0582		-0.0445	-0.0462
		(0.0626)	(0.0613)		(0.0600)	(0.0598)
		[0.376]	[0.343]		[0.458]	[0.440]
Same Site	-0.00928	0.0528	-0.0607	0.0664	0.0822	0.0231
	(0.103)	(0.0839)	(0.122)	(0.101)	(0.0797)	(0.120)
	[0.928]	[0.529]	[0.619]	[0.510]	[0.303]	[0.848]
Same Site × Training trip	0.0921		0.0943	0.0463		0.0490
	(0.0743)		(0.0743)	(0.0733)		(0.0734)
	[0.215]		[0.205]	[0.528]		[0.504]
Same Site × Commute to work		0.0391	0.0364		0.0307	0.0292
		(0.0516)	(0.0498)		(0.0486)	(0.0481)
		[0.449]	[0.465]		[0.528]	[0.545]
Control mean	0.409	0.409	0.409	0.329	0.329	0.329
Observations	677	677	677	677	677	677

Notes: Same regressions and sample as Table A8, with the reported round-trip costs of the two trips, in 1,000 FCFA, instead of their duration. Inverse-propensity weights applied where relevant. Robust standard errors in parentheses; p-values in brackets.

C.3 Co-Commuting Potential

Ideally, assigning a focal jobseeker and their network members to different worksites would prevent them from commuting to work together. In practice, two people traveling to different worksites can still commute together part of the way, for instance if they leave the same neighborhood and go in the same direction. To measure how much of her commute a focal jobseeker could make together with their network members, we take the following steps. We trace public-transport trips with GraphHopper, an open-source routing engine, on Abidjan's transit lines and the OpenStreetMap street network of Côte d'Ivoire as of 2022, the year of the experiment. We compute the measure with CanterburyCommuto, an open-source Python package we developed for this purpose (Shi et al. 2025). The package is fully available online on GitHub and PyPI, and the replication package includes the exact version we use (0.2.5). Figure A9 illustrates the measure for two example pairs.

1. We locate homes and worksites.

   (a) Each focal jobseeker reports the neighborhood they live in at baseline. We assign them the coordinates of the center of that neighborhood.

   (b) Network members who are neighbors of the focal jobseeker are assigned the focal jobseeker's home location. For other network members, we use the neighborhood the focal jobseeker reports for them at baseline.

   (c) Each focal jobseeker can be assigned to either of two candidate worksites. In the Different Site treatment arm, the focal jobseeker is assigned to one worksite and their network members to the other (Section 3.2). We use the coordinates of both.

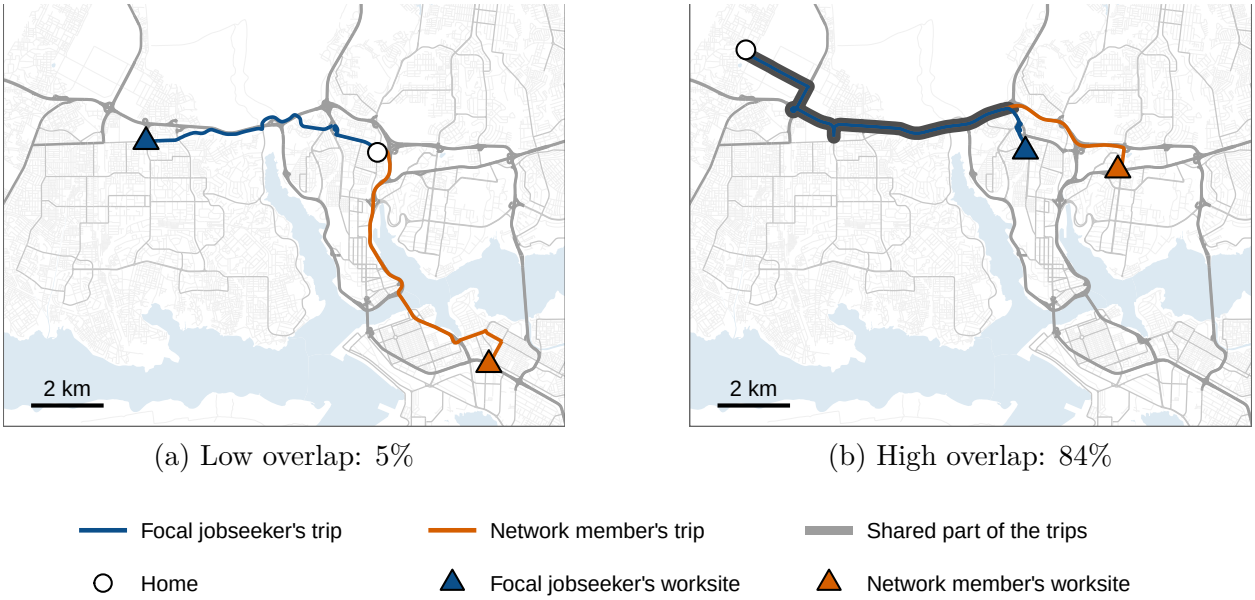
2. For each person and each of the two candidate worksites, we trace the fastest public-transport trip from home to the worksite: walking to a stop, riding *gbakas*, *woro-woros*, buses or boats, changing lines where needed, and walking to the worksite. We route on two open datasets of Abidjan's public-transport lines. The first was produced in 2019–2021 for the Ministry of Transport by Jungle Bus, Systra and OpenStreetMap Côte d'Ivoire, which collected the lines in the field and mapped them in OpenStreetMap (Jungle Bus et al. 2021). Its 2021 version covers 398 lines: the 112 SOTRA bus lines, 7 boat lines, and 279 informal lines (109 *gbaka* and 170 *woro-woro* lines). The second, collected between November 2022 and May 2023 by Ewarren Financial Services and DigitalTransport4Africa, covers informal lines only: 490 lines (149 *gbaka* and 341 *woro-woro*) (Ewarren Financial Services and DigitalTransport4Africa 2023). We assume that these lines already ran in 2022. The lagoon can be crossed only on the scheduled boat lines. Because the fastest combination of lines can change from one departure to the

next, we compute the overlap for four weekday departures between 7:00 and 7:45 and average them. Each trip starts and ends at the point of the street network closest to the home and the worksite.

3. For each pair of routes, we compute their overlap with CanterburyCommuto.
   (a) We draw a 100-meter corridor around each route, so that two people traveling along the same road count as overlapping even when their traced lines do not coincide exactly.
   (b) The overlap is the share of the area of the focal jobseeker's corridor that lies within the network member's corridor. It ranges from 0, when the two routes share no road, to 1, when the focal jobseeker's whole route runs along the network member's.
4. We compute the overlap under two assignments, whatever the arm the focal jobseeker was actually assigned to.
   (a) *Same worksite:* both travel to the focal jobseeker's assigned worksite. Since neighbors share the focal jobseeker's home location, their overlap is 1 by construction.
   (b) *Different worksites:* the focal jobseeker travels to their assigned worksite, and the network member to the other candidate worksite.
5. We aggregate to the focal jobseeker by taking the maximum overlap across their offered network members (up to three). This measures whether the focal jobseeker could commute with at least one of the network members offered the job.

The above steps create a measure of co-commuting potential between 0 and 1 for each focal jobseeker and each assignment. Figure A10 shows its distribution. Table A10 splits the Different Site and Control arms at the weighted median of the different-worksite measure (0.34). With driving routes instead of public-transport trips, the median co-commuting potential under different worksites is 0.29 instead of 0.34, and the results in Table A10 are similar but weaker: the Different Site effect is 25.9 percentage points larger for above-median commutes among focal jobseekers with above-median potential ($p = 0.08$) and 3.5 points larger among those below ($p = 0.80$), against 26.9 ($p = 0.07$) and 1.9 ($p = 0.89$) with public-transport trips.

Figure A9: Examples of Co-Commuting Route Overlap, Sales Experiment

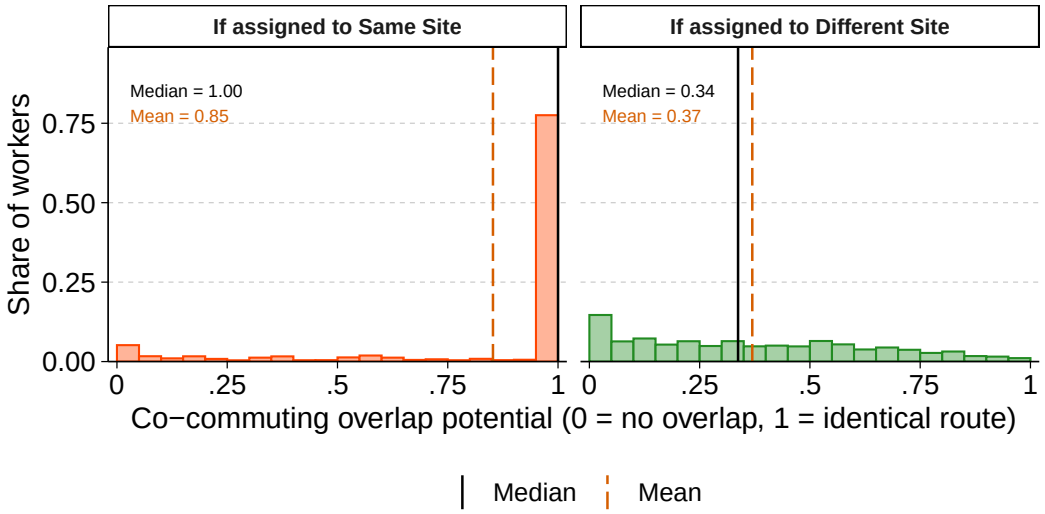


(a) Low overlap: 5%

(b) High overlap: 84%

Notes: Two pairs of a focal jobseeker and a neighbor network member, each traveling to different worksites. Lines trace each person's predicted public-transport trip from home to their worksite leaving at 7:00, and the gray band marks where the two routes could overlap. The potential overlap is the share of the focal jobseeker's route that lies within 100 meters of the network member's route, averaged over four departures between 7:00 and 7:45. Homes are the center of the reported neighborhood, drawn where the trip starts. Map data: © OpenStreetMap contributors (2022); transit lines: Jungle Bus (2021), Ewarren Financial Services and DigitalTransport4Africa (2023).

Figure A10: Distribution of the Co-Commuting Potential, Sales Experiment



Notes: Distribution of the co-commuting potential among the Sales experiment focal jobseekers (main sample, as in Table 3): the maximum, across each focal jobseeker's offered network members, of the share of the focal jobseeker's commute route that runs along the network member's route, from 0 (no overlap) to 1 (identical route). The left panel shows it if the focal jobseeker and their network members were assigned to the same worksite, the right panel if they were assigned to different worksites. Vertical lines mark the median and mean. Inverse-propensity weights applied where relevant.

Table A10: Different Site Effects by Co-Commuting Potential, Sales Experiment

	Covariate: Above-median commute			Covariate: Commute time (hours)		
		Co-commuting potential			Co-commuting potential	
	All	Below median	Above median	All	Below median	Above median
	(1)	(2)	(3)	(4)	(5)	(6)
Covariate	-0.0572	-0.000770	-0.132	-0.0112	-0.00789	-0.00637
	(0.0744)	(0.0972)	(0.116)	(0.0283)	(0.0345)	(0.0512)
	[0.443]	[0.994]	[0.256]	[0.691]	[0.819]	[0.901]
Different Site	-0.0170	0.0471	-0.0738	-0.0419	0.0374	-0.0953
	(0.0564)	(0.0821)	(0.0780)	(0.0915)	(0.132)	(0.134)
	[0.763]	[0.567]	[0.345]	[0.647]	[0.777]	[0.479]
Different Site × Covariate	0.137	0.0187	0.269	0.0335	0.00739	0.0476
	(0.0990)	(0.135)	(0.148)	(0.0377)	(0.0502)	(0.0621)
	[0.166]	[0.890]	[0.070]	[0.374]	[0.883]	[0.445]

Notes: Columns (1) and (4) replicate columns (1) and (4) of Table 3. Columns (2) and (5) and columns (3) and (6) show it separately for jobseekers who have, respectively, below-median or above-median co-commuting potential with their network members, if assigned to the Different Site arm. Inverse-propensity weights applied where relevant.

C.4 Productivity

This subsection describes how we measure the productivity of workers in the sales experiment, and reports additional productivity results. This is specific to the sales experiment: productivity measures were not collected by the factory staff at the time of the factory experiment.

The sales task. Workers hired in the sales experiment canvass on street corners and in market places: they approach passers-by, tell them about the Coin Afrique app, show them how it works, and encourage them to download and use it. The app is a classified-ads platform on which users can create an account, browse goods, set up alerts to be notified when certain goods become available, and request a seller's contact details when they see something of interest. A website also exists, though people are encouraged to primarily use the app; in the measures below, we aggregate the app-based and website-based measures. Each worker carries a personal QR code. An interested passer-by scans it to reach the app store and download the app, which links every new user to the worker who recruited them.

Data sources. We measure productivity from two independent sources.

1. *Evening debrief.* At the end of each workday, the field supervisor gathers the team and asks each worker how many people they approached (*Contacts*), how many of them accessed the platform by downloading the app or visiting the website (*Access*), how many created an account (*Users*), and how many set up an alert (*Alerts*). These numbers are self-reported, but they must remain consistent with what the supervisors observe during unannounced visits throughout the day.

2. *App log.* Coin Afrique's server logs record the behavior of every user who reached the app through a worker's QR code at least once: whether they started using the platform (*New use*), whether they created an account (*Accounts*), how many alerts they set up (*Alerts*), how many times they opened the app (*Sessions*), and how many times they requested a seller's contact details (*Seller contacts*). Seller contacts are what Coin Afrique ultimately values: they are its best proxy for the interactions between buyers and sellers generated by the campaign.

Three dimensions of a worker's output are therefore recorded by both sources, under different names: the people who start using the platform (*Access* in the evening debrief, *New use* in the app log), the accounts created (*Users* and *Accounts*), and the alerts set up (*Alerts* in both). Figure A11 shows each pair on a single row, labeled *New use / Access* and *Accounts / Users*. It is unclear which source measures these dimensions better: the evening debriefs can

suffer from misreporting, while the app logs miss users who did not scan the QR code and downloaded the app directly from the app store. We therefore report both. Contacts are recorded only in the evening debriefs, while Sessions and Seller contacts are recorded only in the app logs. Because the logs follow users beyond their first interaction with a worker, they provide a quality-adjusted measure of productivity.

Index construction. For each data source, we build a standardized productivity index (Anderson 2008; Kling et al. 2007) in three steps.

1. For the evening debrief, we sum each component over the worker's work week. For the app log, we sum it over the day each user was recruited, the app-log counterpart of the evening debrief, and over the 20 days after each user first accesses the platform, for users acquired during the work week; the logs were extracted 20 days after the last workday. This yields three sources: the *App log, 20 days*, the *App log, first day*, and the *Evening debrief*.
2. We standardize each component against the Control arm: we subtract its mean among Control focal jobseekers in the productivity sample (defined below) and divide by its standard deviation in that group, using inverse-propensity weights.
3. We average a source's standardized components and standardize this average again against the Control arm.

Each index is thus expressed in standard deviations of productivity among Control workers. The network members in Panels (b) and (c) of Figure 5 are placed on the same scale, using the Control-arm means and standard deviations of the focal jobseekers.

Sample and missing data. The productivity sample is the 264 focal jobseekers in the estimation sample of Table 3 who took up the job. Since productivity is observed only for those who take up, the estimates combine selection effects (who takes up the job) and interaction effects (how productivity changes when a worker is hired together with network members, under each arm). Evening debrief data are available for all 264. We are missing the app logs for 4% of them (10 out of 264), either because the worker did not use the QR code properly or because their phone was not supported by the system.

Results. Section 3.4 reports the main productivity results. Figure 5 compares the average productivity of focal jobseekers across treatment arms, both overall and among those with above-median commuting time, and compares focal jobseekers with the network members hired through the experiment; Table A12 reports the standard error and p-value of each of

these comparisons. This appendix reports two additional exhibits. Table A11 reports the productivity effects by commuting time. Figure A11 reports the effects on each component of the index, by data source.

Table A11: Productivity Effects by Commuting Time

	Productivity index			Seller contacts	
	App log, 20 days	App log, first day	Evening debrief	App log, 20 days	App log, first day
	(1)	(2)	(3)	(4)	(5)
Panel A: Regression coefficients					
Above-median commute	-0.128	-0.0540	-0.0268	-0.355	-0.137
	(0.286)	(0.296)	(0.286)	(0.239)	(0.253)
	[0.656]	[0.855]	[0.925]	[0.139]	[0.589]
Different Site	-0.0241	-0.0579	-0.137	-0.206	-0.263
	(0.208)	(0.200)	(0.204)	(0.230)	(0.186)
	[0.908]	[0.772]	[0.503]	[0.371]	[0.158]
Different Site × Above-median	0.102	0.0546	0.0967	0.395	0.277
	(0.363)	(0.368)	(0.366)	(0.320)	(0.301)
	[0.778]	[0.882]	[0.792]	[0.217]	[0.358]
Same Site	0.0293	0.0595	-0.181	-0.0878	0.00252
	(0.207)	(0.211)	(0.200)	(0.222)	(0.233)
	[0.888]	[0.778]	[0.366]	[0.692]	[0.991]
Same Site × Above-median	0.0824	0.0298	0.313	0.212	0.0846
	(0.358)	(0.373)	(0.351)	(0.323)	(0.356)
	[0.818]	[0.936]	[0.374]	[0.511]	[0.812]
Panel B: Effects for above-median commuters					
Different Site	0.0782	-0.00331	-0.0400	0.190	0.0139
	(0.297)	(0.309)	(0.304)	(0.222)	(0.236)
	[0.792]	[0.991]	[0.895]	[0.394]	[0.953]
Same Site	0.112	0.0893	0.132	0.125	0.0871
	(0.292)	(0.308)	(0.289)	(0.235)	(0.269)
	[0.702]	[0.772]	[0.648]	[0.595]	[0.746]
Control mean	0.00	-0.00	-0.00	-0.00	0.00
Observations	254	254	264	254	254

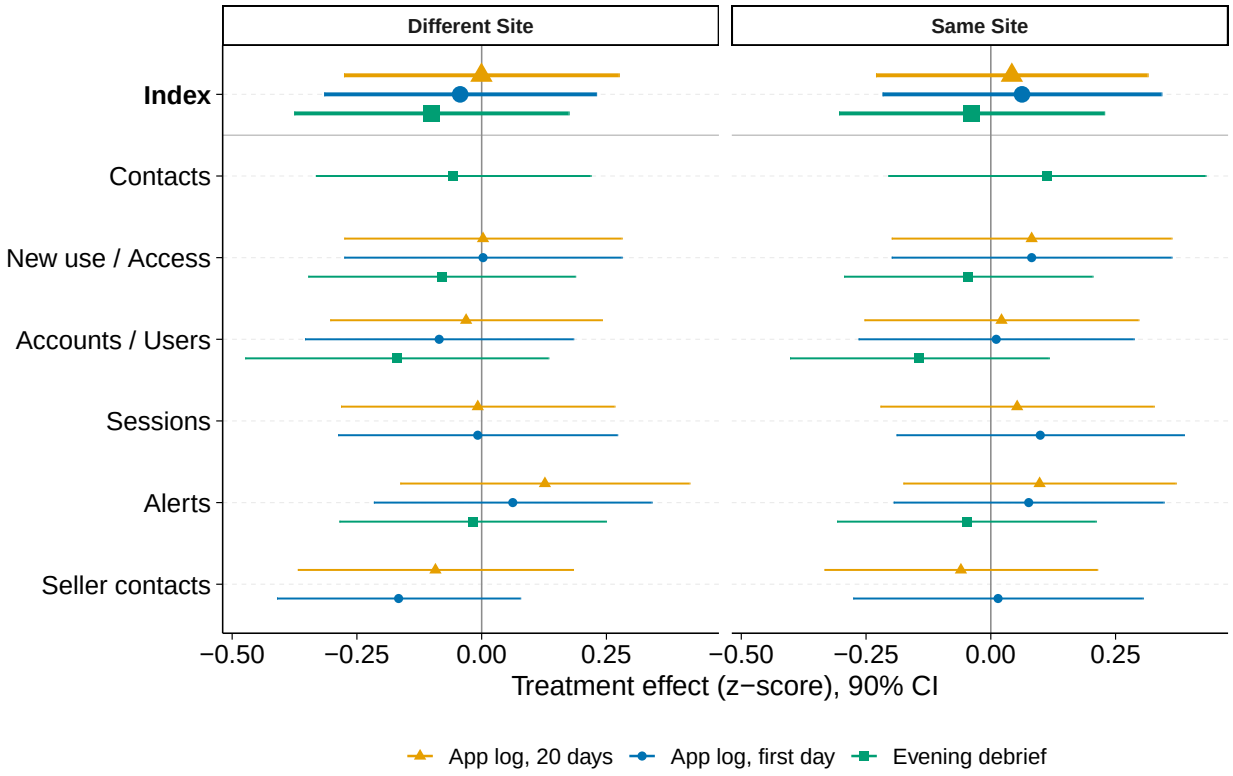
Notes: Each column estimates the specification of column (1) of Table 3, on the focal jobseekers who took up the offered job. Columns (1)–(3) use productivity indices as the outcome, one column per data source; columns (4)–(5) use the number of seller contacts recorded in the firm's app log. Panel B reports the Same Site and Different Site effects among above-median commuters, based on the Panel A estimates (delta-method standard errors). Inverse-propensity weights applied where relevant. Each cell reports the coefficient, the robust standard error, and the p-value in brackets.

Table A12: Productivity Comparisons Shown in Figure 5

	App log, 20 days	App log, first day	Evening debrief
Panel A: Focal jobseekers who took up the job			
Same Site − Control	0.0422	0.0625	-0.0384
	(0.165)	(0.170)	(0.161)
	[0.799]	[0.713]	[0.812]
Observations	108 / 62	108 / 62	111 / 67
Different Site − Control	-0.000454	-0.0431	-0.101
	(0.168)	(0.166)	(0.167)
	[0.998]	[0.795]	[0.549]
Observations	84 / 62	84 / 62	86 / 67
Same Site − Different Site	0.0426	0.106	0.0622
	(0.152)	(0.155)	(0.151)
	[0.780]	[0.497]	[0.680]
Observations	108 / 84	108 / 84	111 / 86
Panel B: Network members and focal jobseekers			
Network members − all focal jobseekers	-0.0281	-0.0359	-0.0583
	(0.0930)	(0.0919)	(0.148)
	[0.762]	[0.697]	[0.694]
Observations	192 / 254	192 / 254	204 / 264
Network members − Control focal jobseekers	-0.0110	-0.0239	-0.106
	(0.157)	(0.157)	(0.187)
	[0.944]	[0.879]	[0.572]
Observations	192 / 62	192 / 62	204 / 67
Panel C: Focal jobseekers with an above-median commute			
Same Site − Control	0.112	0.0893	0.132
	(0.293)	(0.309)	(0.290)
	[0.704]	[0.774]	[0.650]
Observations	52 / 16	52 / 16	53 / 17
Same Site and their network members − Control	0.0226	-0.00647	0.117
	(0.282)	(0.297)	(0.349)
	[0.936]	[0.983]	[0.739]
Observations	96 / 16	96 / 16	100 / 17

Notes: Difference in average productivity between the two groups, in standard deviations of the Control focal jobseekers (the scale of Figure 5). Inverse-propensity weights applied where relevant. Robust standard errors in parentheses, clustered on the focal jobseeker in the rows that include network members; p-values in brackets. *Observations*: first group / second group.

Figure A11: Effects on Productivity, by Component and Data Source



Notes: Same Site and Different Site effects on productivity relative to Control, with 90% confidence intervals, among the focal jobseekers of Table 3 who took up the job. Each component is standardized against the Control group; the *Index* averages a data source's components and is standardized again. Appendix C.4 describes the data sources. Inverse-propensity weights applied where relevant. Robust standard errors.

C.5 Amenity Value Calibration

This subsection expresses the Same Site effect on take-up as the wage increase that would raise take-up by as much. It uses a wage increase that took place between the operational pilot of the sales experiment and the main experiment.

Conceptual framework. Each focal jobseeker has a reservation wage r, the lowest daily wage at which they would take the job, distributed across focal jobseekers with cumulative distribution function F. A focal jobseeker offered wage ω takes the job if $\omega \geq r$, so take-up at wage ω is $F(\omega)$. Let ω_0 and ω_1 be the daily wages of the pilot and of the main experiment, and y_0 and y_1 the take-up rates at each wage, expressed as shares:

$$F(\omega_0) = y_0, \qquad F(\omega_1) = y_1. \tag{A2}$$

The Same Site treatment raises take-up at the main wage from y_1 to $y_1 + \beta$. The wage increase δ that would raise take-up by as much solves

$$F(\omega_1 + \delta) = y_1 + \beta. \tag{A3}$$

We assume that the focal jobseekers offered the job at each wage, and those of the main experiment, share the same distribution F. We consider two distributions, each with two parameters, which the two conditions in (A2) identify exactly.

1. *Uniform distribution.* Let r be uniform on $[a, b]$, with $a \leq \omega_0$ and $\omega_1 + \delta \leq b$. On $[a, b]$, $F(\omega) = (\omega - a)/(b - a)$. Subtracting the two conditions in (A2) gives $y_1 - y_0 = (\omega_1 - \omega_0)/(b - a)$. Subtracting $F(\omega_1) = y_1$ from (A3) gives $\delta/(b - a) = \beta$. Hence

$$\delta_U = \beta \, \frac{\omega_1 - \omega_0}{y_1 - y_0}. \tag{A4}$$

   Take-up then rises linearly with the wage, and the fraction is the cost to the firm of raising take-up by one percentage point when rates are expressed in percentage points, which leaves Equation (A4) unchanged. The bounds a and b are not separately identified, and the result does not require them.

2. *Normal distribution.* Let r be normal with mean μ and standard deviation σ, so that $F(\omega) = \Phi\Big((\omega - \mu)/\sigma\Big)$. The conditions in (A2) become $\Phi^{-1}(y_k) = (\omega_k - \mu)/\sigma$ for $k \in \{0, 1\}$. Subtracting them gives $\sigma = (\omega_1 - \omega_0)/\Big(\Phi^{-1}(y_1) - \Phi^{-1}(y_0)\Big)$, and the second gives $\mu = \omega_1 - \sigma \, \Phi^{-1}(y_1)$. Condition (A3) becomes $(\omega_1 + \delta - \mu)/\sigma = \Phi^{-1}(y_1 + \beta)$, so

that

$$\delta_N = \sigma\,\Phi^{-1}(y_1+\beta) - (\omega_1 - \mu) = \sigma\Big[\Phi^{-1}(y_1+\beta) - \Phi^{-1}(y_1)\Big], \tag{A5}$$

which requires $y_1 + \beta < 1$. Unlike δ_U, δ_N depends on the level of take-up the effect starts from, which we set at y_1.

The wage increase. During the operational pilot, take-up of the job offers was low. Of the 66 focal jobseekers offered the job in its two rounds, 9.1% took it up, and 13.0% of the 23 in the Control arm did. Of the 51 job offers made to their network members in the same rounds, 5.9% were taken up. At such rates, the campaign was not operationally viable: the firm could not reach enough applicants to fill its positions. Anecdotally, applicants raised the cost of commuting as a primary concern. To remedy this, the daily wage was raised from 7,000 to 9,500 FCFA, a 36% increase.

The take-up response. Take-up rose immediately, for focal jobseekers and network members alike. We compare the two pilot rounds with the two rounds of hiring that immediately followed them, to keep the comparison as close in time as possible. Of the 80 focal jobseekers offered the job in these two rounds, 23.8% took it up, and 19.2% of the 26 in Control did. Of the 46 job offers made to their network members in these rounds, 15.2% were taken up.[18]

Three groups of focal jobseekers can measure the response of take-up to the wage. The cleanest is the Control arm, since its focal jobseekers were hired alone: their take-up rose from 13.0% to 19.2%, so that, under Equation (A4), raising take-up by one percentage point costs the firm 404 FCFA per worker-day. In both the Same Site and Different Site arms, the change in take-up also reflects changes in whether network members take up the job. But the Control arm is small, with 3 takers before the wage increase and 5 after. We therefore pool Control with Different Site, since offering jobs to network members at another worksite has no effect on the focal jobseekers' take-up (Table 2).[19] Take-up in the pooled arms rose from 9.3% to 23.4%, so raising it by one percentage point costs 177 FCFA. Pooling all three arms, take-up rose from 8.9% to 23.6%, and the cost is 170 FCFA. The larger the estimated response of take-up to the wage, the smaller the implied value of the amenity, so the estimate that includes Same Site is likely to understate it.

18. A focal jobseeker's round is the date of their baseline call, when the job was offered. The samples apply the eligibility criteria of the main experiment, except the validity of the commute-time answers, which the calibration does not use and which the first pilot round elicited with a different question. A network member's offer belongs to the round in which the firm called them: the firm called some members again in later rounds, and an offer made after the wage increase is not counted as a pilot offer.

19. The pooled take-up rate is the average of the two arms' rates, so that the change reflects the wage rather than changes in the mix of arms across rounds: the first pilot round had no Different Site arm.

Wage equivalent of commuting together. Same Site raises take-up by 12.4 percentage points in the main sample of Table 2 and by 7.1 percentage points in the full sample (columns (1) and (3)). Pooling Control and Different Site, matching these effects with the wage alone would require raising the daily wage by 2,205 and 1,253 FCFA, respectively, if reservation wages are uniformly distributed (Equation (A4)), or 23% and 13% of the 9,500 FCFA wage of the main experiment. If they are normally distributed (Equation (A5)), the required increases are 1,520 and 899 FCFA, or 16% and 9%. Table A13 reports these figures for each group of focal jobseekers: using Control only gives larger values, and using all three arms gives similar ones. Across the two distributions and the two samples, the amenity of commuting together is thus equivalent to a wage increase of 9 to 23%, which is our preferred range.

Table A13: Wage Equivalent of the Same Site Amenity, Sales Experiment

		Take-up (%)		Wage equivalent (FCFA)			
				Uniform		Normal	
Arms	Focal jobseekers (1)	Pilot (2)	Next rounds (3)	Full sample (4)	Main sample (5)	Full sample (6)	Main sample (7)
Control and Different Site	41 / 55	9.3	23.4	1,253 (13.2%)	2,205 (23.2%)	899 (9.5%)	1,520 (16.0%)
Control	23 / 26	13.0	19.2	2,857 (30.1%)	5,028 (52.9%)	2,308 (24.3%)	3,850 (40.5%)
All three arms	66 / 80	8.9	23.6	1,199 (12.6%)	2,111 (22.2%)	849 (8.9%)	1,436 (15.1%)
Same Site effect (pp)				7.1	12.4	7.1	12.4

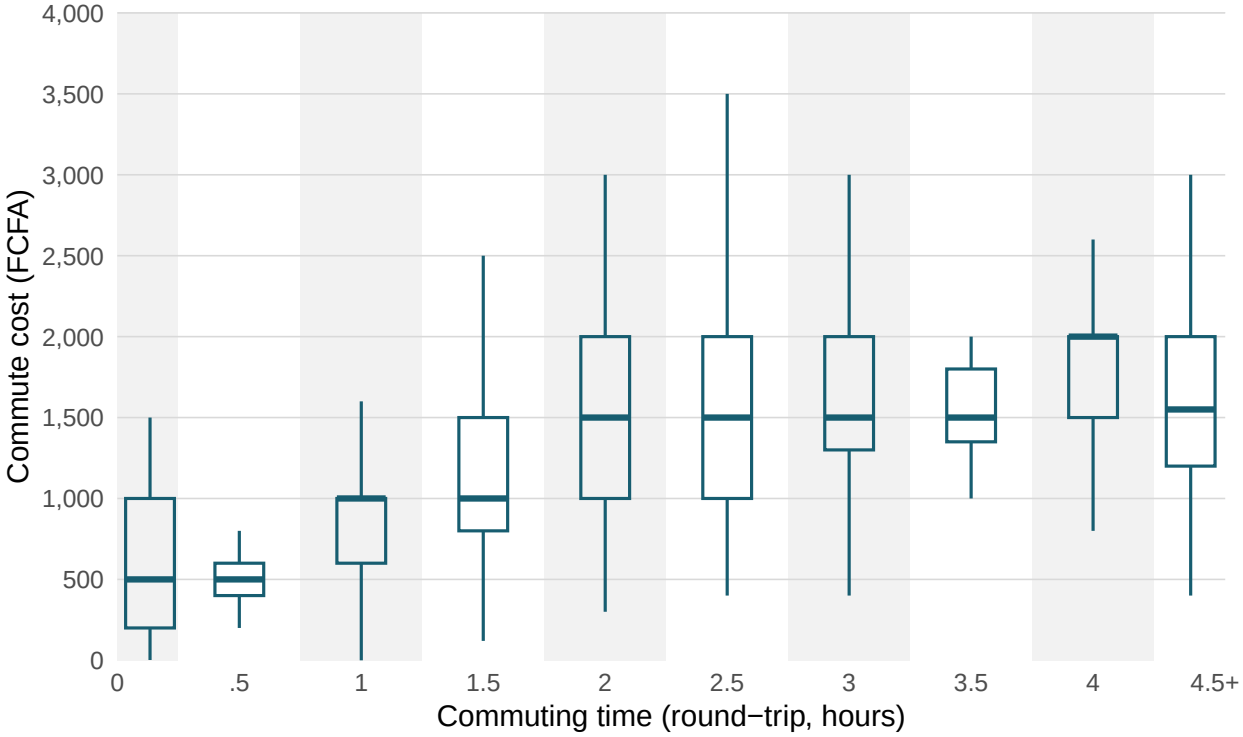
Notes: Wage increase that raises take-up by as much as the Same Site treatment, if reservation wages are uniformly distributed (Equation (A4), columns (4) and (5)) or normally distributed (Equation (A5), columns (6) and (7)). Each row uses the take-up response of a different group of focal jobseekers to the increase of the daily wage from 7,000 to 9,500 FCFA. Column (1) gives the number of focal jobseekers offered the job in the two pilot rounds and in the two rounds that followed. Columns (2) and (3) report their take-up; for several arms, the simple average of the arms' rates. Columns (4) to (7) use the Same Site effect of columns (3) and (1) of Table 2, reported in percentage points (pp) in the last row, and give the wage increase in FCFA per worker-day and, in parentheses, as a share of the 9,500 FCFA wage.

C.6 Mechanisms

Costs. At baseline, each focal jobseeker estimated the round-trip time and cost of their commute to each of the two candidate worksites, before learning which one they were assigned to. Figure A12 shows how the expected cost varies with the expected time. The self-reported cost and time required to travel from home to the offered jobs are positively correlated, but they vary enough independently to separately identify the treatment effect heterogeneity by commuting cost and time. First, longer trips are expected to cost more only up to about two hours: the median expected round-trip cost rises from 500 FCFA for a half-hour trip to 1,500 FCFA for a two-hour trip, and changes little beyond that. Second, the cost varies widely for a given trip length: among jobseekers expecting a two-hour trip, a quarter expect to pay 1,000 FCFA or less, and a quarter 2,000 FCFA or more. Both patterns fit the commuting modes described above: *gbakas* charge a fixed fare per ride, whatever the distance traveled, but trips often require changing *gbakas* along the way, with long waiting and walking times in between. Table A14 adds the commute cost to the specifications of Table 3, alone and together with the commute time.

Safety. Figure A13 shows the effects of the two treatment arms on job take-up separately by how safe focal jobseekers feel coming home at night and by whether they have been harassed on the way home. Table A15 interacts the treatment arms with an indicator for not feeling very safe coming home at night.

Figure A12: Commute Cost by Commute Time, Sales Experiment



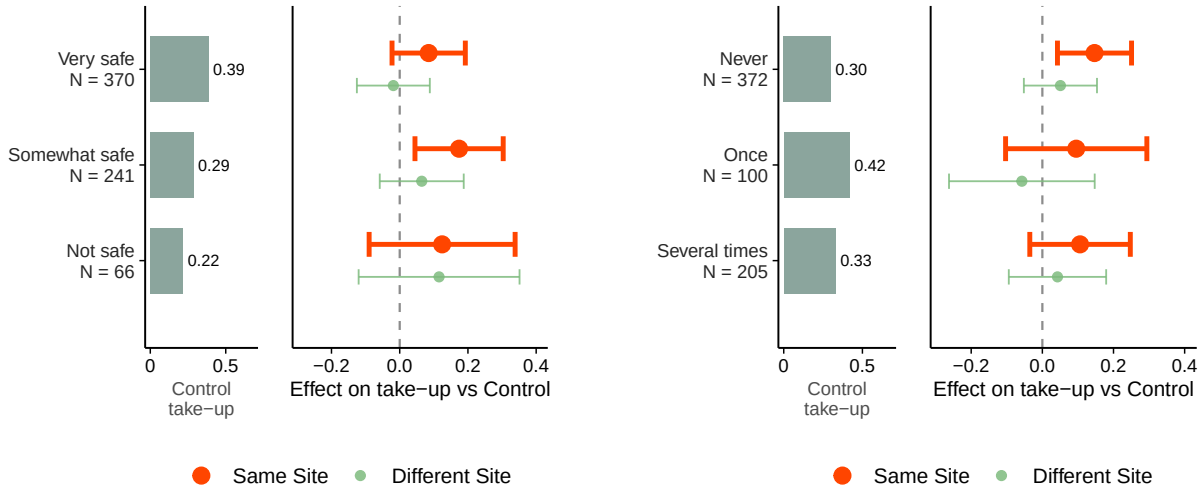
Notes: Baseline answers to "How long do you expect the trip to take?" and "How much do you expect it to cost?", both for the round trip, asked for each of the two candidate worksites before assignment. Each focal jobseeker contributes both answers. Answers are grouped in 30-minute bins centered on the half hour; the last bin groups trips longer than 4 hours 15 minutes. Boxes show the median and interquartile range; whiskers extend to the furthest answer within 1.5 times the interquartile range. The sample is that of Table 3. Inverse-propensity weights applied where relevant.

Table A14: Treatment Effects by Commuting Cost and Time, Sales Experiment

	Above median			Continuous		
	Time (1)	Cost (2)	Both (3)	Time (4)	Cost (5)	Both (6)
Time	-0.0572		-0.0413	-0.0112		-0.0125
	(0.0744)		(0.0803)	(0.0283)		(0.0325)
	[0.443]		[0.607]	[0.691]		[0.701]
Cost		-0.0481	-0.0288		-0.00822	0.00361
		(0.0688)	(0.0741)		(0.0469)	(0.0543)
		[0.484]	[0.697]		[0.861]	[0.947]
Different Site	-0.0170	0.0466	0.0168	-0.0419	0.0906	0.0272
	(0.0563)	(0.0647)	(0.0668)	(0.0915)	(0.0930)	(0.105)
	[0.763]	[0.472]	[0.801]	[0.647]	[0.331]	[0.796]
Different Site × Time	0.137		0.153	0.0335		0.0521
	(0.0990)		(0.105)	(0.0377)		(0.0420)
	[0.166]		[0.144]	[0.374]		[0.215]
Different Site × Cost		-0.0318	-0.0811		-0.0445	-0.0804
		(0.0926)	(0.0982)		(0.0600)	(0.0672)
		[0.731]	[0.409]		[0.458]	[0.232]
Same Site	0.0372	0.0995	0.0535	-0.0270	0.0822	-0.0322
	(0.0574)	(0.0638)	(0.0656)	(0.0910)	(0.0797)	(0.0999)
	[0.517]	[0.119]	[0.415]	[0.767]	[0.303]	[0.747]
Same Site × Time	0.242		0.268	0.0712		0.0693
	(0.0989)		(0.107)	(0.0370)		(0.0405)
	[0.014]		[0.013]	[0.055]		[0.087]
Same Site × Cost		0.0550	-0.0617		0.0307	0.00640
		(0.0943)	(0.101)		(0.0486)	(0.0568)
		[0.560]	[0.543]		[0.528]	[0.910]
Control take-up	0.33	0.33	0.33	0.33	0.33	0.33
Observations	677	677	677	677	677	677

Notes: Specifications of Table 3, columns (1) and (4), with the round-trip commute cost to the assigned worksite ("How much do you expect it to cost?") interacted with both treatment arms. Columns (1) and (4) interact the arms with the commute time only, and reproduce Table 3, columns (1) and (4); columns (2) and (5) with the commute cost only; columns (3) and (6) with both. Columns (1)–(3) use indicators for a cost above the median (1,200 FCFA) and a time above the median (2 hours); columns (4)–(6) use the cost in 1,000 FCFA and the time in hours. Inverse-propensity weights applied where relevant. Each cell reports the coefficient, the robust standard error (in parentheses), and the p-value [in brackets].

Figure A13: Treatment Effects by Baseline Commuting Safety, Sales Experiment



(a) Perceived safety

(b) Ever harassed on the way home

Notes: Each panel splits focal jobseekers into groups by a baseline commuting-safety measure. Within each panel, the right part plots Same Site and Different Site effects on job take-up relative to Control, with 90% confidence intervals. The point estimates and confidence intervals are from regressing take-up on group indicators and their interactions with the Different and Same Site treatment arms. The left part reports the Control-arm take-up rate within each group, with the group's sample size. Panel (a) splits by perceived safety, from "How safe do you feel when you have to come home at night after work?", where *Not safe* pools the three least-safe responses of a five-point scale. Panel (b) splits by harassment, from "Have you ever been bothered or harassed by someone on your way home?". The estimation sample is that of Table 3. Inverse-propensity weights applied where relevant. Standard errors are robust.

Table A15: Treatment Effects by Perceived Safety, Sales Experiment

Commute to work at baseline	All	Yes	No		
Gender	All	All	All	Men	Women
	(1)	(2)	(3)	(4)	(5)
Not very safe	-0.117	-0.0813	-0.160	-0.144	-0.188
	(0.0681)	(0.0936)	(0.100)	(0.127)	(0.167)
	[0.087]	[0.386]	[0.113]	[0.256]	[0.262]
Different Site	-0.0187	0.0213	-0.0654	-0.0548	-0.0566
	(0.0648)	(0.0872)	(0.0974)	(0.126)	(0.159)
	[0.774]	[0.807]	[0.502]	[0.665]	[0.722]
Different Site × Not very safe	0.0970	0.0734	0.129	0.0769	0.200
	(0.0924)	(0.129)	(0.134)	(0.177)	(0.211)
	[0.295]	[0.570]	[0.338]	[0.664]	[0.345]
Same Site	0.0849	0.179	-0.0193	-0.0135	-0.0501
	(0.0652)	(0.0883)	(0.0967)	(0.118)	(0.168)
	[0.193]	[0.043]	[0.842]	[0.909]	[0.766]
Same Site × Not very safe	0.0767	-0.0667	0.235	0.250	0.274
	(0.0938)	(0.129)	(0.137)	(0.175)	(0.223)
	[0.414]	[0.606]	[0.087]	[0.154]	[0.221]
Same Site effect if not very safe	0.162	0.113	0.216	0.237	0.224
	(0.0675)	(0.0944)	(0.0968)	(0.129)	(0.147)
	[0.017]	[0.233]	[0.026]	[0.066]	[0.127]
Control take-up	0.33	0.31	0.35	0.39	0.28
Share not very safe	0.46	0.44	0.48	0.45	0.52
Observations	677	348	329	199	130

Notes: Treatment effects on job take-up by perceived safety. *Not very safe* indicates focal jobseekers who do not report feeling very safe coming home at night, pooling the *Somewhat safe* and *Not safe* groups of Figure A13. The sample is that of Table 3. Column (1) keeps all focal jobseekers. Column (2) keeps those who commuted to an income-generating activity at baseline. Column (3) keeps those who did not, and columns (4) and (5) split them by gender. *Same Site effect if not very safe* is the sum of the *Same Site* and *Same Site* × *Not very safe* coefficients. Inverse-propensity weights applied where relevant. Each cell reports the coefficient, the robust standard error, and the p-value in brackets.

C.7 Gender

These exhibits support the discussion of gender in Section 3.6. Table A16 reports baseline characteristics separately for men and women. Table A17 estimates the treatment effects on take-up by gender, and their heterogeneity by commuting time. Figure A14 splits the safety heterogeneity of Figure A13 by gender.

Table A16: Baseline Characteristics by Gender, Sales Experiment

Variable	Men			Women		
	Mean	SD	N	Mean	SD	N
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Focal Jobseeker Characteristics						
Age	26	(5)	534	24	(4)	314
Years of education	15	(2)	534	14	(2)	314
Has a partner	0.09	(0.28)	534	0.10	(0.30)	314
Has young children	0.14	(0.35)	534	0.23	(0.42)	314
Panel B: Focal Jobseeker Baseline Income-Generating Activities						
Any IGA, past 3 months	0.63	(0.48)	534	0.62	(0.49)	314
Any earnings, past 30 days	0.41	(0.49)	534	0.38	(0.49)	314
Earnings, past 30 days (000s FCFA)	48.7	(107.6)	534	30.4	(54.3)	314
Any earnings, past 7 days	0.22	(0.42)	534	0.20	(0.40)	314
Panel C: Focal Jobseeker Baseline Commute						
Any commute	0.56	(0.50)	534	0.47	(0.50)	314
Commuting time (hours)	1.44	(1.29)	301	1.47	(1.19)	147
Commute with people they know	0.52	(0.50)	301	0.46	(0.50)	147

Notes: Baseline characteristics of focal jobseekers in the sales analysis sample, reported separately by gender. "Mean" and "SD" are the mean and standard deviation within gender, and "N" the number of non-missing observations. Variable definitions and panel structure are as in Table A6. Inverse-propensity weights applied where relevant. Exchange rate in December 2022: 620 FCFA = 1 USD.

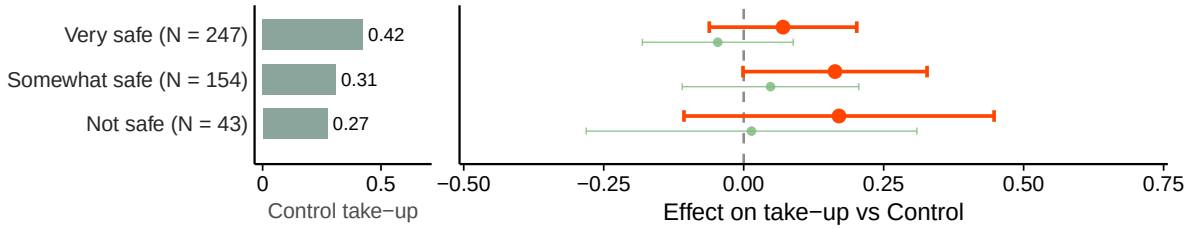
Table A17: Treatment Effects on Job Take-up by Commuting Time and Gender, Sales Experiment

	No commute interaction		Above-median commute		Commute time (hours)	
	Men	Women	Men	Women	Men	Women
	(1)	(2)	(3)	(4)	(5)	(6)
Covariate			0.00890	-0.0983	0.0451	-0.0647
			(0.103)	(0.113)	(0.0408)	(0.0351)
			[0.931]	[0.386]	[0.269]	[0.067]
Different Site	-0.00128	0.0848	-0.0357	0.0261	0.0297	-0.105
	(0.0586)	(0.0758)	(0.0697)	(0.0972)	(0.116)	(0.151)
	[0.983]	[0.264]	[0.609]	[0.789]	[0.798]	[0.489]
Different Site × Covariate			0.0879	0.147	-0.0197	0.0818
			(0.131)	(0.158)	(0.0516)	(0.0542)
			[0.504]	[0.352]	[0.704]	[0.133]
Same Site	0.118	0.126	0.0654	-0.0442	0.142	-0.294
	(0.0584)	(0.0788)	(0.0690)	(0.101)	(0.122)	(0.138)
	[0.044]	[0.112]	[0.344]	[0.663]	[0.246]	[0.034]
Same Site × Covariate			0.153	0.371	-0.0142	0.173
			(0.131)	(0.157)	(0.0560)	(0.0451)
			[0.244]	[0.019]	[0.800]	[0.000]
Control take-up	0.36	0.27	0.36	0.27	0.36	0.27
Observations	444	233	444	233	444	233
p(Men = Women), pooled gender interaction:						
Same Site	0.936		0.371		0.018	
Same Site × commute			0.285		0.009	

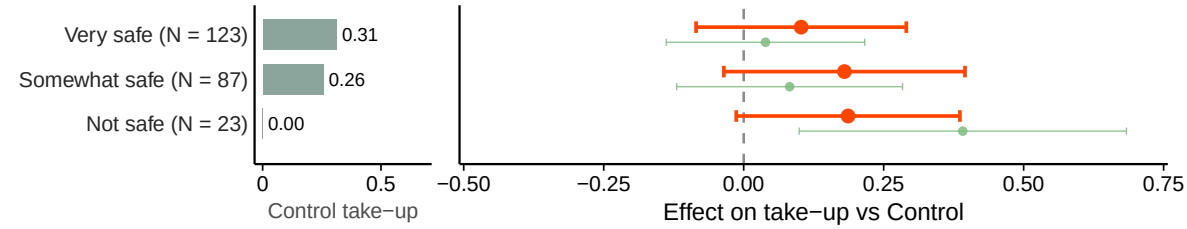
Notes: Estimated separately for men and women, on the sample of Table 3. Columns (1)–(2) regress take-up on the two treatment arms. Columns (3)–(6) follow columns (1) and (4) of Table 3: columns (3)–(4) use the above-median commute indicator and columns (5)–(6) the continuous round-trip commute time. The bottom block reports p-values on the Same Site × Women and Same Site × commute × Women coefficients from a single pooled regression fully interacted with an indicator for women. Inverse-propensity weights applied where relevant. Each cell reports the coefficient, the robust standard error, and the p-value in brackets.

Figure A14: Treatment Effects on Take-up by Baseline Commuting Safety and Gender, Sales Experiment

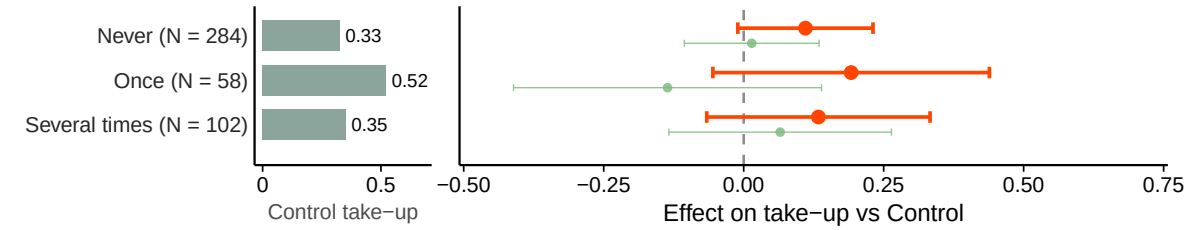
(a) Men: perceived safety



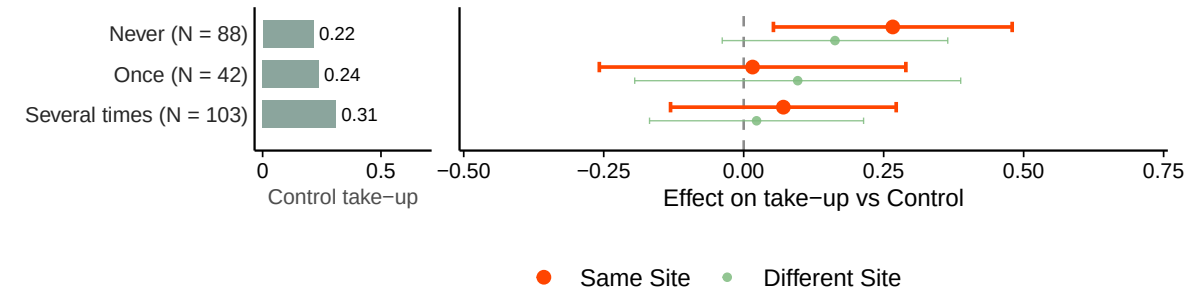
(b) Women: perceived safety



(c) Men: ever harassed on the way home



(d) Women: ever harassed on the way home



Notes: Gender split of Figure A13. Each panel reproduces a panel of that figure on the corresponding gender subsample: perceived safety for men (a) and women (b), and harassment for men (c) and women (d).

C.8 Additional Exhibits

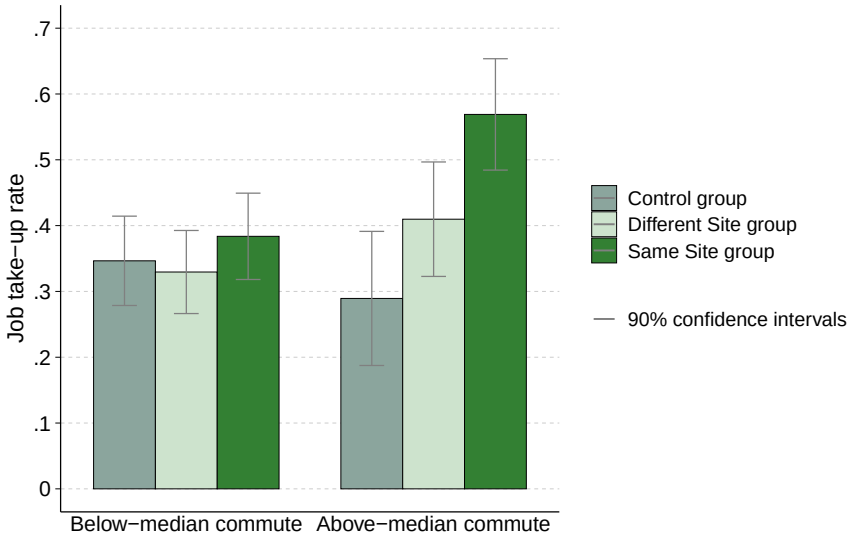
Table A18 re-estimates Table 2 controlling for baseline earnings, the one baseline difference across arms (Table A6). Figure A15 plots column (1) of Table 3 as take-up rates by arm and commute half, and Figure A16 reports the Same Site effect by 30-minute commute bin. Figure A13 reports the treatment effects by baseline commuting safety.

Table A18: Treatment Effects with Baseline Controls, Sales Experiment

	Job take-up (worked ≥ 1 day)			Days worked
	(1)	(2)	(3)	(4)
Different Site	0.0287	0.0240	0.0245	0.0764
	(0.0462)	(0.0435)	(0.0415)	(0.213)
	[0.534]	[0.581]	[0.555]	[0.720]
Same Site	0.121	0.0858	0.0673	0.476
	(0.0470)	(0.0439)	(0.0418)	(0.216)
	[0.010]	[0.051]	[0.108]	[0.028]
p(Same Site = Different Site)	0.039	0.142	0.288	0.051
NMs who could commute together	All	≥ 1	—	All
Control mean	0.329	0.345	0.347	1.52
Observations	677	776	848	677

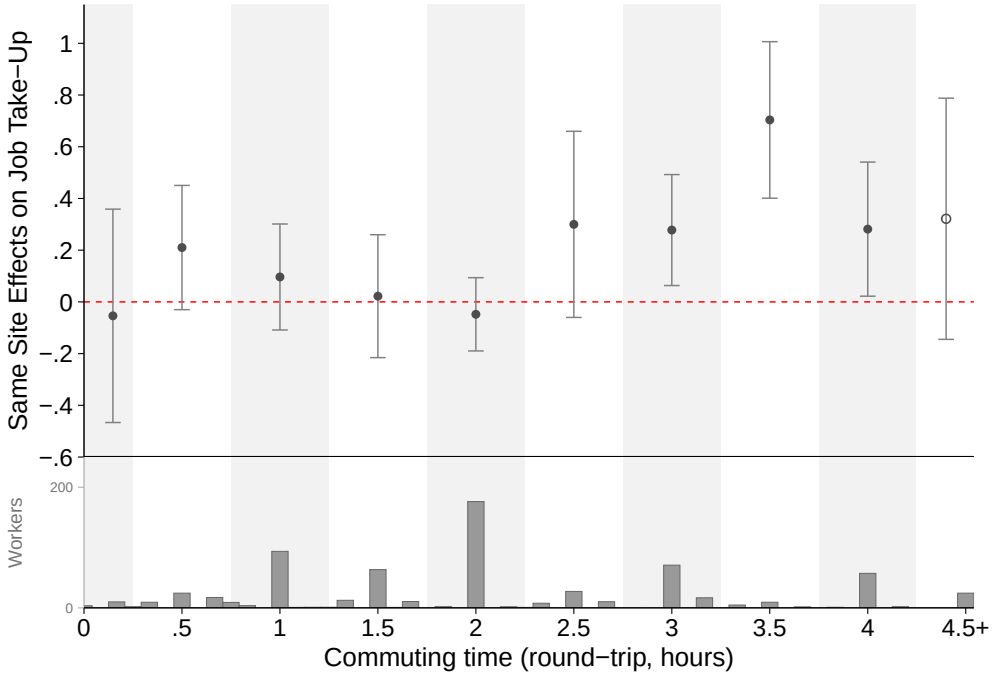
Notes: Same as Table 2, adding a control for the imbalanced baseline value, *Earnings, past 30 days*. Inverse-propensity weights applied where relevant.

Figure A15: Job Take-up by Commuting Time and Treatment Arm, Sales Experiment



Notes: Job take-up rate by treatment arm and commute time. Each bar is the take-up rate in one of six cells defined by treatment arm and an above- or below-median round-trip commute to the randomly assigned worksite, from the specification in column (1) of Table 3. Error bars are 90% confidence intervals from robust standard errors. The estimation sample is that of Table 3. Inverse-propensity weights applied where relevant.

Figure A16: Same Site Effect on Job Take-up by 30-Minute Commuting-Time Bin, Sales Experiment



Notes: Top panel: Same Site treatment effect on job take-up, within each 30-minute bin of the round-trip commute to the assigned worksite, with 90% confidence intervals. Bins are centered on the half-hour (shaded bands): $[0, 0.25]$, $(0.25, 0.75]$, ..., $(3.75, 4.25]$ and above 4.25 hours. The point estimates and confidence intervals are from regressing take-up on bin indicators and their interactions with the Different and Same Site treatment arms. Same specification as column (1) of Table 3, except with ten bins instead of the above-median indicator. The bottom panel shows the distribution of commuting times; its last bar, 4.5+, includes all commutes above 4.5 hours. Inverse-propensity weights applied where relevant.

D Factory Experiment: Supplemental Material

D.1 Implementation Protocols

Hiring process. Olam hires its production workers through a recruitment contractor. The experiment was naturally inserted in the partner company's hiring process, similar to the process described for the sales experiment (Appendix Section C.1 and Figure A8). It differs in two ways.

First, the firm had no list of candidates for us to start from, so we built one. The rosters of focal jobseekers were assembled from lists of candidates provided by placement agencies, by asking current workers at the plant (hired outside the experiment) to recommend people they knew, and by asking community and religious leaders for names of potentially interested jobseekers.[20]

Second, Innovations for Poverty Action called both the focal jobseekers and their network members, and made both of them the job offer. Those who accepted it agreed to share their contact details with the recruitment contractor, which then called them to come to the factory. Candidates who refused to share their contact details could not be hired. At the end of the baseline survey, 148 of the 163 focal jobseekers (91%) accepted the offer, with similar shares across arms.

Timeline. The factory experiment was conducted in two waves, following the expansion of our partner firm's plant over time. For the first batch of hires, 48 focal jobseekers were contacted in August and September 2020.[21] For the second batch of hires, 115 focal job-

20. How a candidate entered the roster does not affect identification: treatment was assigned at random among the candidates on the roster, after their names had been listed. It could matter for interpretation, though, as some focal jobseekers may already know someone working at the plant. Importantly, this is balanced across treatment arms: we are still comparing individuals hired with or without the network members they named at baseline. However, for those who already knew someone at the factory, this variation identifies the effects of changing how many people they know at the factory, not whether they know anyone there. We could thus be increasing the number of potential co-commuters, rather than moving from zero to one. If what matters is to commute with at least one person, this could dampen the estimated effects, relative to a case where the focal jobseekers wouldn't know anyone at the factory to start with.

21. This first wave included a third treatment arm: a "Cash in Networks" arm, where the focal jobseekers received a job offer and were informed that their network members would receive biweekly cash transfers for three months (rather than jobs as in our main treatment arms). The objective was to alleviate redistributive pressure on the focal jobseekers and assess its effects on their job take-up decisions, complementing the work of Carranza et al. (2025) that focuses on productivity and attendance. However, qualitative evidence suggests that pressure arises from a variety of network members, and that offering cash to two of them only would not be sufficient to meaningfully influence pressures to redistribute. Given budget constraints, we therefore dropped this arm from the second, larger, wave of the experiment. Results are similar, albeit noisier, when using Wave 2 for the estimation only (Appendix Table A22). For this reason, the probability of being assigned to a treatment arm varies across waves, and so we apply inverse-probability weighting

seekers were contacted between December 2020 and March 2021. A third wave planned for September 2021 was never implemented by our partner company.

Randomization. Each wave started from a roster with more candidates than positions to fill, so the order in which candidates were called for the baseline survey determined who entered the study. We drew this order at random, but spread out candidates who came through the same person: the order first takes one candidate per person, in random order, then a second one, and so on.[22] We then cut the order into consecutive groups of three candidates in Wave 1 and two in Wave 2. Within each group, one candidate was assigned at random to Control and one to a job offer for their network members; in Wave 1, the third was assigned to Cash in Networks. The candidates assigned to a job offer for their network members were then paired by their position in the calling order, and one of each pair was assigned at random to Same Shift, the other to Different Shift.

In Wave 1, the roster recorded each candidate's gender, and the firm wanted about 80% of its new hires to be women. We therefore grouped women and men separately: each group of three candidates holds only women or only men, and treatment was assigned within these groups. The two lists were then merged into one calling order, alternating twelve women and three men, so that 80% of the first 105 candidates called were women. The Wave 2 roster did not record gender, so neither the calling order nor the assignment used it.

On 20 January 2021, we changed the target shares in Wave 2 to 40% in Control and 30% in each treated arm. Among the Wave 2 candidates not yet called by that date, 30 of those assigned to Control were drawn at random and reassigned, 15 to Same Shift and 15 to Different Shift. The draw depended on whether a candidate had been called, not on whether they had answered. Of the 115 Wave 2 focal jobseekers in our sample, 68 had not yet been called on 20 January 2021.

The probability of each treatment arm therefore differs across focal jobseekers. Setting Cash in Networks aside, it is 1/2 for Control and 1/4 for each treated arm in Wave 1 and for those called in Wave 2 before 20 January 2021. For those called after, it is about 0.35 for Control and 0.32 for each treated arm. As in the sales experiment, we weight each focal jobseeker by the inverse of the probability of the arm they were assigned to. Randomization-inference p-values redraw the assignment within the same groups, and repeat the January

wherever relevant to account for it.

22. In Wave 1, candidates who came through the same current worker with the same type of relationship (friend or family) count together, as do candidates from the same source and neighborhood on the other list; in Wave 2, candidates who came through the same current worker, with the same type of relationship and from the same neighborhood.

2021 reassignment, holding the calling order fixed.

Sample. Table A19 reports baseline characteristics by treatment arm, and Table A20 describes the network members elicited at baseline.

Table A19: Baseline Balance, Factory Experiment

Variable	Control	Different Shift		Same Shift		
	Mean/SD	Mean/SD	P-value	Mean/SD	P-value	N
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Focal Jobseeker Characteristics						
Is a woman	0.72	0.65	0.423	0.81	0.303	163
	(0.45)	(0.48)		(0.39)		
Age	29	28	0.384	31	0.126	163
	(6)	(6)		(5)		
Years of education	8	8	0.613	8	0.892	163
	(5)	(6)		(6)		
Has a partner	0.49	0.52	0.728	0.56	0.432	163
	(0.50)	(0.50)		(0.50)		
Any childcare	0.40	0.20	0.017	0.49	0.375	163
	(0.49)	(0.40)		(0.50)		
Panel B: Focal Jobseeker Baseline Income-Generating Activities						
Any IGA, past 3 months	0.72	0.74	0.592	0.64	0.530	163
	(0.45)	(0.44)		(0.48)		
Any earnings, past 30 days	0.60	0.72	0.083	0.54	0.810	163
	(0.49)	(0.45)		(0.50)		
Earnings, past 30 days (FCFA)	30325	40686	0.152	31672	0.702	163
	(36690)	(51968)		(45402)		
Any earnings, past 7 days	0.46	0.56	0.244	0.42	0.724	163
	(0.50)	(0.50)		(0.49)		

Notes: Baseline covariates of focal jobseekers by treatment arm. Cols. (1), (2), and (4) report the mean and, in parentheses, the standard deviation. Cols. (3) and (5) report randomization-inference p-values for equality of means with the Control group. *Any childcare* equals 1 if the respondent spent any time on childcare in the past seven days. Inverse-propensity weights applied where relevant. Exchange rate in December 2022: 620 FCFA = 1 USD.

Table A20: Elicited Network Members, Factory Experiment

Variable	Mean (1)	SD (2)	N (3)
Panel A: Number of Network Members			
Number of network members	2.92	1.47	163
Named at most 2 network members	0.47	0.50	163
Panel B: Relationship to the Focal Jobseeker			
Share: Friends	0.45	0.41	163
Share: Siblings	0.25	0.33	163
Share: Household members	0.19	0.31	163
Share: Acquaintances	0.11	0.24	163
Share: Cousins	0.11	0.27	163
Share: Parents	0.00	0.03	163
Share: Other relationship	0.06	0.18	163
Panel C: Similarity to the Focal Jobseeker			
Share: Men (among focal men)	0.75	0.36	45
Share: Women (among focal women)	0.87	0.24	118
Similar age (within 2.5 years)	0.38	0.49	163
Similar education (within 1 year)	0.25	0.43	146

Notes: Network members' names and characteristics elicited from focal jobseekers at baseline. One observation per focal jobseeker; shares are computed over all the network members referred by the focal jobseeker. Similar education uses the network members' own baseline survey; it is missing when none of them answered it. Inverse-propensity weights applied where relevant.

D.2 Additional Exhibits

Table A21 re-estimates Table 4 controlling for the two baseline variables that differ across arms (Table A19), and Table A22 restricts it to Wave 2.

Table A21: Main Treatment Effects with Baseline Controls, Factory Experiment

	Days worked (1)	Take-up (2)	Retention (3)	Attendance (4)
Different Shift	12.1	0.0488	25.7	-0.0148
	(17.1)	(0.0864)	(67.0)	(0.0611)
	[0.366]	[0.542]	[0.599]	[0.871]
Same Shift	48.7	0.191	105	-0.0313
	(23.0)	(0.0966)	(60.8)	(0.0583)
	[0.017]	[0.041]	[0.075]	[0.648]
Control mean	39.5	0.284	224	0.710
Observations	163	163	57	57

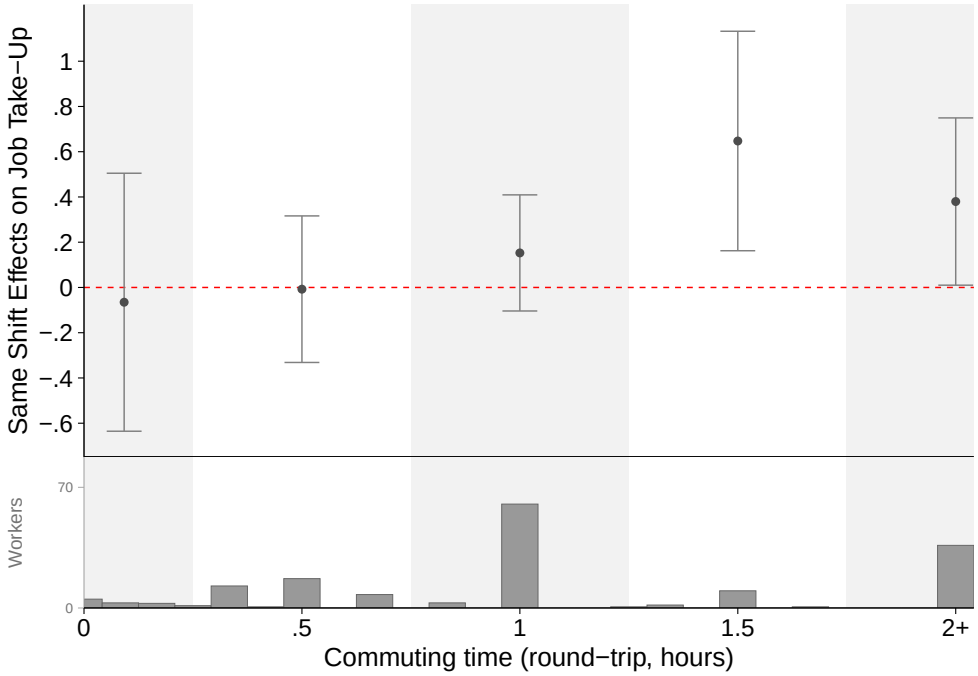
Notes: Same as Table 4, adding controls for the two imbalanced baseline values, *Any childcare* and *Any earnings, past 30 days*. Inverse-propensity weights applied where relevant.

Table A22: Main Treatment Effects, Factory Experiment
(Wave 2 only)

	Days worked (1)	Take-up (2)	Retention (3)	Attendance (4)
Different Shift	7.77	0.0913	-17.0	-0.00320
	(16.2)	(0.101)	(75.7)	(0.0778)
	[0.558]	[0.371]	[0.905]	[0.899]
Same Shift	37.2	0.198	73.6	-0.0602
	(19.7)	(0.109)	(71.6)	(0.0709)
	[0.040]	[0.068]	[0.259]	[0.367]
Control mean	30.0	0.231	237	0.649
Observations	115	115	37	37

Notes: Same as Table 4, estimated on the 115 focal jobseekers of Wave 2 (Appendix D.1). Inverse-propensity weights applied where relevant.

Figure A17: Same Shift Effect on Job Take-up by 30-Minute Commuting-Time Bin, Factory Experiment



Notes: Top panel: Same Shift treatment effect on job take-up, within each 30-minute bin of the round-trip commute to the factory reported at baseline, with 90% confidence intervals. Bins are centered on the half-hour (shaded bands): $[0, 0.25]$, $(0.25, 0.75]$, $(0.75, 1.25]$, $(1.25, 1.75]$ and above 1.75 hours. The point estimates and confidence intervals are from regressing take-up on bin indicators and their interactions with the Different and Same Shift treatment arms. Same specification as Figure A16, with five bins. The bottom panel shows the distribution of commuting times; its last bar, 2+, includes all commutes above 2 hours. Inverse-propensity weights applied where relevant.