Flexibility versus Performance: The Determinants of Labor Contracts in Nairobi, Kenya*

Nathan Barker, Inbar Amit, Alison Andrew, Robert Garlick, Carol Nekesa, Kate Orkin

March 21, 2025

View the Latest Version Here

Abstract

Employment in developing countries is often short and disrupted, generating costly search and limiting the potential for workers to accumulate firm-specific human capital. We study the incentives guiding firms' use of short-term relative to long-term contracts in Nairobi, Kenya, using novel survey data on firms' hiring and contracting practices, and hypothetical vignettes measuring their beliefs and preferences. Our key finding is that the use of short-term labor is governed by a trade-off between managing demand variation versus minimizing adjustment costs and incentivizing worker performance. We first document that firms face considerable variation in demand for goods and services across time, much of which they pass on to workers through short-term contracts: higher demand variation is associated with a greater use of short-term labor. Second, we show that bringing on short-term workers involves adjustment costs: it takes time searching for, hiring, and on-boarding workers, potentially offsetting the gains from flexibility. We show both that median adjustment costs are low, making short-term contracts feasible for many hires, but that hires with greater adjustment costs are more likely to be on long-term contracts. Finally, we show that firms believe contract type incentivizes worker performance: the same worker is expected to perform better when hired on a long-term basis. We incorporate these features—variation in demand, on-boarding costs, and incentives—into a model of firm hiring, through which to interpret contract choice and turnover in low-income countries.

1 Introduction

Labor markets in developing countries are characterized by short, disrupted employment spells. Rather than climbing a job ladder, these disruptions more often reflect workers switching jobs, cycling in and out self-employment, and experiencing frequent spells of unemployment (Donovan et al., 2023). Turnover is especially high at the bottom of the wage distribution. In poor neighborhoods in Nairobi, Kenya, nearly half of young workers had fragmented work spells over the past

*Barker: University of Chicago (corresponding author), barkern@uchicago.edu, Amit: University of Oxford, inbar.amit@economics.ox.ac.uk, Andrew: University of Oxford, Institute for Fiscal studies, alison.andrew@economics.ox.ac.uk, Garlick: Duke University, robert.garlick@duke.edu, Orkin: University of Oxford, kate.orkin@bsg.ox.ac.uk, Nekesa: REMIT Kenya, c_nekesa1@yahoo.com. This study received approval from the University of Chicago's Social and Behavioral Sciences' IRB Office, IRB23-0781, the University of Oxford Department of Economics' Research Ethics Committee, ECONCIA22-23-20, Strathmore University's Institutional Scientific and Ethical Review Committee, SU-ISERC1795/23, and the Government of Kenya's National Commission for Science, Technology, and Innovation, NACOSTI/P/23/28030. This research was funded by the Global Challenges Research Fund UKRI Accelerating Adolescent Achievement Hub and by the Gender, Growth, and Labour Markets in Low Income Countries Programme ($G^2LM|LIC$). Dylan Reich provided excellent research assistance. Obadiah Ogega, Seruya Akhumbi, and the team at REMIT Kenya provided outstanding field management. We sincerely thank the respondents for their time.

two weeks, and 42% of wage workers were not guaranteed regular hours, but were instead called in only when needed.

Disrupted, short-term work can negatively affect both productivity and worker welfare. Workers become more productive as they gain firm- and sector-specific human capital (Becker, 1964; Parsons, 1972; Neal, 1995; Gibbons and Waldman, 2004). High turnover precludes workers and firms from realizing this productivity growth, potentially dampening the overall productivity of the labor force and thus of the economy. Moreover, in a frictional labor market, frequent separations increase the time individuals are without work, and cause them to spend time and money searching for jobs (Caria et al., 2024; Carranza et al., 2022; Franklin, 2018). To the extent separations are not fully predictable, they generate income risk, negatively affecting the welfare of risk-averse individuals (Rosenzweig, 1988; Morduch, 1994).

Given these negative consequences, substantial research and policy efforts have been undertaken to increase the number of good, stable jobs for individuals, in place of the short, disrupted employment commonly available. Much of this work has sought to do so by relaxing supply-side constraints (for example, by training people, or by relaxing information or liquidity constraints), or by increasing the ease with which workers and firms can match (Carranza and McKenzie, 2024).[1] The body of research focused on the demand side is substantially smaller; we lack an overarching understanding of the underlying incentives and constraints governing the types of contracts that firms in developing countries choose to offer.

This paper seeks to fill this gap. We study the question of why firms offer short-term contracts through a survey of the hiring and contracting practices and beliefs and preferences of 601 firms, employing 5,687 workers. These firms are primarily in the retail, hospitality, beauty and manufacturing sectors, and are recruited from a census of firms in a representative sample of enumeration areas from the major business corridors of Western and Central Nairobi.

Our core finding is that firms face a fundamental trade-off when deciding whether or not to offer long-term, stable employment. On the one hand, firms face large and frequent fluctuations in demand for their goods and services. We show that engaging in short-term, high-turnover contracting allows them to dynamically adapt their staffing in response to these fluctuations. However, choosing to staff a business with short-term workers comes at a cost. Every time firms bring in new labor, they have to pay the adjustment costs of searching for and evaluating workers, and in training and getting them up to speed. Moreover, firms perceive a negative incentive effect to short-term contracts: they expect that the same worker will be less reliable and perform worse when only offered short-term work. Our main contribution is providing detailed, empirical evidence using novel data collection methods to show that firms' decisions of whether to commit to their workers on a long-term basis depends on the relative magnitude of these forces.

Evaluating why firms use short-term contracts poses several empirical challenges. Given the

[1]Examples of work that seeks to relax supply-side constraints include: Alfonsi et al. (2020, 2022); Abebe et al. (2021); Cefala et al. (2023); Franklin (2018); Jones and Santos (2022). Studies that seek to improve the efficiency of matching include: Abebe et al. (2021, 2022); Bassi and Nansamba (2022); Carranza et al. (2022); Fernando et al. (2023); Kelley et al. (2022); Groh et al. (2015).

inherent instability of short-term work, surveys measuring its use at a single point in time (or via a short-run panel) will fail to capture ebbs and flows, making it difficult to evaluate how a given firm's staffing changes and why. Surveys with a long recall period risk missing the full universe of short-term spells. Tax-based data will miss informal work (which short-term work tends to be). Labor force surveys offer limited information about firms making the contracting decisions. More broadly, observational data measures the equilibrium provision of contract types; firms' motives also depend on their beliefs about the types of contracts they deem optimal *not* to provide. Multiple mechanisms plausibly operating at once limit the feasibility of identifying the incentives through a single randomized controlled trial.

Testing why firms use short-term work therefore requires data and an empirical strategy that overcomes these issues. We do so by collecting detailed information about firms' hiring and use of different contract types, measuring firm behavior at different levels of "busyness", and testing how firms perceive they would respond to shocks, and how workers would respond to different contract types. Through our data collection and survey design, we provide evidence for three key determinants of when firms choose to rely on short-term employment.

First, we show that short-term contracts give firms the *flexibility* to adjust their labor in responce to demand variation. Our surveys ask firms about the level of demand they face in busy, normal and quiet periods, these periods' relative prevalence, and how firms fill their staffing needs in each of these periods. Our vignettes study the hiring arrangements firms would offer under different scenarios. We document that the median firm faces highly variable demand: the standard deviation in its volume of sales is equal to 70% of its mean. We then show that firms take action to transmit this variation to their workers: they hire workers short-term, adjust worker salaries and days worked, and pay staff on commission. We document a strong, positive correlation between a given firm's product demand variation and the degree to which it passes this variation onto its workers, and show it is especially driven by firms' use of short-term labor. For example, we estimate that firms in the top quintile of sales variation have a coefficient of variation in their short-term staffing more than twice that of firms in the bottom quintile of sales variation. These patterns are robust to a rich set of controls. We then corroborate these results using vignettes: firms report that when faced with a short-term shock, they would be more likely to hire, and in particular, would do so via short-term contracts.

Second, we show that the size of *adjustment costs* also influences contract type. Adjustment costs vary substantially across workers but, for most hires, firms believe there are many appropriate workers that that replacing and on-boarding them has low costs. For example, the median firm reports that if their most recent hire(s) left the business today, they could find three workers the very same day capable of filling in and doing the job. Low adjustment costs therefore make short-term contracts viable for many matches. However, when adjustment costs are high, firms are more willing to commit to workers on a longer-term basis. When comparing short-term and long-term hires, we find that on average, workers on long-term contracts take 88% longer to train, and 55% longer until they are as good as a typical worker within the firm. These results thus suggest both

that short-term hires are feasible for many matches, given low adjustment costs, but also suggest that when adjustment costs are high, there are greater incentives to offer workers stability.

Third, we show that firms perceive there to be *incentive effects* associated with contract choice: they believe that the same worker will be less productive if the firm does not commit to them on a long-term basis. We present hiring managers with "profiles" of hypothetical jobseekers, designed to mimic the content and structure of typical job applications in this context. We measure firms' beliefs about these workers' reliability and performance *if hired*, and in particular how they would perform if hired on either a long-term or short-term basis. Under a model of screening, we might expect that the same worker would be especially motivated under a short-term contract, as it offers them an opportunity to demonstrate their quality and secure longer-term employment. However, we find the opposite. They perceive the same worker will perform worse, be more likely to be absent, and be more likely to quit if just offered short-term employment. These differences thus offer an important motive for firms to offer stable employment to workers, especially in cases where worker performance is especially important.

We also evaluate whether or not two alternative hypotheses for the use of short-term labor—screening and binding labor market regulations—are consistent with the empirical patterns we observe. We find evidence to suggest that while both phenomena are relevant considerations for firms on some margins, neither appears to be driving the key margin of whether to use short- or long-term contacts.

We use our empirical findings to motivate a model of firm hiring that formalizes the trade-offs we observe. In our model, profit-maximizing firms make decisions regarding whether to hire a worker and whether to offer them a short- or long-term contract, while facing uncertainty over their future demand for labor, and whether workers will quit. Whenever a new worker is brought on board, firms pay an adjustment cost. Workers offered only a short-term contract are more likely to leave, and perform less well in their jobs. Firms must therefore balance the benefits of committing to a worker on a long-term basis, with the risk of having too many (few) workers in periods of low (high) demand. On the basis of this framework, we derive comparative statics to identify the factors that would increase the use of long-term contracts in equilibrium. We argue that the trade-off inherent in our model—managing demand variation against minimizing on-boarding and performance costs—offers a useful lens through which to interpret stylized facts about hiring and contracting in labor markets in a broader range of settings. In particular, we discuss the applicability of our findings to cross-country variation in labor market turnover and to the increased use of gig workers in higher-income settings.

Our results and model have important implications for policy debates on generating more stable employment in developing countries. First, they emphasize the critical role that demand-side product variation plays in firms' willingness to commit to their workers. Moreover, our survey examines a narrow form of demand variation—day-to-day and month-to-month variation in customers. Given the broad set of risks that businesses face, a more general set of policies that better allow firms to

manage the variation they face may have positive impacts on worker stability.[2]

Second, they speak to the importance of policies and interventions that increase the demand for specialized labor. To the extent that low training and on-boarding costs reflect relatively low task complexity, the relative prevalence of such jobs in low-income countries is likely to remain a major barrier to more stable employment. Because firms perceive that they can find workers who can step in at a moment's notice, their need to commit to workers is limited. In an economy where a greater share of jobs requires high levels of specialization (and thus on-boarding and training), we might expect to see greater commitment from firms to their workers, and thus reduced employment instability.

Third, our results suggest important dimensions of heterogeneity that should be considered when designing policies that aim to boost employment such as wage subsidies or matching services. Our work suggests that policies are more likely to facilitate long-run increases in employment in contexts where demand is especially stable, or where hiring requires a greater up-front investment (and thus where re-filling the role would require significant adjustment costs).

Lastly, our results and framework have increasing applicability in higher-income countries, as a way to understand the app-based gig economy. Much of the work on the subject has focused on the roles that monopsony power and workers valuing flexibility have played in driving this expansion (Adams-Prassl et al., 2023; Mas and Pallais, 2017). Our work suggests that the growth in gig-based work can also be meaningfully explained by adjustment costs falling. It is plausible that many roles in a firm were previously staffed in-house because the adjustment costs associated with bringing in a worker on-demand were too substantial. Now, the proliferation of apps that facilitate matching have caused adjustment costs to plummet, making it possible to rely on gig workers for specific labor needs. Our results suggest that roles with high demand variation, low adjustment costs, and low perceived consequences to less motivated workers are likely to be ones especially prone to further use of gig-based work.

Our main contribution is to generate novel data that provide the first evidence that the use of short-term contracts in developing countries is driven by a tradeoff between managing variable demand versus minimizing adjustment costs and incentivizing effort. This contribution builds on four strands of existing literature.

First, we extend existing work on employment dynamics and related labor market policies. We build on the important findings of Donovan et al. (2023), who document high rates of labor market turnover in developing economics. We contribute to this agenda by unpacking *why* developing country firms are especially likely to use short-term, high-turnover contracts. We also provide evidence that traditional labor force surveys in developing countries may understate the prevalence of this type of work. Our work can inform policies that boost employment through a reduction in hiring costs or through wage subsidies (e.g., Armand et al., 2020; Card and Hyslop, 2005; De Mel et al., 2019; Galasso et al., 2004; Groh et al., 2016; Katz, 1996; Hardy and McCasland, 2023).

[2]For example, in the period of 2024 after the conclusion of surveying, firms in Nairobi were subject to both flooding, and to closures due to demonstrations against the government. Both of these phenomena further speak to the underlying riskiness that might discourage long-term employment arrangements.

Similarly, our work speaks to research on the organizational constraints shaping firm behavior in developing countries. For example, Walker et al. (2024) provide evidence in (primarily rural, Western) Kenya that an "integer constraint" (i.e., that the amount of labor supplied must be an integer) leads to slack labor within firms. The short-term labor we describe in our paper can be thought of as a way of relaxing this constraint—a labor market that makes it feasible to bring in a worker for a few days (or even a few hours) means this constraint no longer binds. As we show, firms in our setting perceive labor to be readily available; Walker et al. (2024) corroborate that thickness of the labor market is predictive of the degree of slack within firms. Given this, our work has some analog to Bassi et al. (2022), who show that rental markets allow manufacturing firms to collectively share large inputs. The presence of a thick labor market with workers willing to work on a short-term basis has meaningful parallels.[3]

Second, we add to the literature which examines the factors behind firms' choice of employment arrangements, much of which has focused on how firms in higher-income countries use temporary staffing and contracts. For example, Nickell (1986), Goux et al. (2001), Abraham (1988), and Houseman (2001) provide evidence that temporary staffing in the United States and Europe, including staff hired via temp agencies, on-call workers, and short-term work are driven by seasonality and variable demand. Faccini (2014) suggests that temporary contracts can serve as a screening tool for firms. Closely related, Adams-Prassl et al. (2023) explores the link between labor market concentration and the use of zero-hours contracts in the United Kingdom. Much of this work has emphasized the importance of regulation in governing contract choice (Cahuc and Postel-Vinay, 2002; Daruich et al., 2023), including the smaller set of research studying contract choice in low-income countries. For example, Ulyssea (2018) and Meghir et al. (2015) model the decision of whether to hire workers off-the-books as a trade off between the cost savings of not paying taxes against a cost of noncompliance increasing in firm size.

An important difference between these studies and ours is the setting—in these studies, labor market regulations play an important and often binding role in determining the set of feasible contracts, while also narrowing the universe of jobs for which these sorts of contracts are feasible. We show that in our setting labor market regulations do not appear to play a key role in the use of short-term contracts, and a sizable share of firms (59%) use short-term contracts in some capacity. These distinctions speak to a different underlying economic environment, suggesting that it is not ex ante clear the same mechanisms will be operant in our context. Importantly, we find that using a mix of short-term and long-term contracts is common even for firms that do not comply with government regulations *for any workers*. Our work thus suggests that variation in stability and commitment to workers in Nairobi is not primarily driven by labor market regulations.

Finally, our work provides a proof-of-concept for the use of vignettes on hiring and business scenarios as a means through which to measure firm preference and beliefs about workers in low-income countries. While audit studies have been used effectively in many contexts (summarized in

[3]Our work also has some connection to the broader set of studies examining firm size and organization, examples of which include Akcigit et al. (2021); Anderson and McKenzie (2022); Bassi et al. (2023) and Hsieh and Olken (2014).

Baert, 2018), they are often impractical in low-income contexts, given the structure of hiring. In our sample, 75% of hires involved individuals the manager did not know personally, but only 25% were evaluated with CVs. Most were assessed through face-to-face interactions or text exchanges. Instead, we build on Kessler et al. (2019), who had managers rate resumes they know to be fictitious. In addition to asking firms their willingness to hire hypothetical candidates, we collect more detailed information about firms' beliefs, including how their hiring would change in response to demand shocks, and how they believe contract type affects worker performance. The high agreement between our observational and vignette data supports the viability of our profile and vignette-based strategy as a further method of uncovering detailed firm preferences.[4]

2 Context, Sampling and Data Collection

Our primary analysis for this paper depends on in-depth surveys with 601 firms in Western and Central Nairobi with at least long-term regular employees, mostly across the retail, hospitality, beauty and manufacturing sectors. We supplement this data with a survey we administered to 427 young adults in low-income neighborhoods within Nairobi.

2.1 Nairobi

Our study takes place with firms in Western and Central Nairobi. Nairobi is the capital and largest city of Kenya, and is a major hub for commerce, financial services, transportation, and technology in the region. Approximately 20% of Kenya's population lives within either Nairobi County or the Counties adjoining it (KNBS, 2019). The five sectors with the largest contribution to Nairobi County's Gross County Product, per the Kenya National Bureau of Statistics, are manufacturing (25%), retail trade (20%), transportation and storage (12%), real estate (12%), and construction (12%) (KNBS, 2022).

The labor market is comprised primarily of young adults that have completed primary or secondary school, but have no further schooling, and who are engaged in basic occupations. The median Nairobi County resident is age 23; 47% are between the ages of 18 and 35.[5] Among the population aged 18-65, 49% have completed secondary school as their highest form of education, and 37% primary school (6% have a university degree; 12% have post-secondary technical education). Despite the growing share of high human capital jobs, the most common occupations are in the retail, local services, manual labor (such as construction), transportation and manufacturing sectors.

[4] Closely related, Alfonsi and de Souza Ferreira (2022) adopt a similar methodology to examine the extent of gender discrimination in job referral networks, Macchi (2023) uses hypothetical profiles of borrowers to identify the impact of body weight on the probability to be offered a loan in Uganda, and Ndayikeza (2023) examines the returns to basic occupations among college graduates in Burundi.

[5] All numbers in this paragraph are author calculations from IPUMS-International's 1% sample of the 2009 Kenya Population and Housing Census (Ruggles et al., 2024).

2.2 Firm Sample and Questionnaire

Our main sample for this project is our survey of firms, who were identified and recruited via a three-stage sampling procedure. We intend for our sample be representative of "firms in Central and Western Nairobi that hire three of more workers, and that would be approachable for a jobseeker on foot."

We began by identifying major business corridors within Western and Central Nairobi.[6] This was done by a combination of surveys with young adults about where they searched for work, informed discussion with Nairobi residents, and, within relevant neighborhoods, identification of major blocks and streets using Google Maps. Next, we partitioned business corridors into Enumeration Areas (EAs), with an EA representing roughly 0.5 kilometers of a major street in less dense areas, and a city block in the city center and its surroundings. After defining these EAs, in January 2024 we sent surveyors to verify the presence of businesses on these corridors, and to seek necessary approvals to conduct surveying within shopping centers and plazas. Through this initial procedure, we identified 317 EAs across seven neighborhoods as eligible for inclusion in our sample.[7] A map of Enumeration Areas is displayed in Appendix Figure A1.

Next, we randomly sampled 174 of these EAs, and conducted an initial census of all eligible businesses in February and March of 2024.[8] Surveyors were instructed to approach any business with evidence of at least three employees in any of eight target sectors.[9] Upon obtaining consent, surveyors conducted a short survey regarding the total number of employees and their roles within the firm, while also collecting the contact information of the individual(s) at the company responsible for hiring and managing workers.[10] Through the census, we collected the information of 1,624 firms, with 703 being eligible and willing to conduct a longer interview with us at a later date.

Following the initial surveying, we contacted all businesses with at least three employees in our target sectors and attempted to schedule an interview with an individual at the firm responsible for hiring and staffing decisions (usually the owner or a local manager). Our final sample includes

[6] We focused on these areas to represent a sample of firms within a "feasible commuting zone" for young jobseekers in low-income neighborhoods of Western Nairobi.

[7] The seven neighborhoods include two in a core commercial/industrial area: the Central Business District and the Industrial Area, three in high-income areas within Western Nairobi: Westlands/Parklands, Kilimani and Lavington/Kileleshwa, and two in low-income areas within Western Nairobi: Kibera and Kawangware.

[8] We initially targeted a sample of 200 firms across three broadly-defined areas of Nairobi: (1) The city center, encompassing the Central Business District, Westlands and the Industrial Area; (2) high-income neighborhoods in Western Nairobi that should be accessible from low-income neighborhoods in Western Nairobi, encompassing Lavington, Kilimani, and Kileleshwa; (3) low-income neighborhoods neighborhoods in Western Nairobi, encompassing Kibera and Kawangware. Given the unequal distribution of businesses across town (specifically, that there are substantially more businesses in the Central Business District and Industrial Area than in other neighborhoods) we surveyed in all EAs in high-income neighborhoods in Western Nairobi and in low-income neighborhoods, and randomly sampled 76 EAs in CBD, and 13 in the Industrial Area. We chose the number of EAs to sample based on the number of businesses in each area, calculated based on our initial surveying exercise, and an assumption about the likely response rate, estimated on the basis of initial piloting.

[9] (1) Beauty, (2) cleaning, (3) construction, (4) courier, delivery and transportation, (5) food, beverages and hospitality, (6) manufacturing, (7) retail, and (8) security.

[10] In cases where the manager was not available and that we were unable to obtain their contact information, we revisited the location up to two times in an attempt to obtain their information.

601 establishments, surveyed in March through May, 2024.[11]

Our questionnaire is specifically designed to measure how firms make decisions about hiring, staffing, and the types of contracts they offer to their employees. After collecting basic information about the business, we administer five main modules: (1) recruiting and hiring practices, (2) variation in customers and sales and how this affects staffing, (3) vignettes regarding hiring and contracting under hypothetical scenarios, (4) the characteristics of regular employees, and (5) the characteristics of workers brought in on a temporary, on-need basis.

Firm characteristics are detailed in Table 1. The majority of our sample is concentrated in four sectors: retail (44%), food, beverage, and hospitality (24%), beauty services (15%) and manufacturing (11%). On average, firms employ a total of 9 workers.[12] This mean reflects a relatively long right tail—the median firm in our sample employs four workers; 18% employ ten or more. On a typical day at the business, the median business has 15 customers (mean 34), and 278 USD in sales (mean 1,078 USD; both numbers in Purchasing Power Parity terms).

2.3 Vignettes

We complement our survey on firm behavior with vignettes to measure firms' preferences and beliefs about hiring and contracting. Specifically, we ask firms to rate hypothetical candidates, evaluating their performance, likelihood of hiring them, and, if hired, the type of contract and pay they would offer. The vignettes consist of a "base case" and two variations.

In the base case, firms review six hypothetical candidate profiles, each featuring a name,[13] demographics, connection to the firm, and work experience, all varied experimentally. An example profile is provided in Appendix B. Firms are asked to assess each candidate's performance if hired, including their likelihood of being absent or late on a given day, quitting, and overall performance. We also ask if they would make an offer to the candidate under their current demand conditions, and the type of wage and contract they would offer them.

In addition to the base case, we extend the vignettes in two ways. First, for one profile, we present a scenario where the firm expects to be busy for the next three weeks and ask if they would hire the candidate under this condition, and if so, the type of contract and wage they would offer them. Second, for a subset of profiles, firms were asked to consider a scenario in which they hired the candidate under three contract regimes: (1) as a regular worker, paid monthly; (2) as a regular worker, paid daily; and (3) as an on-need worker. For each profile-contract observation, we then

[11] We believe there are four ways in which this represents a departure from the full universe of firms in the equivalent neighborhoods of Nairobi: (1) this excludes firms with 1-2 employees, whom from piloting we believe are relatively unlikely to engage in non-family hiring, (2) firms outside of our target sectors, (3) firms that are not accessible to those approaching on foot, such as government offices or law firms, and (4) firms with a relatively large presence in the city, but where their main office is not in our enumeration areas (most notably, private minibus carriers). Given a lack of data/availability, we are unable to directly compare our data to the full universe of firms. In correspondence with the Kenya National Bureau of Statistics (KNBS), we were told that their Census of Business Establishments had very low coverage, that accordingly the KNBS themselves lack a representative sample of firms in Nairobi, and that their data could not be shared.

[12] For this and all non-bounded continuous variables in this paper we winsorize at the 99th percentile.

[13] Only first names to avoid any concerns about ethnic bias

ask a subset of the same questions about hypothetical worker performance (absenteeism, likelihood of quitting, likelihood they would need to be fired, and overall performance rating).

These vignettes help recover parameters that would be difficult or impractical to obtain from observational data or randomized evaluations. For instance, understanding whether firms use short-term labor in response to demand shocks would require both exogenous shocks and high-frequency surveying to capture temporary hiring. Similarly, studying the effects of contract regime on worker performance would require experimentally manipulating the types of contracts firms use, and mandating that firms hire workers (to observe the performance of workers they do not want to offer employment to, and the performance of workers under contracts they do not want to offer). Given the challenges inherent in implementing these procedures (and the limited scope that could be done in any one study), vignettes are a useful tool to measure these parameters, to the extent that they accurately measure firm beliefs and preferences.

To evaluate the credibility of our results, we consider two questions in Appendix B: whether respondents understood the survey and whether the results are influenced by experimenter demand effects. We conclude that the vignettes were understood and taken seriously, and that demand effects are likely to be small.

2.4 Labor Force Survey in Low-Income Neighborhoods

We supplement firm-level data with a Labor Force Survey (LFS) of 427 young adults in two low-income neighborhoods in Nairobi, collected in August, 2024. Details on sampling are in Appendix A. We aim to address two gaps with this data: (1) capturing labor market attributes—short-term jobs, contract type, and worker preferences—that we hypothesis are missing or under-reported in traditional surveys, and (2) assessing the prevalence and importance of short-term work for an important, defined population: young adults in low-income neighborhoods.

In Appendix A, we also document partial overlap between the Firm and LFS samples, with 78% of wage work episodes in the LFS taking place in neighborhoods that form part of our firm survey, and 88% in the same sector.[14]

This survey shows several key features of the economic environment for low-income youth, shown in Table 2 and Figure 1.

We show that employment is frequently fragmented and irregular. Although 89% of individuals have done some work in the last two weeks, for many individuals, this does not reflect stable, long-term employment. As shown in Figure 1, 39% of individuals who worked in the last two weeks report having worked less than full-time (instead working 1-9 days). Among individuals who worked for a wage, 60% did some work on an "on-need basis," where they were brought in only as needed, rather than as a job where they have regular hours. Even those who manage to achieve full-time employment do so by cobbling together work: 42% of workers worked in multiple jobs.

The fact that young adults have disrupted and irregular work behavior largely does not reflect

[14] One notable difference is that work in the LFS is more likely to take place in respondents' neighborhoods and less likely to take place in commercial areas.

a preference for part time work. When asked for their ideal amount of work in the last two weeks (i.e., if work were available, but they were not required to go in everyday), 95% of individuals with no or part-time work report wanting to work more than they currently are, and 78% report wanting full-time employment.

Moreover, we show that to a meaningful degree, workers in our sample are not specializing, but rather accumulating experience across a wide range of sectors and jobs. 74% of respondents report a work history that spans multiple sectors, including 27% who worked in multiple sectors in the last two weeks. While this sample is very likely to be in the labor force, and desires full-time, stable employment, a large share are unable to do so, motivating our study of short-term contracts and firms' incentives to offer them.

2.5 Contract Type

2.5.1 Classification of Workers

We use both the firm and labor force surveys to collect detailed information about employment arrangements. The dichotomy that forms the basis of the paper is between long-term, *regular* employment versus short-term, *on-need* employment. In our firm questionnaire, we define a regular worker as one "that is expected to come into business on a regular basis," (including part-time) and an on-need employee as someone who is brought in specifically when needed ("this could be just a few hours or up to three months"), but where the default arrangement is *not* that they are coming into the business day-to-day.

This classification is motivated by three features of the economic environment. First, compliance with mandated labor market regulations is low (we estimate that 20% of firms are paying into the tax system for their most recent regular hire). Given this, much of firms' variation in their commitment to workers is about whether they offer consistent, regular work, rather than formal contracts. Second, both firms compliant and non-compliant with labor market regulations use a mix of regular and on-need contracts, suggesting the decision to use short-term employment is driven by non-regulatory incentives. Third, whether a worker's jobs are regular or on-need correlates strongly with their total labor supply, suggesting this distinction matters for worker welfare.

2.5.2 Regular and On-Need Worker Descriptives

In this section, we briefly outline characteristics of on-need versus regular workers, including the degree to which firms use each, and for which roles, their relative pay, and the stability these contract types offer to workers. We also present these statistics in Table 3.

Prevalence: In total, 59% of firms in our sample use on-need workers; all use regular workers.[15] For firms that use any on-need work, 31% of their workers in the last month were brought in on a on-need basis. In the last month, short-term workers have worked an average of 7.9 days for the

[15]This second fact is by construction—firms need at least three regular employees to be eligible for inclusion in our sample.

firm. Short-term employment is not generally a single, one-off interaction with a firm, but rather an arrangement where workers are called in from time-to-time (albeit with considerable uncertainty as to how much, and exactly when). Of the short-term workers in the last month, 71% had worked for the firm in some capacity previously, and firms expect to re-use 87% at some point in the future.

Roles: In nearly all worker roles, we observe a mix of on-need and regular workers. In Figure 2, we show the share of workers that are working on an on-need basis at the role level, for all roles in which we observe at least 50 workers.[16] The three roles with the lowest share of staff on an on-need basis are those for which either specialized or firm-specific human capital is necessary: managers (0%), specialized construction workers (e.g., crane operators; 1%), and administration and finance (3%). Loading/unloading and carrying is by far the single role most likely to be on an on-need basis (81% of workers in this role are hired on-need); the next two roles with the highest share of workers on need are semi-skilled craftwork (e.g., painting, or polishing furniture; 34%), and welding and metalwork (34%). We show in Section 3.2 and Appendix Table A2 that a substantial share of differences in hiring and adjustment costs between on-need versus regular workers reflect *within-*, rather than between-role variation.

Compensation: We find that salaries are similar between regular and on-need workers, though payment frequency and payment towards government-mandated contributes differ. On average, on-need workers earn $18.08 USD PPP per day (median $11.60). Regular workers are paid $17.30 USD PPP per day on average (median $14.43); we cannot reject the null hypothesis that short-term and long-term workers have equal wages. On-need work is overwhelmingly paid daily (91%), and without contributions made to social security or health insurance funds (just 4% are). Regular workers are roughly equally likely to be paid daily (43%) or monthly (45%), government contributions are paid for 20% of regular workers.

Stability to workers: A relevant question is how *stable* on-need jobs are from the perspective of workers. A context in which workers are called in less than full-time, but in a consistent way (e.g., every Saturday) is different from a scenario in which a given job only offers work intermittently, and in concentrated stints. The latter poses more risk to workers, and requires more time and resources cobbling jobs together. Our data suggests that on-need work is much more the latter, i.e., that it is not especially stable.

We find that for just 41% of on-need jobs, workers report that they will "definitely" or "likely" have any work in the job in the next three months (shown in Appendix Figure A5). Similarly, workers anticipate that an on-need job will provide less work in the future than it has in the recent past. A given on-need job provided 66% of the work of a regular job in the last two weeks (7.1 versus 10.8), but workers expect on-need jobs to provide a smaller share, 42% of the amount of work of a regular job, in the next two weeks (4.6 versus 11.0). That is, while regular workers expect the next two weeks will look like the last two weeks in terms of how much work will be available to them, on-need workers expect the work available to them to decline over time.

[16] This represents 96% of all workers we observe.

2.6 Relation to existing data sources

A strength of our questionnaire is that it is explicitly designed to capture workers that may not be accounted for in typical firm and labor force data.

Administrative firm data, such as the dataset constructed by Wiedemann et al. (2023) from the Kenya Revenue Authority's VAT records, only records formal employees. Employer-employee datasets from administrative registries, like Brazil's Annual Report of Social Information ("RAIS"), similarly only record formally hired workers (Ulyssea, 2018). Short-term work, however, is almost entirely informal, paid in cash, and without contracts or benefits. While certain firm surveys, such as the World Bank Enterprise Survey, record temporary and seasonal employment, they typically consider only full-time workers who "work 40 hours or more per week for the term of their contract" (World Bank, 2013), plausibly missing short-term workers, whose work in a given week may not exceed 40 hours.

Labor force surveys may similarly miss short-term employment because short-term workers themselves may not self-identify as "employed." Our survey data provides support for this hypothesis. We first asked a common question to classify employment in labor force surveys: "in the last 14 days, did you work for a wage, salary, commission or any payment in-kind (including paid domestic work), even if it was for only one hour?"[17] We then ask about any short-term work, using the Swahili colloquial term of "small small work:" "in the last 14 days, did you perform any 'small small' work for payment in cash or in-kind, even if it was for only one hour?". We found that 37% of those who reported that they had not done any work to the standard LFS-style question subsequently reported doing "small small" work when asked more specifically. This pattern suggests that many short-term jobs may go under-reported in surveys of staffing and labor force participation, and thus that the empirical patterns we document are present in other low-income contexts, but are not observed in more standard questionnaires.[18] [19]

3 Empirical Results

In this section, we provide empirical evidence suggesting that the three mechanisms we highlight—business variation, adjustment costs, and behavioral responses to contracts—all play an important role in guiding the types of contracting arrangements that firms offer to their workers. First, we show that firms experience considerable variation in their demand; businesses with greater variation are more likely to engage in strategies that pass this variation onto their workers and involve less ex ante commitment on the part of firms. Second, we show how adjustment costs affect the contracting and

[17] The wording of this question was taken directly from the South African Quarterly Labour Force Survey. In our survey, the question was asked in Swahili.

[18] Labor force surveys also typically focus on primary employment, and often collect relatively limited information on additional jobs. Since many workers hold more than one short-term job, focusing only on primary employment likely leads to an under-estimation of the prevalence of this work.

[19] One notable exception, which measures Kenyan firms on a monthly basis and makes a concerted effort to capture all employees, is Kempis et al. (2023), who also document substantial variation in employment patterns within a given firm across time.

staffing decisions of firms. Overall, firms perceive that they can quickly and easily replace workers, limiting their incentive to commit to employees. However, when firms do commit for longer to workers, these tend to be individuals for whom adjustment and training costs are greater. Finally, we show that firms perceive that the same worker will perform better and be less likely to leave when offered greater stability, partially offsetting the benefits of a more flexible stock of labor.

We complement these core results with consideration of other possible mechanisms driving firm contracting decisions, including the importance of candidate screening, and the role that labor market regulations might play. Our evidence suggests that while both phenomena affect certain firm decisions, neither appears to be a key driver of the decision of whether or not to commit long-term to workers.

3.1 Response to Demand Variation

Our first set of results relate to business variation. We document three main facts: (1) demand variation is quantitatively important, (2) firms transmit much of their demand variation onto workers, rather than absorb it themselves, and (3) short-term labor arrangements are a primary mechanism through which firms transmit this variation.

Our primary evidence for these claims comes from a survey module explicitly tailored to measure variation in business demand, and the way that firms respond to it. We use this data to measure the magnitude of demand variation, and whether and how it gets transmitted to workers.

First, we ask managers to report their total number of customers and sales on days that are normal, busy and quiet. We then measure the relative prevalence of these days in the last year by showing firms a laminated sheet with the three types of days (busy, normal, quiet) and asking them to allocate 20 coins between the three types of days, each representing 5% of days. This information gives us three points in a distribution (the values of sales/customers in the busy, normal, and quiet bins) and three densities (the share of days that belong to each category). We use these values to estimate a log-normal distribution of sales and customers for every day, and to recover the mean and variance parameters, $\hat{\mu}$ and $\hat{\sigma}$ for each.[20] These estimates of $\hat{\mu}$ and $\hat{\sigma}$ allow us to measure the total degree of variation firms face, and also to examine how variance-transmission strategies co-vary with a given firm's demand.

We also ask detailed information about the strategies firms engage in response to variation. Specifically, for busy, normal, and quiet days, we ask if they bring in short-term workers, adjust the number of regular workers who come in on a given day, and adjust the salaries they pay to their regular workers. We also collect information about the number of regular workers, the number of short-term workers, and the total salaries paid to each type during normal, busy, and quiet days.

This survey module allows us to capture adjustments that are plausibly missing from standard surveys of firms in low-income countries that ask for a single, average estimate of staffing and

[20] That is, given our assumption of log-normality, the mean and standard deviation of the underlying normal distribution. Our choice of a log-normal distribution is motivated by a comparison of the out of sample fit of several distributions. We discuss our estimation procedure and the performance of different distributions in more detail in Appendix C.1.

hiring. This data is plausibly also missing from infrequent panels (e.g., yearly surveys of firms in a panel or randomized evaluation), because we ask about short-term adjustments (which might not be reported when firms attempt to aggregate over a long interval) as well as very short-run high-frequency panels (for whom the bulk of days might all fit into a single bin).

Our first takeaway is that firms face meaningful demand variation. We estimate a coefficient of variation in sales for the median firm equal to 0.70. Our estimates imply that for the median firm, its 75th percentile day of sales is approximately 2.4 times its 25th percentile day. Given this variation, firm's staffing needs will differ meaningfully from day to day.[21] Moreover, we find evidence that the high level of demand variation is not concentrated in a single sector or neighborhood, but rather a broad phenomenon across businesses. For each of the four sectors where businesses are explicitly customer-facing (and thus where we are able to ask about variation in customers and sales), beauty, hospitality, manufacturing, and retail, and for each of the seven neighborhoods we find the same broad pattern of sales variation (shown in Appendix Figure A4).

In Panel A of Table 4, we examine the degree to which firms transmit this variation to their workers, and the margins through which they do so. We examine three variance-transmission strategies (hiring staff on-need, adjusting the number of regular staff, adjusting regular staff salaries), and one variable measuring labor composition (the share of staff hired on-need). For each behavior, we consider both the overall prevalence, and how the use of these strategies co-vary with a given firm's variation in customers. (We show the same empirical test with variation in sales rather than customers in Appendix Table A3).

Specifically, we estimate:

$$y_{ins} = \beta\ \hat{\sigma}_{ins} + \gamma_n + \theta_s + \delta\ \text{emp}_{ins} + \epsilon_{ins} \tag{1}$$

where y_{ins} is an outcome for firm i in neighborhood n, engaged in sector s, $\hat{\sigma}_{ins}$ is the estimated $\hat{\sigma}$ term from our log-normal-fitted distribution of customers, γ_n are neighborhood fixed-effects, θ_s sector fixed-effects, and emp_{ins} the total number of regular workers within the firm. Standard errors are Huber-White heteroskedastic.

Panel A includes two main findings. First, the use of variance-transmission strategies is common. Forty-five percent of firms explicitly use short-term contracts in response to demand variation, 29% adjust the number of regular staff at work on a given day, and 32% adjust staff salaries. The share of staff who engage in none of these strategies is 37%. Second, there is a strong and statistically significant correlation between a given firm's $\hat{\sigma}$ and the degree to which they transmit variation onto their workers, suggesting that firms for whom this issue is especially acute are especially likely to engage in these strategies. For example, a one $\hat{\sigma}$ (roughly equivalent from moving from the top to bottom quintile of the distribution) is associated with a 62 percent (28 percentage point) increase in the probability firms use on-need labor to manage this variation.

We plot these patterns specifically for on-need workers (their use on the extensive margin,

[21] This result is qualitatively consistent with citewalker2024slack's estimate that the median firm in one of four Nairobi markets had 2.7 times as much revenue in its busiest week than on its least busy week in the last month.

and the share of staff hired on-need) in Figure 3. We present a binned scatter plot, binned at the vigintile level. Throughout the distribution, there is a strong and positive association between business variation (measured by customers or by sales) and whether they choose to use on-need staff, and the share of their staffing filled by on-need work.

In Panels B and C of Table 4, we further test the degree to which firms pass variation onto their workers, and the margins through which they do so. Here, we regress coefficients of variation of (a) all workers, (b) on-need workers, (c) regular workers paid daily, and (d) regular workers paid monthly onto our estimates of $\hat{\sigma}$ (in levels of staff in Panel B, and in total wage bill in Panel B). These estimates both tell us qualitatively how much variation gets passed onto workers, and specifically, the degree to which different strategies are used. We find both that a substantial share of variation is transmitted from firms to workers (for example, a one $\hat{\sigma}$ increase in customer variation is associated with a 0.24 increase in the coefficient of variation in firm wage bills) and that short-term contracting is the primary margin through which firms achieve this. For example, a one $\hat{\sigma}$ increase in customers is associated with an 0.32 increase in the coefficient of variation of on-need staff wages, as compared to 0.10 for regular staff, paid daily.

These results involve positive correlations that we hypothesize reflect a causal relationship: firms transmit variation onto their workers. An important consideration is whether some third variable can explain both phenomena instead. It could be, for example, that it is simply the case that poorer neighborhoods have less regular demand and depend more on less regular, or that some (non-demand variation) difference across sectors (for example, in the production process) correlates with both demand variation and the type of contract used. Given these, in all regressions, we control for firm fixed effects, neighborhood fixed effects, and the size of the firm.

In Appendix D, we show the sensitivity of all results to the controls used (for both this section, and subsequent sections). For all of our results on demand variation, there is a consistent pattern. The inclusion of sector fixed effects (i.e., what we present in our main tables) involves coefficients that are slightly smaller than raw estimates of the covariance between $\hat{\sigma}$ and the variance-transmission strategies; the other covariates do not appear to affect our estimates. Our three key facts: (1) firms face meaningful demand variation, (2) firms transmit this variation to their workers, and (3) short-term labor is a primary mechanism through which they do so, are robust to the specific controls we include.

The fact that our results are robust to the inclusion of sector, neighborhood and size controls mitigates our concerns that the positive association we observe between business variation and staffing/wage bill adjustment strategies reflects some third, unobserved variable causing both phenomena. However, we nonetheless supplement our results with a vignette to examine how firms anticipate that they would respond to shocks in demand. We do so to gauge the hiring and contracting margins firms believe they would use in respond to an increase in demand.

Specifically, we use the worker profiles outlined in Section 2.3. We show firms six hypothetical candidates, and ask about their willingness to hire the candidate (and if so, the type of contract they would use) under the current levels of demand. For one of the six profiles, we ask firms to

consider a scenario where they experienced a positive demand shock. In particular, they are asked to consider a situation where "[they] expect the next 3 weeks to be made up entirely of busy days. That is, days where a large number of customers come to [their] business," and to then again rate whether they would be willing to hire the worker, and if so, the type of employment arrangement they would offer them.

In Figure 4 and in Appendix Table A4, we show that firms perceive this demand shock would increase their willingness to hire, almost entirely through an increase in short-term labor. Firms are 42 percentage points more likely to hire a given worker, including a 38 percentage point increase in their willingness to offer the worker a short-term contract.

We therefore document that firms with greater demand variation are more likely to use short-term labor, depend on short-term labor to a greater extent, and that demand shocks increase their willingness to hire via short-term labor. Managing demand variation is an important consideration for firms, and short-term is a key mechanism through which to manage it.

3.1.1 Predictable versus Unpredictable Demand Variation

One relevant question in interpreting our results is whether the variation we observe reflects ex ante predictable demand variation (for example, restaurants anticipating weekend crowds), versus unpredictable risk. We discuss this question in Appendix C.4, including the degree to which this should affect our understanding of the firms' problem. Specifically, we present evidence on firms' reported abilities to properly adjust to staffing, worker beliefs, and the degree to which busyness ebbs and flows are predictable by sector or neighborhood. Our findings collectively suggest a meaningful degree of the variation we document does not appear to be ex ante predictable.

3.2 Adjustment Costs

Our second main empirical result is that the ease with which a worker can be found, hired, and trained plays an important role in governing the degree to which firms are willing to commit to their workers on a long-term basis. Our data suggests that many roles are easily filled, and require minimal investment on the part of firms, plausibly contributing to the high share of fragmented, high-turnover employment in Nairobi. However, it also shows that there is meaningful variation on this margin, and that longer-term roles tend to be ones where firms perceive adjustment costs to be higher.

Our evidence for this section comes primarily from a survey module related to hiring and onboarding workers. We ask firms about the last time they hired a regular, long-term worker, and (for the 59% of firms that also use short-term labor) about their last time hiring a short-term worker, providing us with 963 firm-hire observations. For both types of hires, we collect information about the time to hire and replace the worker, the time required until they reach the performance of an average worker, and the characteristics of the most recent hires.

In Table 5, we present mean statistics of the most recent hires conducted by the firm. The key takeaway from these statistics is that labor is perceived to be readily available, and the costs of

searching, training and on-boarding are low. If the most recent worker at the firm left, only 5% firms report that they would not be able to find an individual today who would be both qualified and interested in replacing the worker; the median firm reports that they could find three such individuals. Firms report that 40% of hires would take less than a day to train (the median time is 0.6 weeks); similarly, the median hire would be as good as a typical worker within a week's time.

While these statistics reflect an environment in which adjustment costs are low for a large share of jobs, we also find that there are substantial differences in adjustment costs among jobs offered on a long-term versus short-term basis. To compare these types of workers, we estimate:

$$y_{hins} = \beta \text{ long-term}_{hins} + \gamma_n + \theta_s + \delta \text{ emp}_{ins} + \epsilon_{hins} \tag{2}$$

where y_{hins} is an outcome for hire h within firm i within neighborhood n in sector s, and long-term$_{hins}$ is an indicator variable for whether the hire was made on a long-term, regular basis.

We estimate that firms perceive there to be higher adjustment costs for staff hired on a long-term role for the dimensions we measure. Firms report that long-term hires will take 2.3 weeks longer to train on average, and 2.7 weeks longer until they become as good as the typical employee. They believe that these employees will remain with the firm 47 weeks longer, and would take 2.3 weeks longer to replace.

We document in Appendix Table A2 that these differences reflect both *between*- and *within*-role variation in adjustment costs, with a greater share attributable to within-role variation. That is, roles that tend to have especially low adjustment costs (e.g., loading and unloading, semi-skilled craftwork) tend to be ones hired on-need, but also, it is especially the case that the same role involves less investment on the part of firms when they hire on a short-term basis (e.g., the same carpenter will receive less training if only brought in on a short-term basis). We discuss this decomposition in greater depth in Appendix E.

These results thus suggest an offsetting factor for firms. While short-term work allows firms to manage demand variation, for hires where a high degree of training or investment is necessary, firms are more likely to rely on regular, long-term labor.

3.3 Perceived Incentive Effects of Contract Choice

Firms might also use the structure of the employment arrangements they offer to workers as a way to incentivize particular sorts of behavior. We find evidence that this is indeed the case: firms believe that the same individual will perform better when offered a more stable contract than when hired on an on-need basis.

We test the hypothesis that contract choice affects worker behavior through our profile rating exercise, in which firms report their interests and beliefs regarding hypothetical candidates. In this exercise, hiring managers are asked to rate candidates on multiple criteria, including (i) how the hiring manager would rate them on a Likert scale after a month on the job, (ii) how frequently they would not show up at a given day on the job, (iii) how likely they would be to try and leave the job

within a month's time, and (iv) how likely the manager would be to ask them to leave the job. We asked each of these two questions about two of the hypothetical candidates, and in each case, asked the hiring manager to estimate the same candidate if offered regular work, paid monthly, regular work, paid daily, and short-term work on an on-need basis.

With these ratings, we estimate the perceived behavioral response of candidates to different candidates, reported in Table 6 . We estimate regressions of the form:

$$y_{ipc} = \beta \text{ daily}_{ipc} + \zeta \text{ on-need}_{ipc} + \delta_{ip} + X_{ipc}\theta + \epsilon_{ipc} \tag{3}$$

Where y_{ipc} is the rating for firm i, evaluating profile p, under contract-regime c. daily_{ipc} and on-need_{ipc} are indicator variables for when the rating is asked about for daily and on-need contracts (i.e., being offered a long-term monthly contract is the omitted group). δ_{ip} are firm-profile fixed effects,[22] and X_{ipc} is a vector of characteristics, manipulated for each profile (for example, the gender of the respondent, the neighborhood they come from, and their experience profile). By asking firms to evaluate the same worker, in the same role, at the same firm under different contract regimes, we are avoiding many potential forms of unobserved selection in our comparison of candidates.

For each of the characteristics we ask about, firms perceive that the same worker would perform worse if offered a short-term, on-need contract than if offered a stable, regular employment arrangement. The same individual is expected to not show up to work in a given 20-day period 0.74 additional times if only offered short-term employment (27% more than the mean of those offered longer-term contracts). They are believed to be 9.7 percentage points (23%) more likely to quit, and 3.1 percentage points (8%) more likely to be let go by the employer. On a 1-7 Likert Scale, the same employee is ranked 0.22 points (5%) lower when offered a short-term contract (effect size of 0.16). In comparison to the other traits we manipulate, these effects are relatively large. For example, our estimated coefficient of having any experience in the same sector as the job on the worker rating (relative to no experience, in Appendix Table A5) is 0.26, or very similar to our estimate of the difference between being a regular worker versus one brought in on an on-need basis.

It is important to note that there are ex ante theoretical reasons to think that short-term contracts might be expected to improve worker performance. If, for example, firms primarily use short-term contracts as a mechanism through which to gauge worker quality, and that workers have a high probability of being converted into longer-term workers, we might expect that firms will expect the same individuals to perform better, before offered the stability of longer-term employment. Instead, our results suggest that managers perceive that workers will effectively have a foot out of the door, missing more work and performing worse (and potentially searching for a better opportunity) when only offered short-term employment. We explore evidence in favour of the use of short-term contracts as a screening mechanism in the subsequent subsection.

These results collectively suggest that in addition to on-boarding costs, that workers' responses under different potential arrangements is an important consideration as they decide how and when to commit to their workers. The benefits of flexibility need to be weighed against the fact that they

[22]For example, a fixed effect might be "Profile 6 showed to firm 101."

perceive the very same worker to perform worse if not offered the commitment of a longer-term employment arrangement.

3.4 How much does screening matter?

Research in other low-income countries has provided evidence that firms have a challenge observing candidate quality, that improving the ability of candidates to signal their quality can affect both jobseeker and firm behavior (for example, Carranza et al. (2022); Abebe et al. (2021); Bassi and Nansamba (2022)), and that *apprenticeships* in low-income countries offer a way for firms to evaluate worker quality (Alfonsi et al., 2020, 2022; Hardy and McCasland, 2023). Given this, a relevant question to consider is whether the use of the short-term employment we document is similarly driven by screening motives.

The evidence we collect both leads us to believe that screening has some relevance in the economy, but also that it is not the key driver of variation in contract choice we document. Our assessment is driven by three pieces of evidence.

First, in Section 3.2, we document that roles with lower training and lower investment required by firms tend to be the ones that are more frequently offered on a short-term basis. In an environment in which a key motive to hire short-term is to first observe candidate quality, we might expect that the more complicated, investment-intensive roles would be the ones in which short-term hiring is especially prevalent; instead we observe the opposite. In a similar vein, we observe that network-based hiring (a key mechanism through which to overcome limited information) is more common for *short-term*, rather than long-term roles.

Second, in our-module on short-term employment, we ask about the number of short-term hires that have worked for the business previously, and that the firm expects to re-use in the future. The numbers for both are high: 71% of short-term workers had previously done work with the business; firms expect to use 87% in the future. They expect to promote a lower percent, 31% to longer-term employment at some point. If the use of short-term hiring were primarily driven by screening, we might expect to see a relatively high rate of "up or out" hiring—once candidate quality is revealed, the good candidates get promoted, and the bad candidates are not re-used. Instead, the high re-use of candidates speaks to the primacy of other mechanisms at play, most notably demand variation.

Lastly, our results from the vignettes on worker performance under different contracts appears inconsistent with the hypothesis that screening is a key driver of hiring behavior. Firms perceive that on average, workers will be worse when offered a short-term contract. If short-term hiring is primarily a trial period, we should expect to see workers especially motivated in this period (as an opportunity to show their high quality and ensure a longer-term contract). Therefore, the fact that firms think the same individuals will perform worse would seem to suggest that workers do not perceive short-term contracts to function as a trial period.

3.5 What role do labor market regulations play?

Another important consideration is the degree to which binding labor market regulations play a role in dictating the type of commitment that firms offer to their workers. One hypothesis is that firms are required to offer rigid, inflexible commitments to its workers due to Kenya's labor market laws, but that offering short-term contracts offers them an avenue through which to circumvent these laws. We find evidence that labor law governs some aspects of contracting, but that it does not appear to be a key driver of the variation in short-term and long-term contracts.

The relevant legislation in Kenya is the Employment Act of 2007, which states that employees must receive social security and health care contributions, have tax contributions deducted, and receive sick leave and maternity leave. The Law also states that any casual labor working in excess of one month is automatically converted to regular employment, with salaries paid monthly, and all other benefits applicable.[23] Thus, if these laws have actual bite in the economy, work-term stints of less than one month would offer a way to circumvent these laws.

We measure formality indirectly, given concerns that firms might be reluctant to directly admit to non-compliance. First, we ask firms one-by-one about whether they offer particular benefits to their employees, including both common, informal benefits (meals, transportation allowances, loans), and formal benefits (National Health Insurance Fund and National Social Security Fund contributions). Second, we ask firms one-by-one about required materials for applicants (e.g. secondary school credentials, references), including whether they require a Kenya Revenue Authority (KRA) PIN (required to pay taxes and to make contributions). We find that just 14% of firms both make mandated contributions (to health insurance and social security) and require a Kenya Revenue Authority PIN; when we use the less stringent criteria that they make contributions, 20% of firms appear to be broadly in compliance with government regulations with respect to staffing.

Our first piece of evidence that labor market regulations do not play a key role in the stability offered by firms to workers is in our comparison of compliant and non-compliant firms. We are unable to reject the null that the two types of firms are equally likely to use short-term contracts. This finding is inconsistent with the idea that regulations are the driving force—even firms who do not comply with the law still find it optimal to use a mix of short-term and long-term contracts, suggesting that other incentives must be a driving factor.

Second, as outlined in Section 3.4, we see a high rate of renewals among short-term workers, a behavior incompatible with Labor Law. In our module on short-term employment, we ask about the number of workers used in a given role within the firm, and of this number, how many had previously worked for the firm, how many the firm expects to re-use in the future. The Kenya Employment Act of 2007 states that any employee who works for more than a month consecutively automatically converts to a regular employee (with the full required benefits, i.e., including social security and health contributions). However, we find that 71% of employees are re-used, and 87%

[23] The specifics vary from country-to-country, but policies that (a) mandate benefits for formal workers, and (b) impose a maximum point of casual employment after which casual workers are converted to formal workers are both fairly common in low-income countries. For example, just in East Africa, Tanzania, Rwanda, Uganda and Ethiopia all have similar laws (Kuddo and Kuddo, 2018).

are expected to be re-used in the future, a strategy incompatible with short-term use of workers as a way to avoid labor law. These two facts suggest to us that the use of short-term work is not primarily in response to binding labor market regulations.

However, we do see evidence of at least two margins (the degree to which salaries get adjusted in response to demand variation, whether employees are paid daily versus monthly) that do appear quite different between compliant and non-compliant firms. These results suggest that labor market regulations do affect certain behaviors, but do not appear to be the key driver of the use of short-term and long-term labor arrangements in Nairobi.

4 Conceptual framework

Hiring on a short-term contract presents firms with a trade-off. As we've shown in Section C, short-term contracts allow firms to manage variable product demand and hence reduce slack at the firm. At the same time, the vignette evidence we present in Section 3.3 demonstrates that firms expect workers on short-term contracts to be less productive, which directly reduces the firm's profits. Lastly, as we show in the same section, workers are more likely to be absent and quit under a short-term contract, which exposes firms to the risk of not having a worker at the firm when they need one. This creates an added cost to the firm, since each time a firm hires a worker they must pay an on-boarding cost. In this section, we aim to formalize this relationship with a simple conceptual framework.

4.1 Set-up

We consider an infinite horizon, discrete time setting with discount rate β. In every period a firm can either: (1) employ a short-term worker, (2) employ a long-term worker, or (3) not employ a worker. If a firm employs a worker, their per-period payoff is composed of the sum of a productivity term, θ, a worker quality term, η, net of the wages they pay the worker, w. If a firm does not employ a worker, their per-period payoff is zero.

Variable demand: The model centers on two primary frictions: (1) it takes time and money for a firm to find and train a worker, and (2) firms face uncertain demand for their goods and services. To capture the latter friction, our framework assumes that in every period firm j independently draws its productivity, θ from a normal distribution with mean μ_j and variance σ_j.[24] Firms hold correct beliefs about the parameters characterizing their productivity distribution, and form expectations about the sequence of future productivity draws. Uncertainty about future productivity draws is costly for firms offering long-term contracts as it limits their ability to adjust their employment level to their productivity, resulting in periods where firms do not have enough

[24] In our empirical analysis, we focused on shocks to product demand and showed that these translate into increased demand for labor. As such, we simplify our analysis by focusing directly on shocks to labor demand. θ, together with the quality of a worker, constitute the marginal product of labor. As a result, a higher value of θ, which could be driven, for instance, by a greater number of costumers, directly translates into higher demand for labor.

workers to meet demand or have idle workers.[25]

Search frictions and on-boarding costs: To capture the first friction, we assume that filling a vacancy takes place over multiple periods of time, and that, in every period, a vacancy is filled with probability $h < 1$. The probability that a firm fills a vacancy is an exogenous reduced form parameter that reflects the intensity of search for workers. After a vacancy is filled, firms have to pay a one-time on-boarding cost, c, which may be heterogeneous across contract type. For simplicity, we do not allow offers to be rejected or bargaining to take place.

Contract choice: When a worker and a firm match, the firm offers the worker one of two types of contracts: (1) a long-term contract which commits it to paying the worker a per-period wage of w indefinitely and commits the worker to showing up for work indefinitely, and (2) a short-term contract which commits it to paying a worker for one-period and the worker to showing-up to work for one period. Consistent with the evidence that firms and workers often have repeated relationships, a firm can choose to re-hire a worker in every future period unless the worker has left the firm, which occurs with probability $q < 1$. When a firm re-hires a worker for a period, they commit to paying them the per-period wage w, in a period where they choose not to hire a worker, they do not pay them anything. For simplicity, we do not allow long-term contracts to be destroyed, consistent with the evidence that workers offered these contracts are much less likely to be fired or disappear.

Heterogeneity in worker quality: When a worker and firm match, a worker independently draws her quality, η from a quality distribution $F(\eta)$. A worker's quality directly impacts her productivity and therefore the firms' profits in every period. Consistent with the evidence that the same worker is expected to be more productive when offered a long-term contract, we allow workers of equal quality to be more productive under a long-term contract relative to a short-term contract. We model this by assuming that, for a given worker, productivity would be γ percent higher with a long-term contract.[26]

What the model misses: We make two key simplifications in order to keep our model tractable. First, we do not allow firms to adjust the wages they offer workers. While in principle, wages could differ between short-term and long-term contracts, we assume that firms do not optimize over the wages they offer workers. Second, we abstract away from any general equilibrium considerations.[27]

[25] In the model we assume that firms are risk neutral, however allowing firms to be risk averse makes long-term contracts costly for a second reason: relative to short-term contracts which allow firms to pass product demand risk onto workers, long-term contracts make firms bear the risk, increasing the spread of the distribution over potential profits in every period.

[26] This set-up allows workers under long-term contracts to be more productive for three reasons: (1) the direct "incentive" effect of long-term contracts, (2) the selection of higher quality workers into long-term contracts, (3) the higher use of long-term contracts for more productive firms.

[27] In particular, comparative statics that lead firms to hire more workers on long-term contracts may make filling a vacancy harder, or may decrease the average quality of job seekers that firms encounter. Since firms report that filling vacancies with qualified workers is very easy, we do not consider these effects to be first-order.

4.2 Solving the firm's problem

We derive comparative statics by focusing on the choice of a firm that initially does not have a worker, and draws productivity θ_0, and a worker of quality η_i. We define three value-functions: V_0, the value of not having a worker, corresponding to the choice not to hire; V_1, the value of hiring the worker on a short-term contract; and V_2, the value of hiring the worker on a long-term contract.

4.2.1 Value functions

Short-term contracts: The value of offering a short-term contract is composed of three terms:

$$V_1(\theta_0, \eta_i) = \underbrace{\theta_0 + \eta_i - w - c}_{\text{First period flow value}} + \underbrace{\sum_t (\beta(1-q))^t E_\theta \max\{\theta + \eta_i - w, 0\}}_{\substack{\text{Discounted expected value} \\ \text{conditional on not separating}}} + \underbrace{\sum_t \beta^t (1-q)^{t-1} q V_0}_{\substack{\text{Discounted expected value} \\ \text{conditional on separating} \\ \text{in a future period}}} \quad (4)$$

The first term corresponds to the first-period flow value associated with hiring a worker under a short-term contract. It is the sum of the firm's initial productivity draw, θ_0 and worker-quality draw, η_i, less the wages the firm pays the worker, w, and the costs associated with hiring a new worker, c. The second term corresponds to the discounted expected future value of the worker, conditional on not separating from the firm. It is the product of the discount factor, β, the probability a worker does not separate from the firm, $1-q$, and the expected value from the option to re-use a worker in a future period. The last term corresponds to the discounted expected value associated with losing a worker and ending up in a state in which the firm does not have any workers.

Under our parametric assumption about the distribution of the productivity term, we can re-write the value function as:

$$V_1(\theta_0, \eta_i) = \theta_0 + \eta_i - w - c + \frac{\beta(1-q)\sigma[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w})]}{1-\beta(1-q)} + \frac{\beta q}{1-\beta(1-q)} V_0 \quad (5)$$

where $\tilde{w} = \frac{w-\eta_i-\mu}{\sigma}$.

Long-term contracts: The value of a long-term contract, V_2, is simpler, and given as follows:

$$V_2(\theta_0, \eta_i) = \underbrace{\theta_0(1+\gamma) - w + \eta_i - c}_{\text{First period flow value}} + \underbrace{\sum_t \beta^t E_\theta[\theta(1+\gamma) + \eta_i - w]}_{\text{Discounted expected future value}} \quad (6)$$

As before, the first term corresponds again to the first period flow value of the worker. We assume that workers hired under a long-term contract will be more productive, and so we assume that a worker's productivity is given by $\theta(1+\gamma)$, where $\gamma > 0$ captures the "productivity boost" associated with a long-term contract.[28] The second term corresponds to the discounted expected future value

[28] We model γ as a boost to the productivity term, θ. Our comparative static results would not change if we modeled it as a boost to worker quality, η. From equation, 7 we can see that the impact of the productivity boost is simply

of the worker. Note that this differs from the corresponding expression for short-term contracts in two ways. First, we assume that workers do not separate under long-term contracts, and so the firm does not take into account the possibility of ending up without a worker. Second, the firm can no longer choose whether to work, and pay, a worker on a particular day.

Under our parametric assumptions, we can re-write this as:

$$V_2(\theta_0, \eta_i) = \theta_0(1+\gamma) + \frac{1}{1-\beta}(\beta\mu(1+\gamma) + \eta_i - w) - c \tag{7}$$

Not hiring: Finally, the value of being without a worker is given as follows:

$$V_0 = \underbrace{0}_{\substack{\text{First period} \\ \text{flow value}}} + \underbrace{\beta h E_{\theta,\eta} \max\{V_0, V_1, V_2\}}_{\substack{\text{Discounted expected value} \\ \text{conditional on matching} \\ \text{with a worker}}} + \underbrace{\beta(1-h)V_0}_{\substack{\text{Discounted expected value} \\ \text{conditional on not matching} \\ \text{with a worker}}} \tag{8}$$

where the first term corresponds to the first period flow value, which is simply 0 as the firm does not produce or pay workers. The second is the discounted expected flow-value conditional on matching with a worker. It is the product of the discount factor, the probability to match with a worker, h, and the expected value associated with matching with a worker. The latter term corresponds to the maximum associated with hiring a worker under a long-term contract, V_2, a short-term contract, V_1, and not hiring at all, V_0. The last term corresponds to the discounted expected value associated with not matching with a worker today.

4.2.2 Optimality conditions

The firm's optimality conditions are characterized by a set of three cut-off values with respect to quality: (1) the minimum quality such that a firm hires a worker under a short-term contract, relative to not hiring them at all, η_{01}; (2) the minimum quality such that a firm hires a worker under a long-term contract, relative to not hiring them at all η_{02}; and (3) the minimum quality such that a firm hires a worker under a long-term contract, relative to hiring them under a short-term contracts, η_{12}. Formally,

$$V_1(\theta, \eta_{01}(\theta)) = V_0 \tag{9}$$

$$V_2(\theta, \eta_{02}(\theta)) = V_0 \tag{10}$$

$$V_2(\theta, \eta_{12}(\theta)) = V_1(\theta, \eta_{12}(\theta)) \tag{11}$$

Note that these cut-offs map directly into the shares of workers hired on different contract types. As η_{01} and η_{02} decrease, hiring becomes more likely overall as a larger share of job-seekers a firm meets end up above the threshold. As η_{01} decreases and η_{12} increases, hiring under a short-term contract becomes more likely. Finally, as η_{02} and η_{12} decrease, hiring under a long-term contract becomes more likely.

to increase the value of a contract by some constant amount.

Proposition 1 (Cut-offs) The model yields unique cut-off values in η that characterize whether a firm wants to hire a worker, and if so, whether under a short-term contract or long-term contract.

Proof in appendix F.1

4.3 Comparative statics

As the spread of product demand increases, the value of a short-term contract increases, as there is more to be gained from the option to re-optimize on the extensive margin. Since firms are risk neutral, there is no impact on the value of long-term contracts. For a given level of quality, the value of short-term contracts increases relative to the value of long-term contracts, increasing the cut-off value to hire on a long-term contract, η_{12}, and increasing the share of short-term contracts.

Prediction 1 (Product demand variability): A firm which faces more variable product demand will have a higher quality threshold at which they are willing to offer long-term contracts, leading to a smaller share of long-term employment.

Proof in appendix F.2

Our framework also predicts that increasing the on-boarding costs associated with long-term (short-term) contracts, will directly decrease their value, which will increase the value of short-term (long-term) contracts. An increase in on-boarding costs also decreases the value of not having a worker at the firm, as it implies a higher cost to be borne out in a future period. Since short-term contracts expose firms to the risk of not having a worker at the firm, a uniform increase in on-boarding costs will decrease the value of short-term contracts by a proportionally larger amount than the decrease in the value of long-term contracts, thereby increasing the share of long-term contracts.

Prediction 2 (Hiring frictions and on-boarding): A decrease in hiring costs for a particular contract type, increases employment on that contract type. As hiring frictions and on-boarding costs decrease *overall*, firms become more willing to hire on a short-term basis.

Proof in appendix F.3

We illustrate how the share of employment on different contract varies with adjustment costs and variable demand in Figure 6. As firms face more variable demand, the share of short-term contracts monotonically increases. At the same time, when adjustment costs are high, the share of long-term contracts is weakly higher than when adjustment costs are low.

Finally, an increase in the impact of long-term contracts on the productivity of workers directly increases the value of long-term contracts, while leaving the value of short-term contracts unchanged, thus increasing the share of long-term contracts.

Prediction 3 ("Productivity effect" of long-term contracts): An increase in the productivity of long-term contracts will reduce the quality cut-off to be hired on a long-term basis, increasing the share of long-term employment.

Proof in appendix F.4

5 Discussion

We have documented that firms use short-term contracts in response to variable demand, and that adjustment and performance costs dampen the degree to which this is a feasible strategy. Our model of firm hiring formalizes the idea that demand variation, adjustment costs, and performance costs are all parameters governing the optimal type of contract for a firm to use.

In this section, we consider the degree to which these results and this framework can be used to broaden our understanding of economic phenomena demonstrated elsewhere. In particular, we detail how demand variation and adjustment costs can be used to interpret frequent turnover in low-income countries, how reduced adjustment costs can be used to interpret the growth of the gig-based economy, and what our results might suggest about the possibility of a "productivity trap," limiting the ability of workers and firms to specialize.

5.1 Labor Market Dynamics

The findings of this paper offer a lens through which to interpret the high rates of turnover present in low-income countries. Donovan et al. (2023) (DLS) demonstrate a strong negative correlation between GDP per capita and labor market turnover. In low-income countries, there are substantially more separations, concentrated early in workers' tenure, disproportionately in low-wage jobs.

The types of short-term, on-need labor we observe in our setting is a plausible pathway through which this high turnover occurs. Employment under a short-term contract does not allow workers to accumulate tenure at the firm and leads to frequent turnover early into a match. As a result, firms' choice of employment arrangements maps directly onto the average level of turnover in a labor market—a high use of short-term contracts facilitates frequent separations. Moreover, to the extent that low-adjustment cost hires also tend to be low-wage hires, our results are consistent with DLS' finding that high-turnover contracts tend to be among low-wage workers. Our results are consistent with the authors' hypothesis that low-income countries may have a higher share of low-value, high-turnover jobs.

To the extent that short-term contracts are a key driver of high turnover, our results suggest that demand variation, adjustment, and performance costs should be considered when modeling and assessing the prevalence of high turnover in labor markets. Our results suggest two candidate hypotheses consistent with the patterns documented by DLS. First, firms in low-income counties face greater demand variation (or have a greater tendency to pass their demand variation onto their workers). Second, firms in low-income countries have more matches with low adjustment and training costs than firms in high-income countries.

Moreover, our finding that short-term jobs are likely to be under-counted in traditional labor force surveys suggests that differences in labor market dynamics across the development spectrum may be greater than those documented by Donovan et al. (2023). We hypothesize that short-term work is likely to be more prevalent in lower-income countries, and hence more likely to be under-counted at the bottom of the development spectrum. Since short-term work is associated with lower employment duration and greater turnover than long-term work, this suggests that Donovan et al. (2023) plausibly estimate a lower-bound of the gradient between labor market turnover and development.

5.2 Adjustment Costs and the Gig Economy

Our framework can also be used to interpret a more recent phenomenon in higher-income labor market: firms' increased use of gig-based hiring as a way of managing variable labor demand. Our framework would suggest that adjustment and hiring costs previously limited the feasibility of on-demand hiring; apps that facilitate task-based hiring have done so in part by dramatically reducing these hiring costs.

Consider the case of a restaurant that depends in part on deliveries for its business. Prior to the proliferation of app-based ordering, these firms faced a distinct trade-off with respect to staffing. On the one hand, having a delivery driver staffed in house meant that there would be windows of time where the driver would be idle, but nonetheless paid for the work. But it was not feasible to hire staff exactly and only when needed (i.e., on an order-by-order basis) because of the adjustment costs associated with hiring and bringing on-board a worker. This fundamentally changed in response to apps that allow contracting on an on-need basis. Workers can be hired immediately and temporarily, workers require no training, and their performance is observable. Given this, restaurants are more likely to rely on apps rather than on staffing this role in-house, and accordingly, they do not have to pay for a worker's idle time.

While fundamentally distinct in many respects, the labor market in Nairobi is similar insofar as adjustment costs and training for many roles are minimal. Just as apps have made it feasible to hire labor on an on-need basis, firms' reported ability to bring in multiple workers the same day who could adequately fulfill a role make on-need hiring a feasible strategy.

Our framework thus suggests that business variation and on-boarding / adjustment costs are key parameters to consider when assessing the degree to which future sectors or occupations might transition towards a greater use of zero hours contracts and gig-based hiring.

5.3 Low-Commitment Equilibrium

Our results are consistent with, albeit not dispositive of, the possibility of a "productivity trap," in which workers and firms are stuck in a low commitment equilibrium. We provide evidence in this paper that there is limited commitment on both sides of the market. Firms perceive adjustment and training costs for their workers to be low, suggesting that firms make limited investments in their hires. In turn, firms expect that when they use short-term contracts, workers will be less committed

to them. They perceive that when a given worker is only offered work on a short-term basis, they are more likely to be absent from work in a given day, more likely to quit, and will perform worse on the job. It is thus possible that the same matches could be more productive if both sides could credibly commit to each other, but that in the absence of such a mechanism, commitment remains low, and matches remain low-commitment, low-specialization. This pattern is broadly consistent with the work of Atencio-De-Leon et al. (2023), who find that workers in Peru generally do not specialize to the same extent as workers in higher-income countries.

Our study therefore closely relates to important work by Cefala et al. (2023), who find that Burundian farm owners underinvest in training laborers on their farms, because the owners cannot ensure that they will capture the full returns to this training. The authors note that there are particular stages in the agricultural production process which labor demand is especially high (e.g., during the planting stage); our results suggest that this demand variation is an important element driving the inability of firms to commit to workers (and plausibly driving this underinvestment). Our results that urban firms across many sectors also experience substantial demand variation suggests that the phenomenon they document may be broadly applicable in low-income countries, rather than being a phenomenon specific to agriculture.

6 Conclusion

We examine the underlying incentives that guide the types of labor contracts firms in Nairobi, Kenya offer to their workers. We provide evidence that firm decision-making is governed by a key trade-off facing firms: offering workers short-term arrangements allows firms to dynamically adjust their labor in response to varying demand for their goods and services. However, frequently bringing in new workers involves costs associated with training and on-boarding workers, and risks employees being less reliable or productive. Firms must therefore trade-off the value of flexibility against the cost of worse performance; the relative magnitude of these costs dictates the optimal contract for a given firm.

These results therefore speak to the importance of considering demand variation, adjustment costs, and performance costs when considering programs and policies that seek to generate stable employment. So long as firms in low-income countries experience considerable variation in their demand, or remain inclined to pass this variation onto their workers, there are likely to be substantial challenges in ensuring worker stability. Similarly, in labor markets where a sizable share of the demand for labor is in the form of tasks with limited training or specialization required, firms will plausibly face limited consequences associated with only committing to their workers on a short-term basis.

Finally, our results and model provide a framework through which to consider the degree of firms' investment in workers and employment stability more generally. When demand variation is substantial or firms face little constraints to passing this variation onto workers, short-term labor arrangements are plausibly a strategy that many firms will deem optimal, in the absence of strong

offsetting incentives or regulations.

References

Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, and S. Quinn (2021). Anonymity or distance? job search and labour market exclusion in a growing african city. *The Review of Economic Studies 88*(3), 1279–1310.

Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, S. Quinn, and F. Shilpi (2022). Matching frictions and distorted beliefs: Evidence from a job fair experiment.

Abraham, K. G. (1988). Flexible staffing arrangements and employers' short-term adjustment strategies.

Adams-Prassl, A., M. Bagalova, M. Qian, and T. Waters (2023). Firm concentration & job design: the case of schedule flexible work arrangements. *IFS working paper*.

Akcigit, U., H. Alp, and M. Peters (2021). Lack of selection and limits to delegation: firm dynamics in developing countries. *American Economic Review 111*(1), 231–275.

Alfonsi, L., O. Bandiera, V. Bassi, R. Burgess, I. Rasul, M. Sulaiman, and A. Vitali (2020). Tackling youth unemployment: Evidence from a labor market experiment in uganda. *Econometrica 88*(6), 2369–2414.

Alfonsi, L. and P. de Souza Ferreira (2022). Whom would you rather work with? gender discrimination in the referral system in uganda.

Alfonsi, L., M. Namubiru, and S. Spaziani (2022). Meet your future: Experimental evidence on the labor market effects of mentors.

Anderson, S. J. and D. McKenzie (2022). Improving business practices and the boundary of the entrepreneur: A randomized experiment comparing training, consulting, insourcing, and outsourcing. *Journal of Political Economy 130*(1), 157–209.

Armand, A., P. M. Carneiro, F. Tagliati, and Y. Xia (2020). Can subsidized employment tackle long-term unemployment? experimental evidence from north macedonia.

Atencio-De-Leon, A., M. Lee, and C. Macaluso (2023). Does turnover inhibit specialization? evidence from a skill survey in peru.

Baert, S. (2018). *Hiring discrimination: An overview of (almost) all correspondence experiments since 2005*. Springer.

Bassi, V., J. H. Lee, A. Peter, T. Porzio, R. Sen, and E. Tugume (2023). Self-employment within the firm. Technical report, National Bureau of Economic Research.

Bassi, V., R. Muoio, T. Porzio, R. Sen, and E. Tugume (2022). Achieving scale collectively. *Econometrica 90*(6), 2937–2978.

Bassi, V. and A. Nansamba (2022). Screening and signalling non-cognitive skills: experimental evidence from uganda. *The Economic Journal 132*(642), 471–511.

Becker, G. S. (1964). A theoretical and empirical analysis, with special reference to education. In *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, pp. 22–0. NBER.

Cahuc, P. and F. Postel-Vinay (2002). Temporary jobs, employment protection and labor market performance. *Labour economics 9*(1), 63–91.

Card, D. and D. R. Hyslop (2005). Estimating the effects of a time-limited earnings subsidy for welfare-leavers. *Econometrica 73*(6), 1723–1770.

Caria, A. S., G. Gordon, M. Kasy, S. Quinn, S. O. Shami, and A. Teytelboym (2024). An adaptive targeted field experiment: Job search assistance for refugees in jordan. *Journal of the European Economic Association 22*(2), 781–836.

Carranza, E., R. Garlick, K. Orkin, and N. Rankin (2022). Job search and hiring with limited information about workseekers' skills. *American Economic Review 112*(11), 3547–3583.

Carranza, E. and D. McKenzie (2024). Job training and job search assistance policies in developing countries. *Journal of Economic Perspectives 38*(1), 221–244.

Cefala, L., P. Naso, M. A. Ndayikeza, and N. Swanson (2023). Do short-term contracts limit investments in training workers?

Daruich, D., S. Di Addario, and R. Saggio (2023). The effects of partial employment protection reforms: Evidence from italy. *Review of Economic Studies 90*(6), 2880–2942.

De Mel, S., D. McKenzie, and C. Woodruff (2019). Labor drops: Experimental evidence on the return to additional labor in microenterprises. *American Economic Journal: Applied Economics 11*(1), 202–235.

De Quidt, J., J. Haushofer, and C. Roth (2018). Measuring and bounding experimenter demand. *American Economic Review 108*(11), 3266–3302.

Donovan, K., W. J. Lu, and T. Schoellman (2023, 05). Labor Market Dynamics and Development*. *The Quarterly Journal of Economics 138*(4), 2287–2325.

Faccini, R. (2014). Reassessing labour market reforms: Temporary contracts as a screening device. *The Economic Journal 124*(575), 167–200.

Fernando, A. N., N. Singh, and G. Tourek (2023). Hiring frictions and the promise of online job portals: Evidence from india. *American Economic Review: Insights 5*(4), 546–562.

Franklin, S. (2018). Location, search costs and youth unemployment: experimental evidence from transport subsidies. *The Economic Journal 128*(614), 2353–2379.

Galasso, E., M. Ravallion, and A. Salvia (2004). Assisting the transition from workfare to work: A randomized experiment. *Ilr Review 58*(1), 128–142.

Gibbons, R. and M. Waldman (2004). Task-specific human capital. *American Economic Review 94*(2), 203–207.

Goux, D., E. Maurin, and M. Pauchet (2001). Fixed-term contracts and the dynamics of labour demand. *European Economic Review 45*(3), 533–552.

Groh, M., N. Krishnan, D. McKenzie, and T. Vishwanath (2016). Do wage subsidies provide a stepping-stone to employment for recent college graduates? evidence from a randomized experiment in jordan. *Review of Economics and Statistics 98*(3), 488–502.

Groh, M., D. McKenzie, N. Shammout, and T. Vishwanath (2015). Testing the importance of search frictions and matching through a randomized experiment in jordan. *IZA Journal of Labor Economics 4*, 1–20.

Haaland, I., C. Roth, and J. Wohlfart (2023). Designing information provision experiments. *Journal of economic literature 61*(1), 3–40.

Hardy, M. and J. McCasland (2023). Are small firms labor constrained? experimental evidence from ghana. *American Economic Journal: Applied Economics 15*(2), 253–284.

Houseman, S. N. (2001). Why employers use flexible staffing arrangements: Evidence from an establishment survey. *Ilr Review 55*(1), 149–170.

Hsieh, C.-T. and B. A. Olken (2014). The missing "missing middle". *Journal of Economic Perspectives 28*(3), 89–108.

Jones, S. and R. Santos (2022). Can information correct optimistic wage expectations? evidence from mozambican job-seekers. *Journal of Development Economics 159*, 102987.

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics 129*(2), 597–652.

Katz, L. F. (1996). Wage subsidies for the disadvantaged.

Kelley, E. M., C. Ksoll, and J. Magruder (2022). How do online job portals affect employment and job search? evidence from india.

Kempis, M., T. Ogden, A. Heyer, and T. Cook (2023, May). Kenya country data overview: Data from the small firm diaries. Small Firm Diaries Firm Profile, Financial Diaries Research.

Kessler, J. B., C. Low, and C. D. Sullivan (2019). Incentivized resume rating: Eliciting employer preferences without deception. *American Economic Review 109*(11), 3713–3744.

KNBS (2019). 2019 kenya population and housing census: Volume i. Technical report, Kenya National Bureau of Statistics (KNBS).

KNBS (2022). Gross county product 2021 (gcp). Technical report, Kenya National Bureau of Statistics (KNBS).

Kuddo, A. and A. Kuddo (2018). *Labor regulations throughout the world: an overview*. World Bank Washington, DC.

Macchi, E. (2023). Worth your weight: Experimental evidence on the benefits of obesity in low-income countries. *American Economic Review 113*(9), 2287–2322.

Mas, A. and A. Pallais (2017). Valuing alternative work arrangements. *American Economic Review 107*(12), 3722–3759.

Meghir, C., R. Narita, and J.-M. Robin (2015). Wages and informality in developing countries. *American Economic Review 105*(4), 1509–1546.

Morduch, J. (1994). Poverty and vulnerability. *The American Economic Review 84*(2), 221–225.

Ndayikeza, M. A. (2023). Underemployment of college graduates: is doing anything better than doing nothing? In *Pacific Conference for Development Economics 2023*.

Neal, D. (1995). Industry-specific human capital: Evidence from displaced workers. *Journal of labor Economics 13*(4), 653–677.

Nickell, S. J. (1986). Dynamic models of labour demand. In *Handbook of Labor Economics*, Volume 1, pp. 473–522. Elsevier.

Parsons, D. O. (1972). Specific human capital: An application to quit rates and layoff rates. *Journal of political economy 80*(6), 1120–1143.

Rosenzweig, M. and H. Binswanger (1993). Wealth, weather risk and the composition and profitability of agricultural investments. *The Economic Journal 103*(416), 56–78.

Rosenzweig, M. R. (1988). Labor markets in low-income countries. *Handbook of development economics 1*, 713–762.

Ruggles, S., L. Cleveland, R. Lavaton, S. Sarkar, M. Sobek, D. Burk, D. Ehrlich, Q. Heimann, and J. Lee (2024). Integrated public use microdata series, international: Version 7.5 [dataset]. IPUMS-International, University of Minnesota.

Ulyssea, G. (2018). Firms, informality, and development: Theory and evidence from brazil. *American Economic Review 108*(8), 2015–2047.

Walker, M. W., N. Shah, E. Miguel, D. Egger, F. S. Soliman, and T. Graff (2024). Slack and economic development. Technical report, National Bureau of Economic Research.

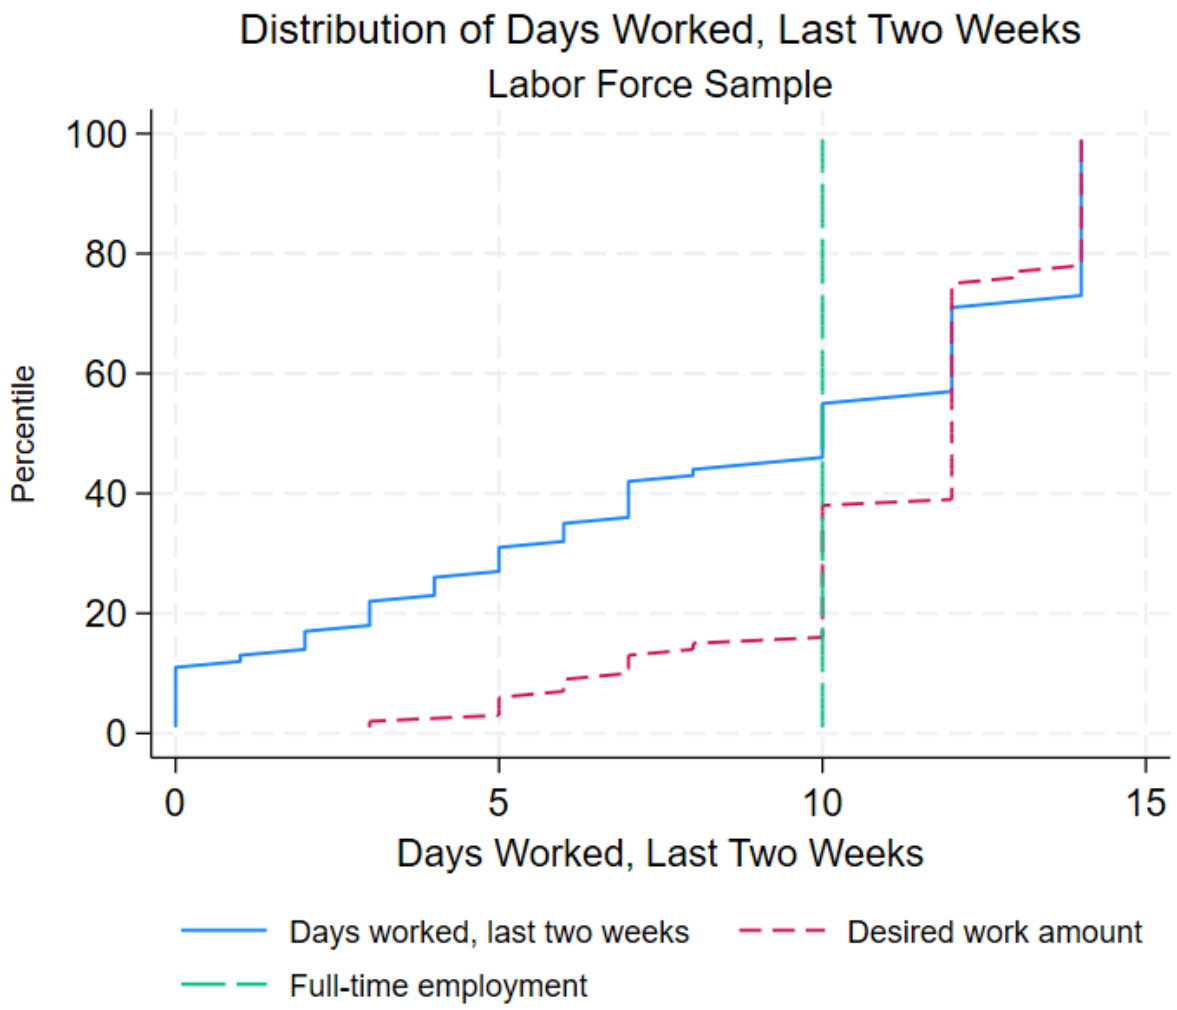
Wiedemann, V., P. Koschmieder, R. Santos, and M. Cervellati (2023). Spatial inequality and informality: Evidence from 22 developing countries. Technical report, CESifo Working Paper No. 10354.

Wiswall, M. and B. Zafar (2021). Human capital investments and expectations about career and family. *Journal of Political Economy 129*(5), 1361–1424.

World Bank (2013). *Enterprise Surveys: Kenya 2013 – Questionnaire Manual*. World Bank. Accessed: 2024-09-26.

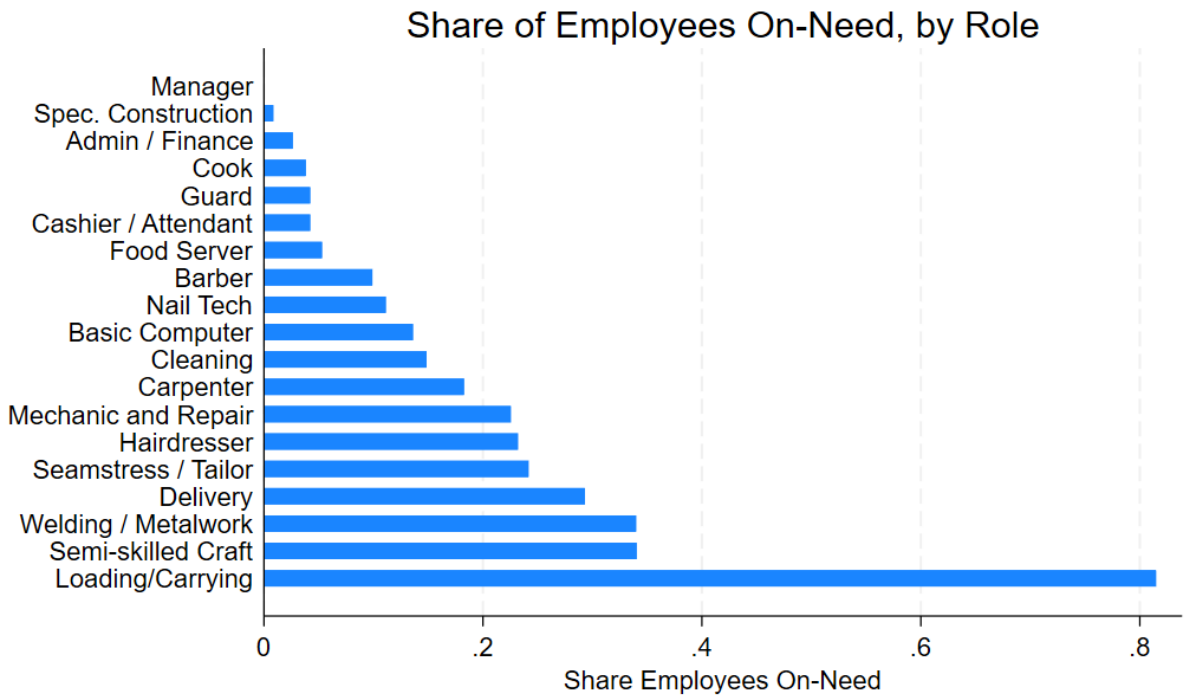
7 Figures

Figure 1: Distribution of Days Worked, Labor Force Sample



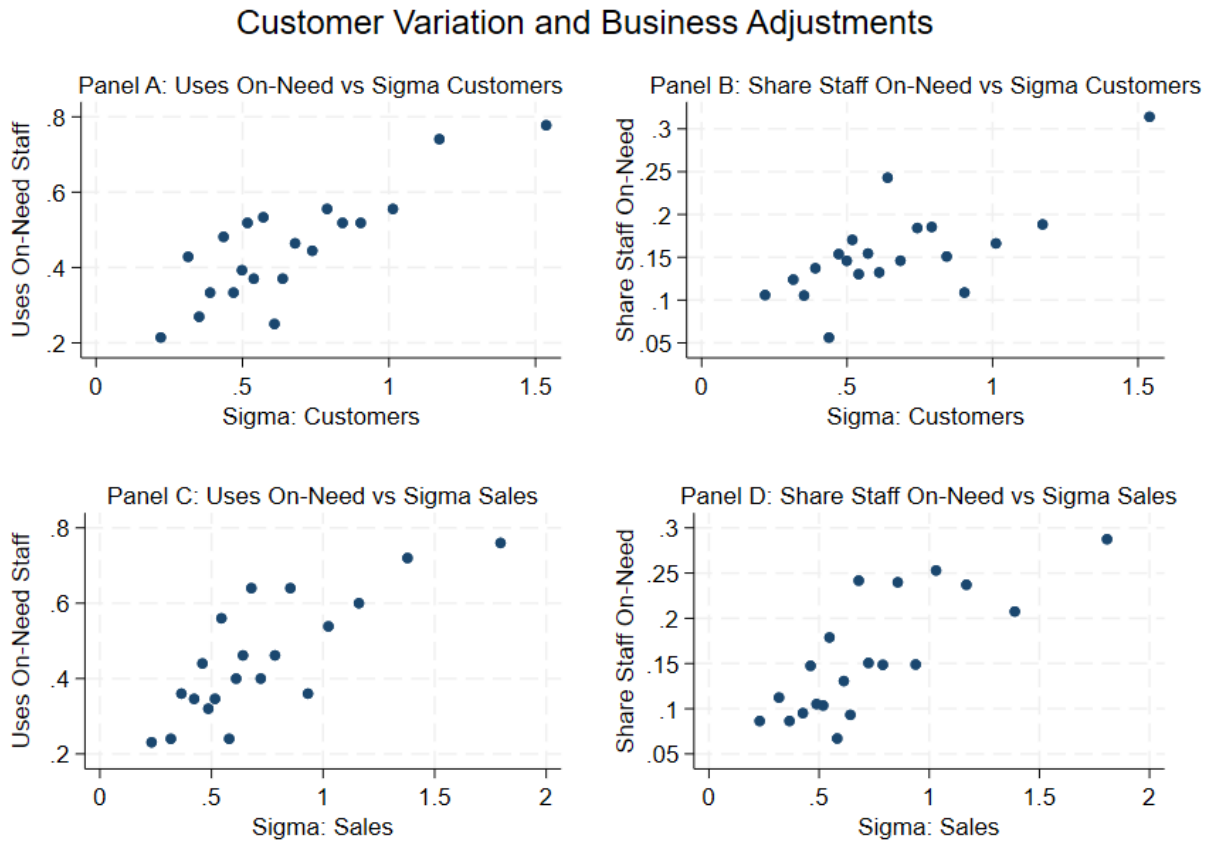
Notes: Data comes from our Labor Force Survey of 427 young adults in low-income neighborhoods. The blue line depicts the cumulative distribution function of workers' actual work in the last two weeks, and the red line, their ideal amount of work, if work was available and they could choose the number of days they worked.

Figure 2: Breakdown of Employment Type, By Role



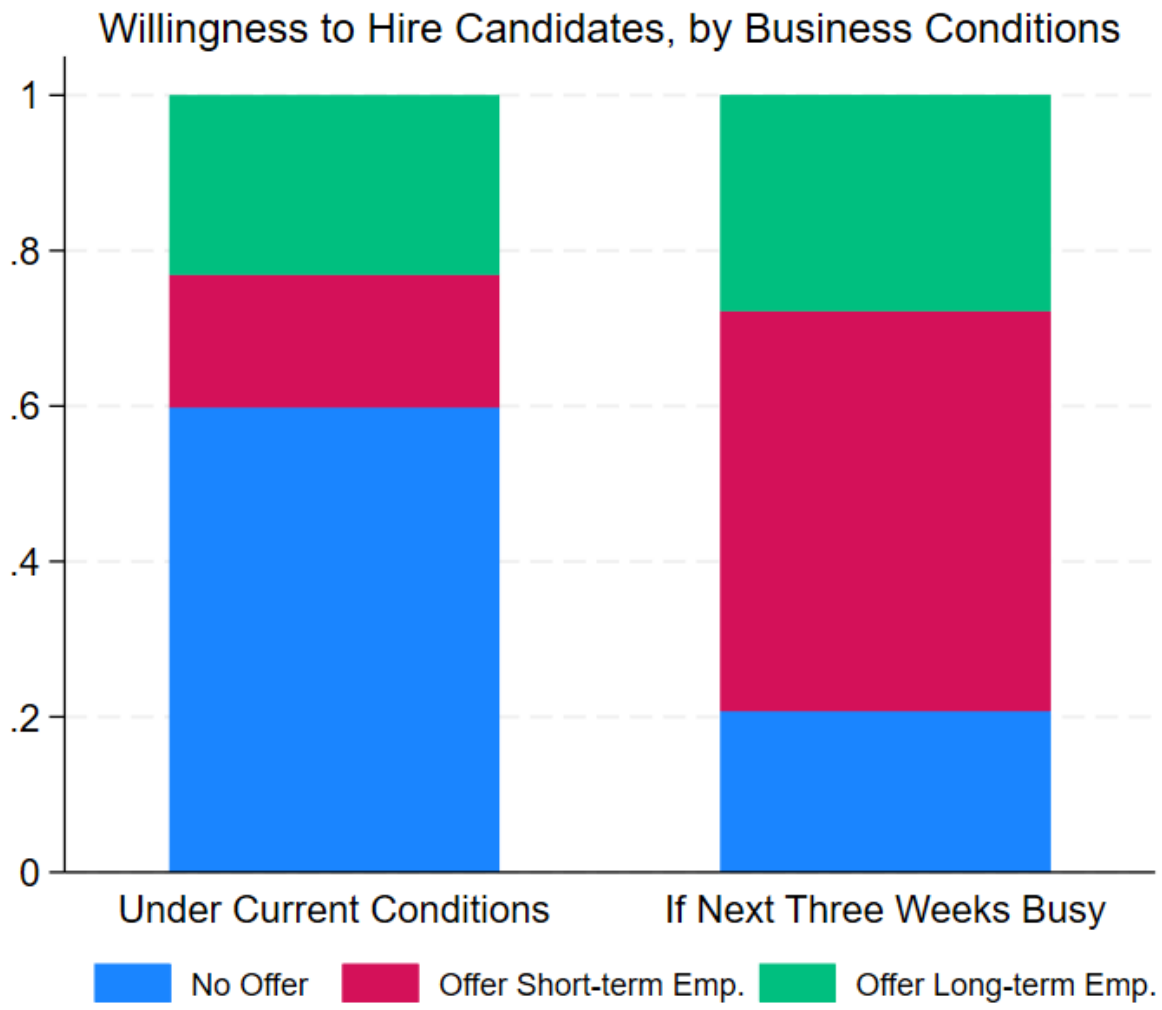
Notes: Data comes from our Firm Survey. Each bar shows the share of staff in each role that are employed on a short-term, on-need (rather than regular) basis, for all roles in which we observe 50 or more workers in our data. The number of employees in any given role-contract type-firm is winsorized at the 99th percentile.

Figure 3: Relationship between customer variation and firm staffing variation



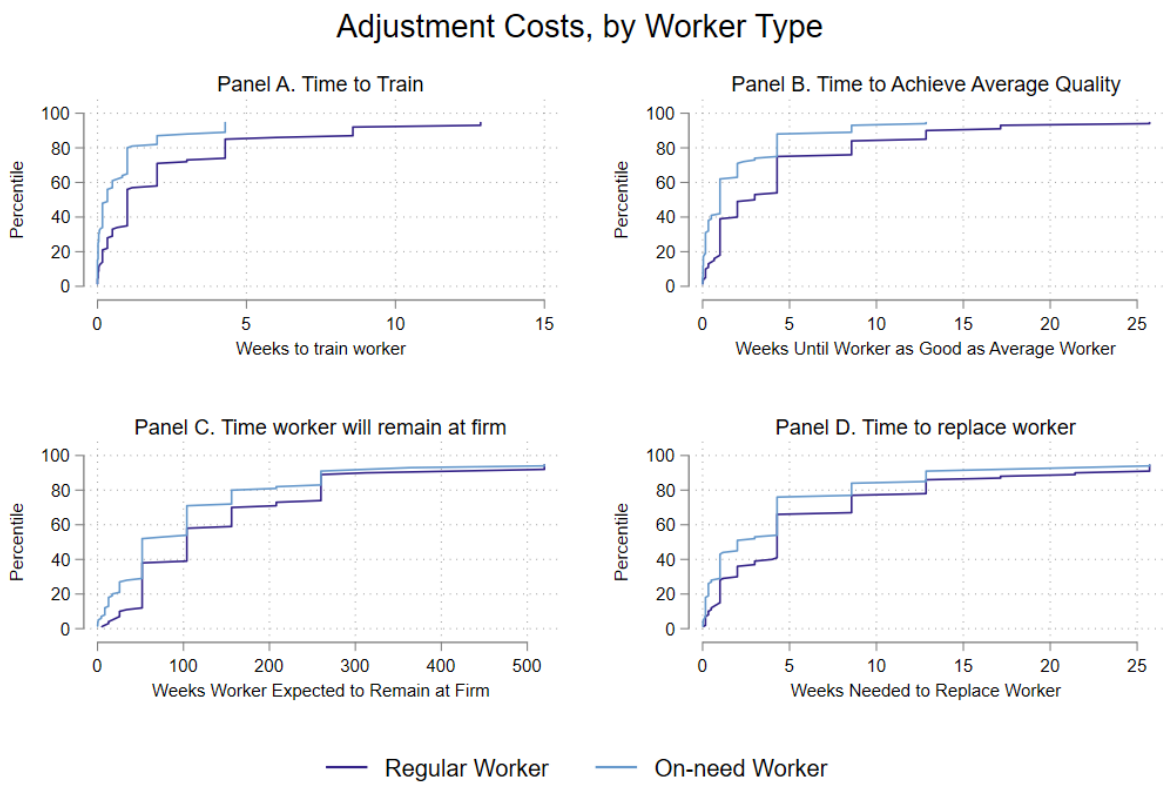
Notes: Data comes from our Firm Survey. Results are binned at the vigintile level. On the x-axis are our two measures of variation in staffing: $\hat{\sigma}$ of our estimated distributions in customers (top row) and sales (bottom row). We plot these against whether a firm uses on-need staff in response to business variation (left column) and the share of staff hired on-need (right column). These results are presented in regression form in Table 4 and Appendix Table A3.

Figure 4: Firm Responses to Demand Shocks: Hypothetical Hiring



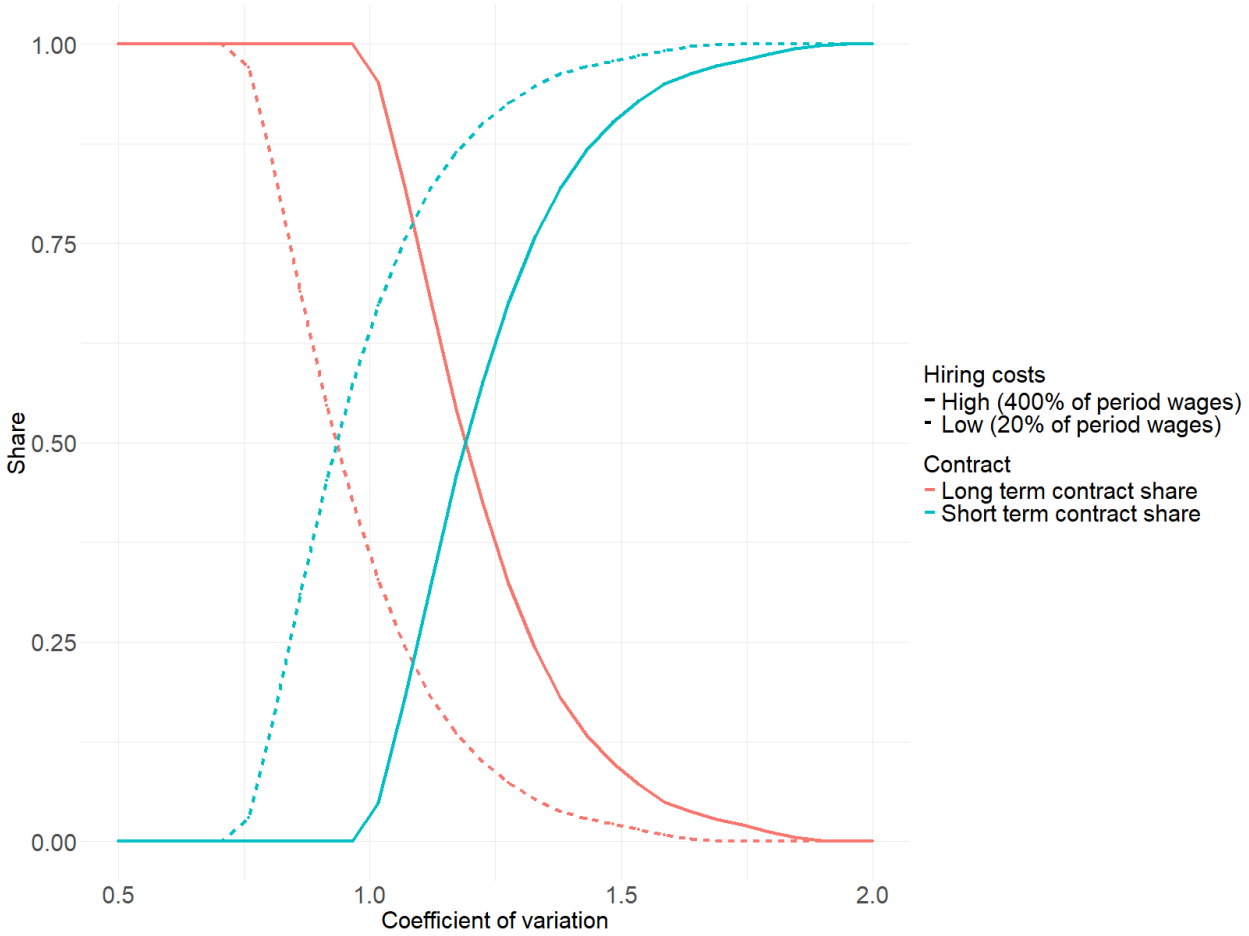
Notes: Data comes from our Firm Survey, in which we ask firms in our vignettes to consider whether they would hire a candidate under their current circumstances (left bar) and if the next three weeks were busy (right bar). Among firms that said they would hire the hypothetical job candidate, we asked whether they would hire them on a short-term, on-need basis, or as a regular worker. Results are presented in regression form in Appendix Table A4.

Figure 5: Adjustment Costs by Worker Type



Notes: Data comes from our Firm Survey, in which we ask firms about the costs associated with on-boarding and replacing their most recent hire(s): their most recent worker hired on a regular basis, and for the 59% of firms who hired workers on a short-term basis, for their most recent short-term hire. Each graph shows a CDF of the adjustment costs associated with these workers. Estimated differences between the types of hires are reported in Table 5.

Figure 6: Contract choice probability



Notes: The model is estimated for the following parameter combination: w = 5, $\theta_0 = 10$, $\beta = 0.9$, q = 0.05, $\mu = 10$, $\gamma = 0.1$, h = 0.9. The distribution of worker quality is assumed to be exponential with rate equal to 1.

8 Tables

Table 1: Summary Statistics, Firms

Panel A: Firm Characteristics		
	Mean	Median
Number of Employees	9.49	4
Employs 10 or More Workers	0.18	
Retail	0.44	
Hospitality	0.24	
Beauty	0.15	
Manufacturing	0.11	
Daily Customers: Typical Day	34	15
Daily Sales (KSH): Typical Day	46482	12000
Daily Sales (USD PPP): Typical Day	1078	278
Panel B: Contracting Arrangements		
	Mean	
Employs Regular Workers, Paid Monthly	0.55	
Employs Regular Workers, Paid Daily	0.61	
Employs Workers on Short-Term Basis	0.59	
Share of Staff Used on Short-Term Basis	0.16	
Most Recent Hire Receives Government Benefits	0.20	

Notes: Statistics here are from our sample of 601 firms. The mean values for number of employees, daily customers and daily sales are winsorized at the 99th percentile. We convert from KES to USD PPP using 43.12, the World Bank's conversion rate for 2023.

Table 2: Summary Statistics, Labor Survey

Panel A: Young Adult Characteristics		
	Mean	
Completed Secondary School	0.78	
Female	0.53	
Age	26	
Married	0.28	
Has Children	0.59	
Has a Certificate, Degree or Diploma	0.30	
Has Done any Work in the Last Two Weeks	0.89	
Panel B: Work Behavior		
	Mean	Mean, Worked = 1
Employed Full-Time, Last Two Weeks	0.54	0.61
Days Worked, Last Two Weeks	8.54	9.63
Desired Days Worked, Last Two Weeks	11.04	11.17
Desired Work Amount Is Full-Time	0.84	0.86
Work in the Last Two Weeks Involves Multiple Jobs	0.36	0.41
Work History Spans Multiple Sectors	0.74	0.79

Notes: Means reported here come from our Labor Force Survey of 427 young adults in Low-Income neighborhoods of Nairobi; the second column of Panel B shows the results for the sub-sample of workers who have done any work in the last two weeks.

Table 3: Comparison of Regular and On-Need Workers

	Regular Workers	On-Need Workers
Paid Daily	0.43	0.91
Paid Weekly	0.11	0.07
Paid Monthly	0.45	0.01
Daily Wage (USD PPP)	17.31	18.08
Days worked, last month	23.22	7.87
Worker has written contract	0.19	-
Offers government-mandated benefits	0.20	0.04
Last month was first time working with employee	-	0.29
Expects to Re-Use Employee	-	0.87
Expects to Promote Employee	-	0.31
Interviewed Individual Before Hiring Them	0.63	0.41
Required CV as part of application	0.31	0.14
Job is in Retail	0.42	0.35
Job is in Hospitality	0.26	0.13
Job is in Beauty	0.15	0.12
Job is in Manufacturing	0.11	0.13

Notes: This table compares regular workers (i.e., those expected to come in on a regular basis) against those brought in on a short-term, on-need basis. Rows 6, 7, 11 and 12 come from our module on hiring, in which we ask about the most recent regular and short-term hires. The rest come from our survey modules of regular workers (in which we record detailed information about up to four regular workers) and on-need workers (in which we collect information about all on-need workers used in the last month). A "-" indicates that the question was not asked about that group of workers. Daily wages are winsorized at the 1st and 99th percentiles.

Table 4: Relationship between Margins of Adjustment in Customers and Variation in Business Demand

	Panel A: Relationship between Margins of Adjustment in Customers and Variation in Business Demand			
	(1) Uses On-Need Staff in Response to Business Variation	(2) Varies the Number of Regular Staff Across Days in Response to Business Variation	(3) Varies the Salary Regular Staff Receive in Response to Business Variation	(4) Share of Staff Hired on On-Need Basis
$\hat{\sigma}$ - Customers	0.277***	0.126*	0.217***	0.073**
	(0.068)	(0.067)	(0.064)	(0.033)
Observations	541	541	541	530
Mean of Outcome	0.45	0.29	0.32	0.16
SD of Outcome	0.498	0.453	0.465	0.214
	Panel B: Total Staffing			
	(1) Coefficient of Variation (CV): Total Staff	(2) CV: On-Need Staff	(3) CV: Regular Employees, Paid Daily	(4) CV: Regular Employees, Paid Monthly
$\hat{\sigma}$ - Customers	0.146***	0.318***	0.013	0.037*
	(0.039)	(0.099)	(0.030)	(0.019)
Observations	531	541	532	540
Mean of Outcome	0.17	0.56	0.06	0.03
SD of Outcome	0.26	0.69	0.17	0.10
χ^2: Coefficient = Column (2)			9.5	8.6
p-value: Coefficient = Column (2)			0.002	0.003
	Panel C: Total Wage Bill			
	(1) CV: Total Staff Wage Bill	(2) CV: On-Need Staff Wage Bill	(3) CV: Regular Employees, Paid Daily Wage Bill	(4) CV: Regular Employees, Paid Monthly Wage Bill
$\hat{\sigma}$ - Customers	0.239***	0.323***	0.097**	0.078***
	(0.049)	(0.099)	(0.041)	(0.029)
Observations	457	537	491	509
Mean of Outcome	0.28	0.55	0.12	0.06
SD of Outcome	0.35	0.69	0.25	0.19
χ^2: Coefficient = Column (2)			5.1	6.4
p-value: Coefficient = Column (2)			0.024	0.012

Notes: Each column is from a regression of firms' use of variation transmission strategies (Panel A), and coefficients of variation of their staffing levels (Panel B) and wage bills (Panel C) on the estimated $\hat{\sigma}$ of the firm's distribution of sales. All regressions include sector fixed effects, neighborhood fixed effects, and controls for the total number of (regular) employees. Standard errors are Huber-White heteroskedastic. The differences in sample sizes reflect don't know(s) or refusal responses from respondents; aggregates are only calculated for cases with all non-missing values. The χ^2 tests reported in columns 3 and 4 are from a seemingly unrelated regression test of the null hypothesis that the coefficients in those columns are equal to the coefficients in column 2, i.e., that firms adjust their regular daily and monthly staff (in terms of staffing and wages, respectively) to the same degree as they do their on-need staff. We present the same results with $\hat{\sigma}$ in terms of sales as our independent variable in Appendix Table A3.

Table 5: Comparison of Adjustment and On-Boarding Costs, by Worker Type

	(1) Weeks to Train Worker	(2) Weeks Until Worker Was as Good as Average Worker	(3) Weeks Worker is Expected to Remain at Firm	(4) Weeks Needed to Replace Worker
Long-Term Worker	2.293***	2.734***	46.749***	2.294***
	(0.434)	(0.554)	(7.956)	(0.548)
Observations	923	933	873	938
Mean of Outcome	2.61	4.93	140.04	7.12
R-squared	0.054	0.070	0.066	0.069
SD of Outcome	6.71	9.28	145.03	10.86

Notes: Each column is from a regression comparing the training and adjustment costs of hires at a firm on an indicator for whether the hire was for a long-term, regular (rather than short-term) role. All regressions include sector fixed effects, neighborhood fixed effects, and controls for the total number of (regular) employees. There are up to two observations per firm. Standard errors are clustered at the firm level. The differences in sample sizes reflect don't know or refusal responses from respondents.

Table 6: Perceived Behavioral Response to Contract

	(1) Absenteeism perception	(2) Likelihood of quitting	(3) Likelihood of being fired	(4) Rating
Regular, Paid Daily	-0.368***	0.111	0.150*	-0.015
	(0.110)	(0.098)	(0.083)	(0.047)
On-Need	0.730***	0.972***	0.314***	-0.223***
	(0.154)	(0.112)	(0.091)	(0.047)
Observations	3565	3567	3421	3561
Mean of Outcome	2.67	4.14	4.05	4.89
R-Squared	0.580	0.578	0.715	0.735
SD of Outcome	2.97	2.43	2.38	1.39
F-Test (Daily = On-Need)	54.459	61.229	3.914	17.864
P-Value (Daily = On-Need)	0.00	0.00	0.05	0.00

Notes: Each column is from a regression comparing firms' perceptions of how hypothetical candidates would perform if hired as a regular employee paid monthly (omitted category), a regular employee paid daily, and on on-need basis. Firm-profile fixed effects (e.g., firm 101's rating of candidate 3) are included. Outcome (1) corresponds to the number of days the respondent perceived the worker would miss in a month if hired under each of the three arrangements. Outcome (2) and (3) correspond to a scale from 1 to 10 where 1 is a extremely low likelihood and 10 is certainty. Outcome (4) corresponds to a scale from 1 to 7 where 1 corresponds to "Very poorly" and 7 corresponds to "Very well." Standard errors are clustered at the firm level.

Supplemental Appendix

A Sampling

A.1 Firm Sampling

Appendix Figures A1, and A2, show the Enumeration Areas (EAs) included in our study. Appendix Figure A1 shows all EAs, while Appendix Figures A2a and A2b show zoomed-in examples of how our EA construction varied by neighborhood type. Appendix Figure A2a shows the Central Business District (CBD)—here, each city block forms an EA. Appendix Figure A2b shows Kilimani, a less dense area, in which approximately 0.5 kilometers of a major road or a shopping center constitutes an EA.

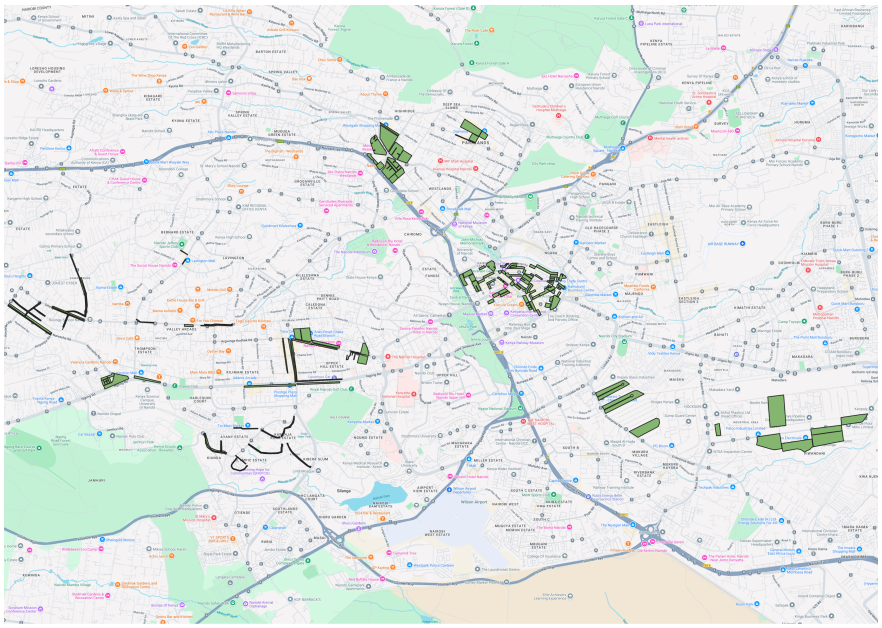
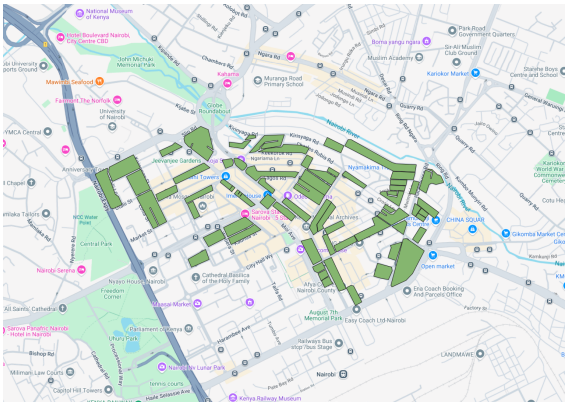
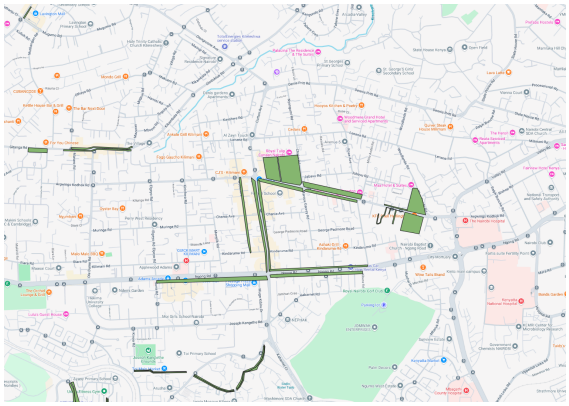


Figure A1: Enumeration Areas, Full Sample



(a) EAs - Central Business District



(b) EAs - Kilimani

Figure A2: Enumeration Areas by Neighborhood

A.2 Labor Force Survey Sampling Procedure

Our Labor Force Survey comprises 427 young adults residing in two low-income neighborhoods, Kibera and Kawangware. Both neighborhoods are in Western Nairobi and are within commuting distance of all of the neighborhoods in our firm sample. This survey was conducted in August 2024, over 15 days (12 of these days were weekdays; 3 were done on Saturdays). We are unable to reject the null that those surveyed on weekdays and weekends have the same characteristics; means for labor force participation if anything are slightly higher for those survey. (We find that 89% of those surveyed on weekdays worked in the last two weeks, and have worked an average of 8.7 days in the last week. Of those surveyed on weekends, 86% have worked in the last two weeks, and have worked an average of 8.0 days).

Each day, enumerators began at a distinct central point within the neighborhood (for example, outside the Town Centre office). They began walking in a randomly selected direction, were told to speak with every third person they encountered. Respondents were eligible for inclusion in the survey if they (a) were aged 18-35, (b) were not a full-time student, and (c) consented to participate in the study.

A.3 Firm and Labor Survey Survey Overlap

We estimate that 78% of wage work episodes observed in our LFS sample took place in neighborhoods also in our firm survey, and 88% in sectors we survey, reflecting partial but incomplete overlap. One important difference is that a substantial share of work episodes completed in the LFS reflect work done by individuals their neighborhoods of residence, rather than across all of the neighborhoods in which we survey. In our firm sample, 31% of firms are in these low-income neighborhoods, as compared to 71% of work episodes in our Labor Force Survey.[29]. A second difference is that the workers in our LFS have a greater share of employment done on an on-need (rather than regular basis). In our firm sample, 18% of matches are on an on-need basis, and 31% among firms that use any short-term labor. In our LFS sample, among the work episodes that involve working for a wage (rather than self-employment), 57% of work episodes are on an on-need basis.

While there are differences between our samples, we believe that the Labor Force Survey still adds meaningful value, for two reasons. First, collecting detailed information from workers for a (partially) overlapping sample allows us to collect information from workers not necessarily known by firms, including workers' full employment history, engagement with other jobs (including contemporaneously), and expectations about the future. Second, the Labor Force Survey lets us document the very high prevalence of short-term, high-turnover contracts for a specific, defined population: young adults in low-income neighborhoods.

[29] Of the workers in our LFS sample who leave their home neighborhood for work, roughly 1/3 work in another neighborhood we sampled, and 2/3 in one we did not

B Vignettes

B.1 Sample Vignette

The following is an example of a vignette that firms saw, in a base case:

Now let's look at Profile Number 2.

You are considering them for the role of:
Security Guard

Their details are:

Name:
Evans

Age:
24

How did you meet?
Referred to you by a friend

What is his experience in the last five years?
Guard - Mlinzi Mkuu security company (2019-2020)
Guard - Simba Security Ltd (2022-2023)

Does he have a family?
Yes, married with two children

Does he have all the necessary documents?
Yes

Does he have any references?
Yes, from their last two employers

Where does he stay?
Born and stays in Kibera

We randomized all elements, except for whether the individual had the necessary documents (for which the answer was always yes).

B.2 Evaluation of Vignettes

In this subsection, we consider two questions related to our use of vignettes to uncover firm beliefs and preferences: (1) whether respondents understood our questions and took them seriously, and (2) whether the results are likely to be biased by experimenter demand effects.

On the first count, we believe our results suggest that respondents understood the questions, and took them seriously, for three reasons.

First, we did extensive piloting with respondents prior to the data collection, and eliminated measures where we received qualitative information that the questions were not well understood or easily answered.[30] Second, we observe meaningful, statistically significant patterns in the hypotheses we test (i.e. how firms respond to demand shocks, and how contract type affects worker performance). In a scenario where respondents were confused or just responding randomly, we might expect our results to be attenuated towards zero; the fact we detect results points against the hypothesis we are simply capturing noise. Third, we can also test whether firms respond sensibly on margins that are not the core focus of this paper. In Appendix Table A6, we show how firm rating and willingness to consider a candidate varies with their work experience, and in Appendix Table A7, how their reported willingness-to-use short-term contracts covaries with their actual use of short-term contracts. Reassuringly, we observe both that firms value experience (and especially same-sector experience), and that firms who depend to a greater degree on short-term labor are more likely to report that the contract they would offer a hypothetical worker would be on a short-term basis.

Moreover, other research that relies on the collection of hypothetical beliefs and preferences (without incentivization) has found that these results correlate reasonably. For example, work exploring New York University student beliefs and preferences about different college majors has been found to correlate positively with students' actual behavior years later (Wiswall and Zafar, 2021).

Our evidence with respect to experimenter demand effects depends on secondary references (i.e., we do not explicitly test for the presence of demand effects). Research in development economics has broadly found that demand effects are generally relatively small, with tight bounds (De Quidt et al., 2018). Work on information experiments, primarily in higher-income countries, finds that demand effects are most likely to be present in contexts where respondents believe there is a "right answer" the research team wants to hear (Haaland et al., 2023). In our case, it is not clear such a phenomenon is likely to be present—it does not seem clear to us that respondents would report that (a) demand shocks lead to increases in specifically short-term hiring, or (b) on-need workers do a worse job because this is a result the research team would want. That being said, as the use of vignettes and "randomized resume"-type approaches grow, we believe that there is value in methodological work that continues to evaluate the presence or absence of demand effects.

[30] Two examples we discarded were: (1) showing firms a path of past sales over time, and observing how this affected subsequent hiring, and (2) predicting the reservation wage of hypothetical workers.

C Demand Variation

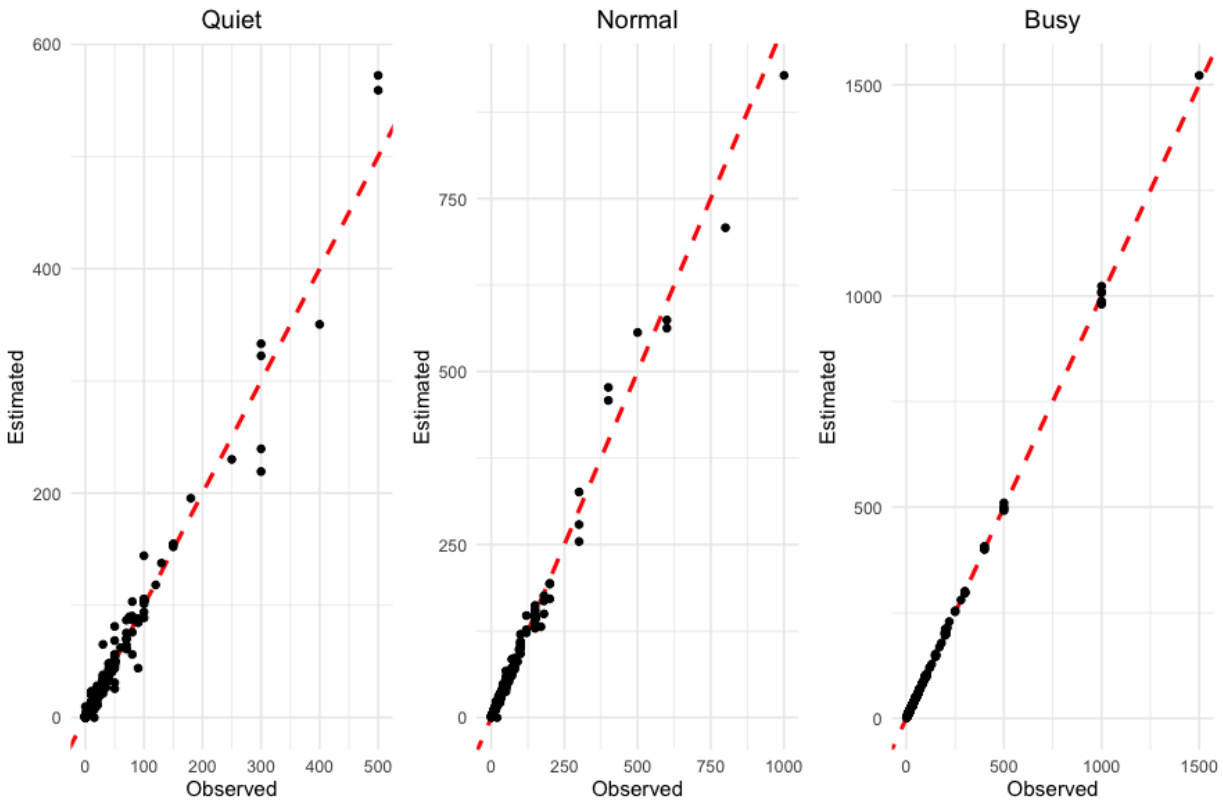
C.1 Distribution Estimation Procedure

For each firm we obtain the share of days that are quiet, normal, and busy, as well as the corresponding volume of sales and costumers on those days. We convert the shares into percentiles, P, and represent the volume of sales and costumers as the associated quantiles, Q(P).

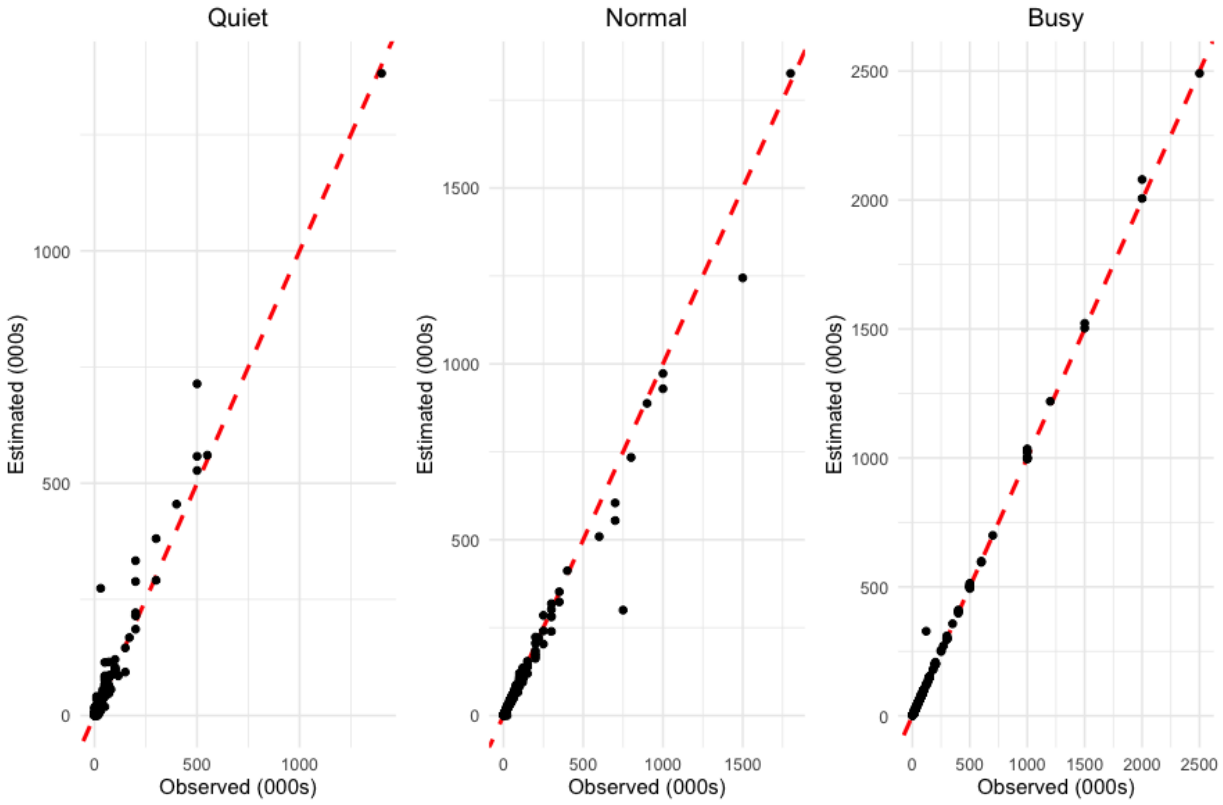
For a given distribution, F, parameterized by θ, we recover θ through a "method of moments" procedure. Formally, we find θ such that:

$$\theta := \arg\min \left(Q(P) - F^{-1}(P;\theta)\right)'\left(Q(P) - F^{-1}(P;\theta)\right) \tag{12}$$

Figure A3 visualizes the fit of the distribution by comparing estimated quantiles, $F^{-1}(P;\hat{\theta})$ on the x-axis, with actual quantiles on the y-axis. The dashed red-line corresponds to the 45% line.



(a) Customers



(b) Sales

Figure A3: Fit of log-normal distribution for customers and sales

C.2 Comparison of different distributions

To identify which distribution matches the data best, we begin by fitting five distributions: Poisson, Exponential, Negative Binomial, Truncated Normal and Log-normal. We use all of the three quantiles to fit the distributions. We then compute residuals by calculating the distance between the estimated quantiles and the true quantiles. Table A1 summarises how well each distribution fits our data using two measures of fit: (1) mean of squared residuals (MSE), corresponding to the mean of squared residual, and (2) mean percent deviation (MD), corresponding to the mean absolute value of the residuals, scaled by the true quantile. The results highlight that both the truncated normal and the log-normal distribution fit the data considerably better than other distributions.

Distribution	MSE (Customers)	MD (Customers)	MSE (Sales)	MD (Sales)
Log normal	50.22	0.09	280772776.76	0.14
Truncated normal	68.46	0.09	201467192.99	0.13
Exponential	1822.91	0.38	9720974441.81	0.47
Poisson	2232.94	0.46	9220055943.76	1.12
Negative Binomial	1128.27	0.18	9141196000.11	1.09

Table A1: Comparison of out of sample fit

C.3 Distribution of $\hat{\sigma}$

Appendix figure A4 plots the distribution of $\hat{\sigma}$ by sector and neighborhood. It shows that while there exist differences across sector (for example, that manufacturing firms appear to have a greater dispersion of customers) that much of the differences in variation we are capturing reflect within-sector and within-neighborhood variation. It is not simply the case that certain sectors or neighborhoods are responsible for the variation we observe.

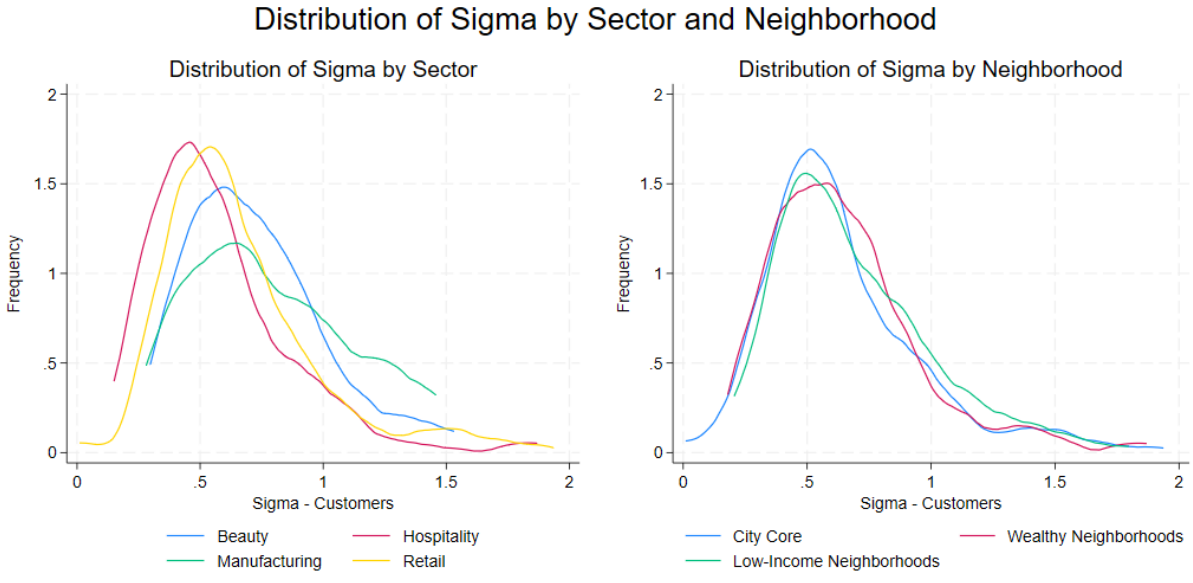


Figure A4: Distribution of $\hat{\sigma}$ - Customers

C.4 Testing the Predictability of Variation

In this subsection, we consider the evidence regarding the degree to which the demand variation we observe represents predictable versus unpredictable variation.

First, we have two pieces of direct evidence: the degree to which firms are able to ensure they have the right number of employees, and whether workers know in advance whether they will be working at a firm. In both cases, we see evidence consistent with the hypothesis that demand variation is not fully predictable.

We ask firms about whether on busy, normal, and quiet days whether they have "too many," "too few," or the "right number" of workers. We find that 23% of firms report that on busy days they have too few workers. While this is consistent with firms being able to predict quickly, but simply being unable to adjust, it is also consistent with the idea that they are not correctly forecasting their busyness levels.

Similarly, on the worker side, we measure their beliefs about future employment. Appendix Figure A5 shows results from our Labor Force Survey, in which we measure jobseekers' confidence regarding how likely their most recent work is to continue. For each job one of our survey respondents worked in the last two weeks, we ask them to report how likely it is that they will do any work in the same job in the next three months. This figure shows respondent beliefs, separated by whether the job was as a regular employee, or whether they were only called in when needed. This pattern is thus consistent with demand being uncertain, although it is also consistent with the variation being known to the firm, but that firms do not communicate these needs to their workers.

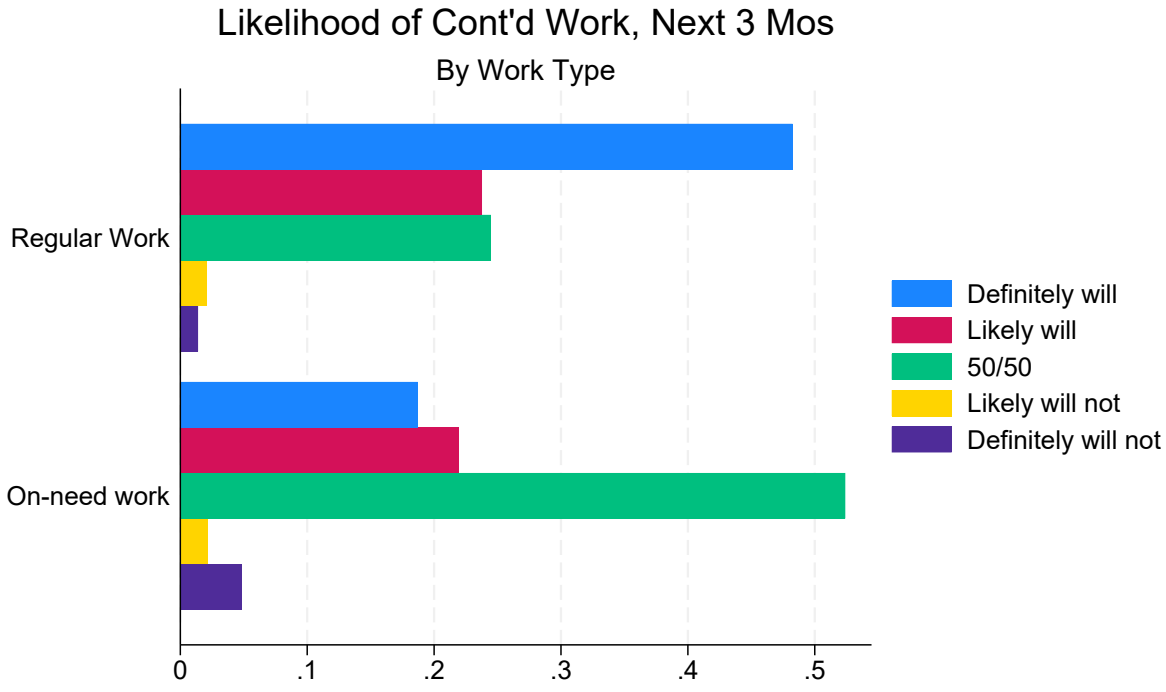


Figure A5: Expectation of Future Employment by Job Type

We did not directly ask firms about whether the demand variation they faced was predictable or unpredictable. However, we do have the ability to examine the degree to which recent variation for businesses represents common (by sector or neighborhood) or business-specific variation. Our hypothesis is that the more demand variation is correlated by sector or neighborhood, the more likely it is that this variation is ex ante predictable (for example, that restaurants are busy on weekends).

We measure whether each day in the last week was normal, busy, or quiet, and whether any on-need staff were brought in on each of those days. We use this to calculate the intraclass correlation (ICC) of busyness at the sector-day, neighborhood-day, and sector-neighborhood-day levels.[31] We calculate an ICC of 0.07 at the sector-day-level, 0.08 at the neighborhood-day-level, and 0.10 at the sector-neighborhood-day level. These results are visualized by sector-day and neighborhood-day in Appendix Figures ?? and A7 below. We interpret these results as suggesting a modest degree of common variation in the past week, but that a substantial share is specific to the firm.

Note that these results are also compatible with a hypothesis in which demand variation is business-specific, but predictable for a given firm. We view additional work that examines the predictability of demand, for example via a high-frequency panel in which firms are asked to forecast their future demand and staffing before reporting the actual realization in subsequent waves, as a promising contribution to this research agenda.

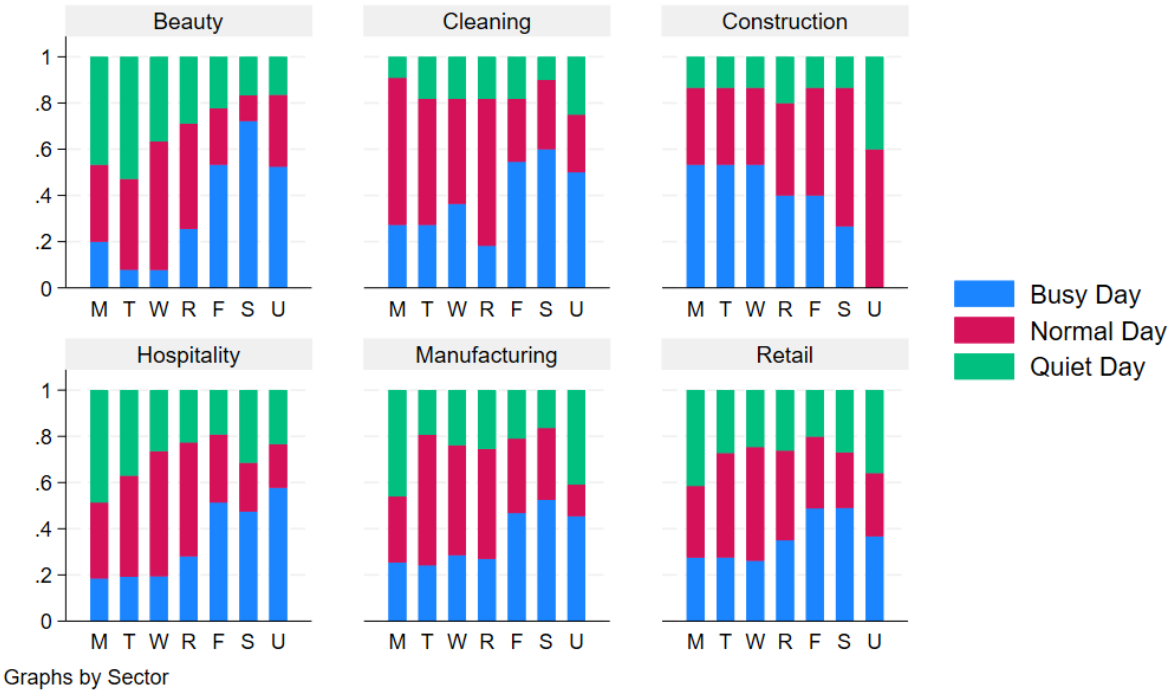


Figure A6: Busyness Last Week, By Sector and Day

[31] For example, we treat "retail, Monday" as a group when calculating the ICC at the sector-day-level. We code busy days as equal to 1, normal as 0.5, and quiet as 0.

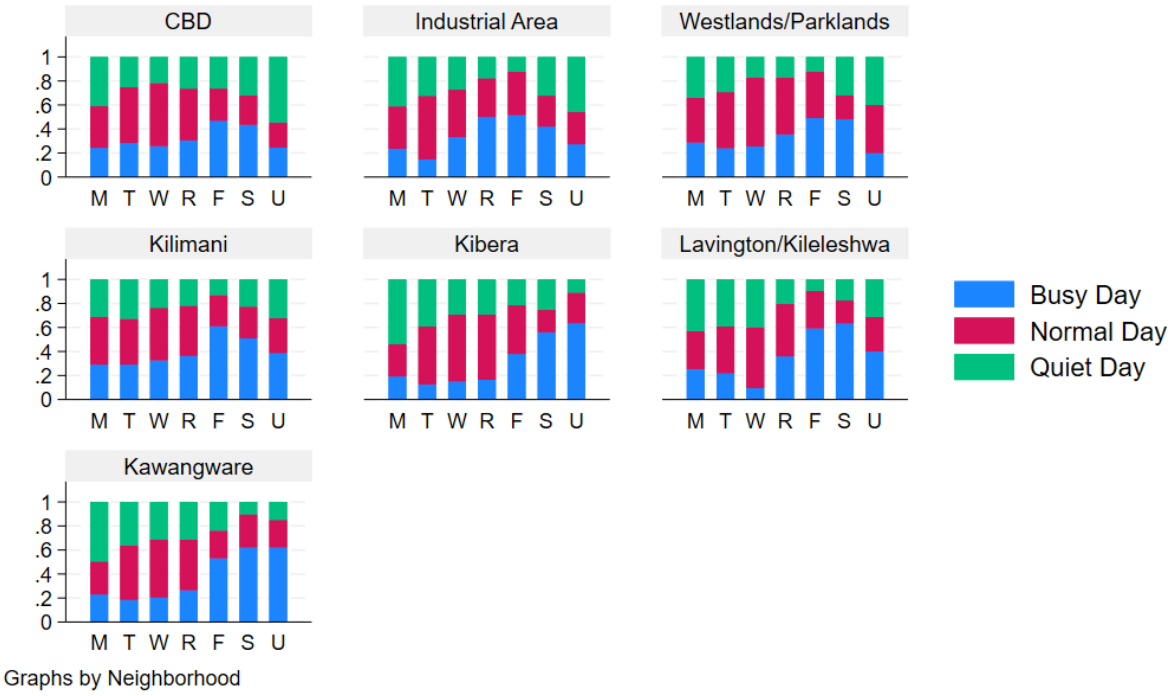


Figure A7: Busyness Last Week, By Neighborhood and Day

These patterns collectively suggest a very specific pathway through which this demand variation could be fully known. If demand variation is fully known, it would need to be the case that (a) some share of firms are still not able to adjust to it, (b) it is not communicated to workers, and (c) it is to a meaningful degree known but idiosyncratic to the firm, rather than predictable by sector or neighborhood characteristics. While we cannot rule this specific pathway out, our interpretation of these patterns is that a nontrivial amount of variation is ex ante unknown to the firm.

It is important to note that our hypothesis that demand variation affects firms' contracting decisions is not inherently contingent on whether this variation is known or unknown ex ante. Both known and unknown variation imply that the optimal level of staffing will vary between busy, normal, and quiet periods. If firms pass this variation onto their workers, both types of variation imply that short-term contracts will be an optimal response to certain levels of demand variation.

However, there are at least two ways in which the distinction between known and unknown demand matters. First, the predictability of demand variation matters for the welfare of risk averse workers. If businesses can accurately predict their future demand, use this information to determine their staffing, and communicate their plans to workers, workers will have less uncertainty about their income, and those with slack labor may be able to fill these gaps with other work. Second, this distinction matters for whether firms are staffing at the level that maximizes profits. In our model, firms are risk-neutral. However, if firm owners are instead risk averse, or face the possibility of exit (in a way that generates disutility for firm owners), a utility-maximizing owner will commit to staffing below the profit-maximizing level. Unpredictable demand variation might therefore affect

the optimal investment of risk-averse managers in their total staffing in much the same way that risk has been shown to shape the optimal investment behavior of farmers (Rosenzweig and Binswanger, 1993; Karlan et al., 2014).

D Regression Robustness Figures

The following figures show the robustness of our regression estimates to the inclusion of various controls. In each case, the version in bold font is the version included in our Main Tables.

Figure A8: Robustness Check, Table 4, Panel A

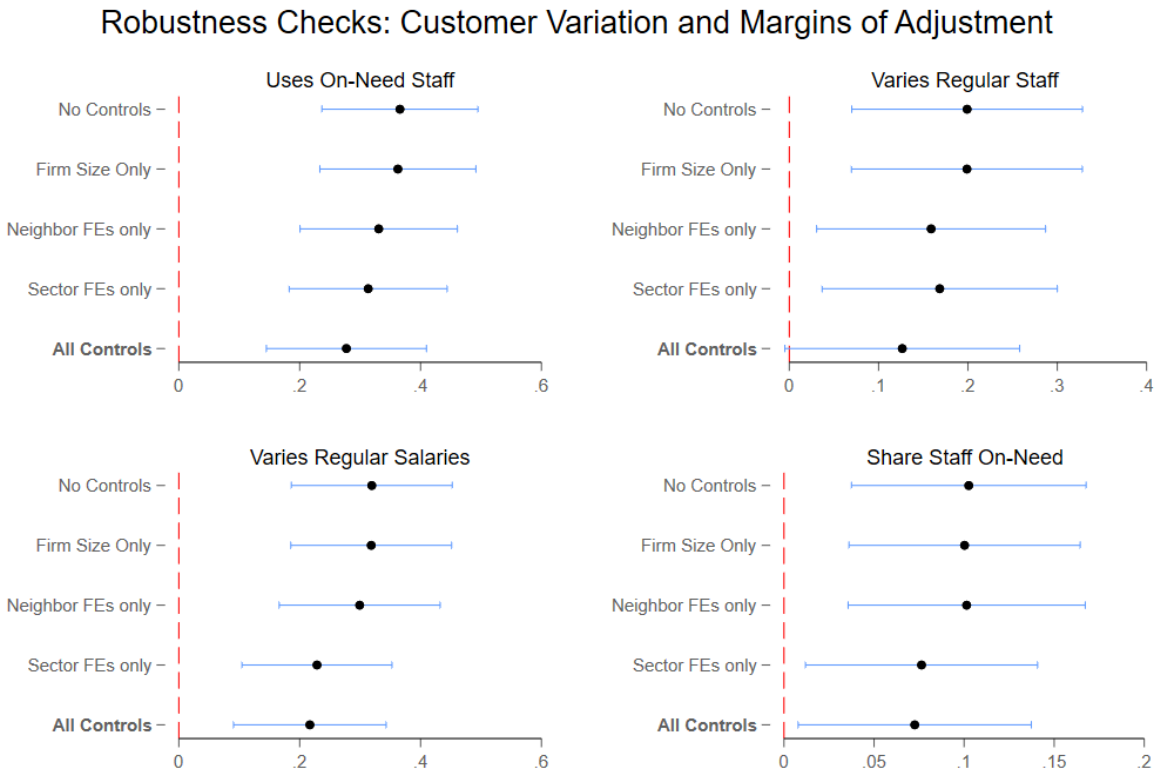


Figure A9: Robustness Check, Table 4, Panel B



Figure A10: Robustness Check, Table 4, Panel C

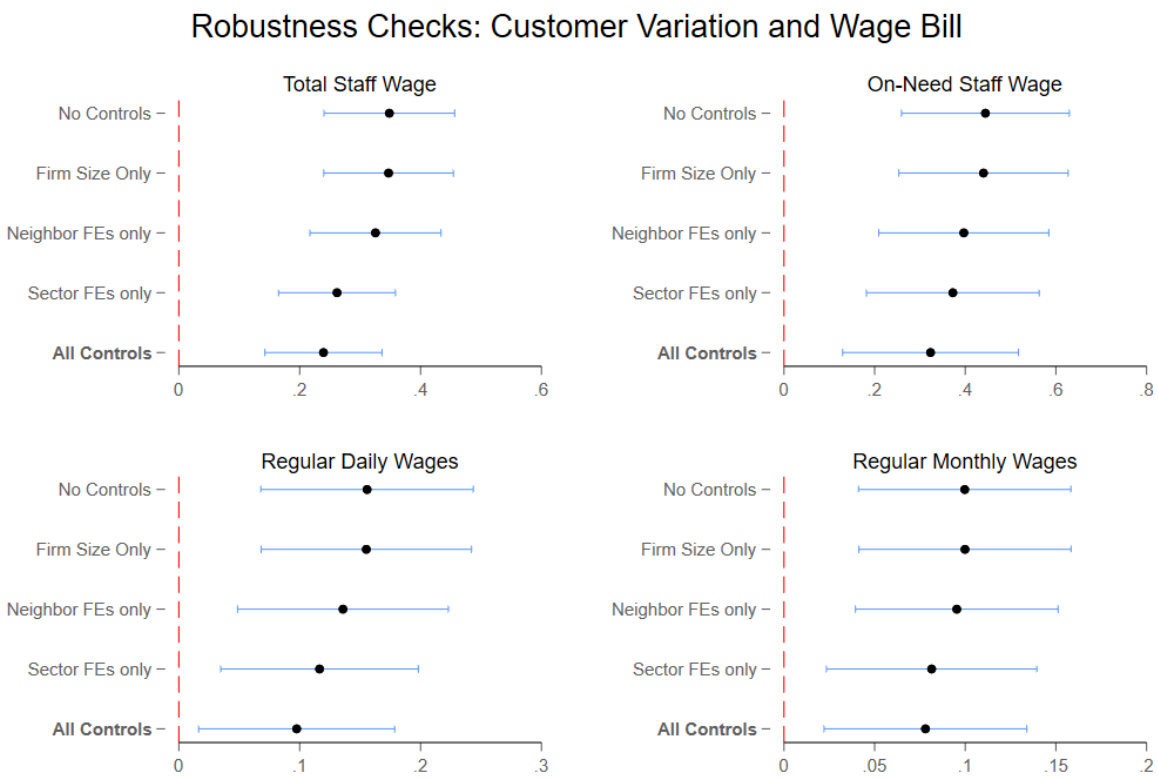


Figure A11: Robustness Check, Table 5

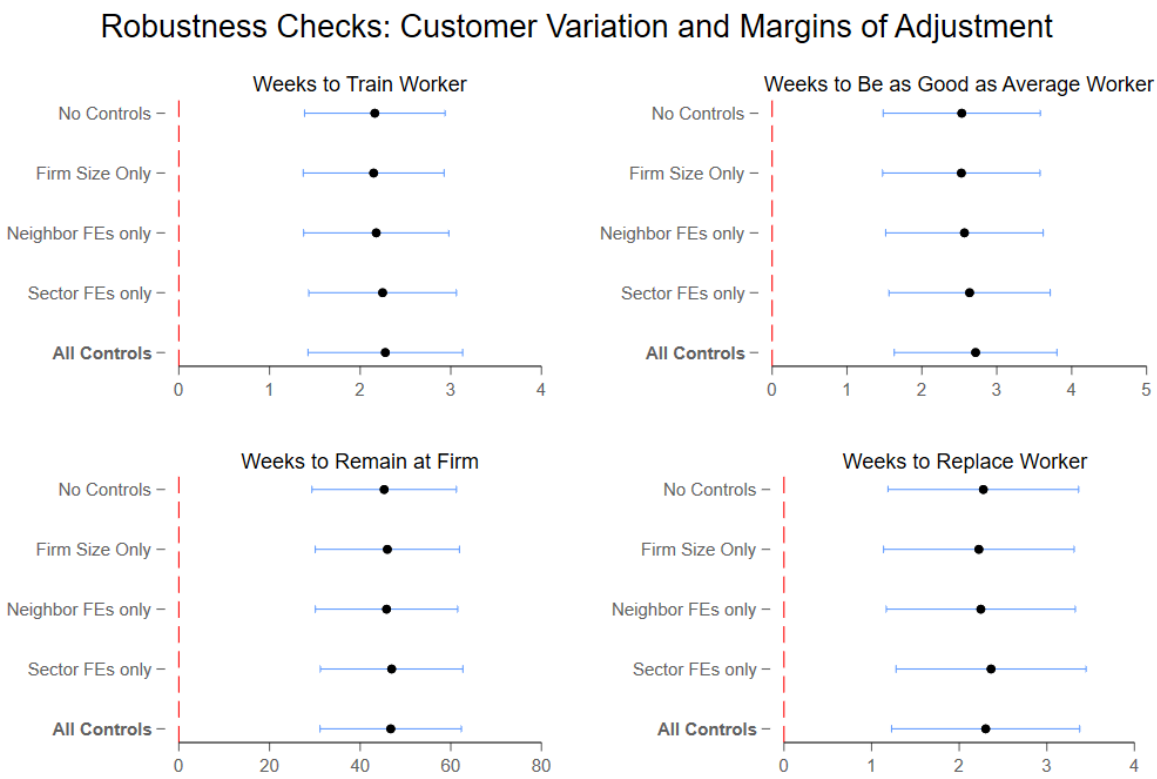


Figure A12: Robustness Check, Table 6, On-Need Workers

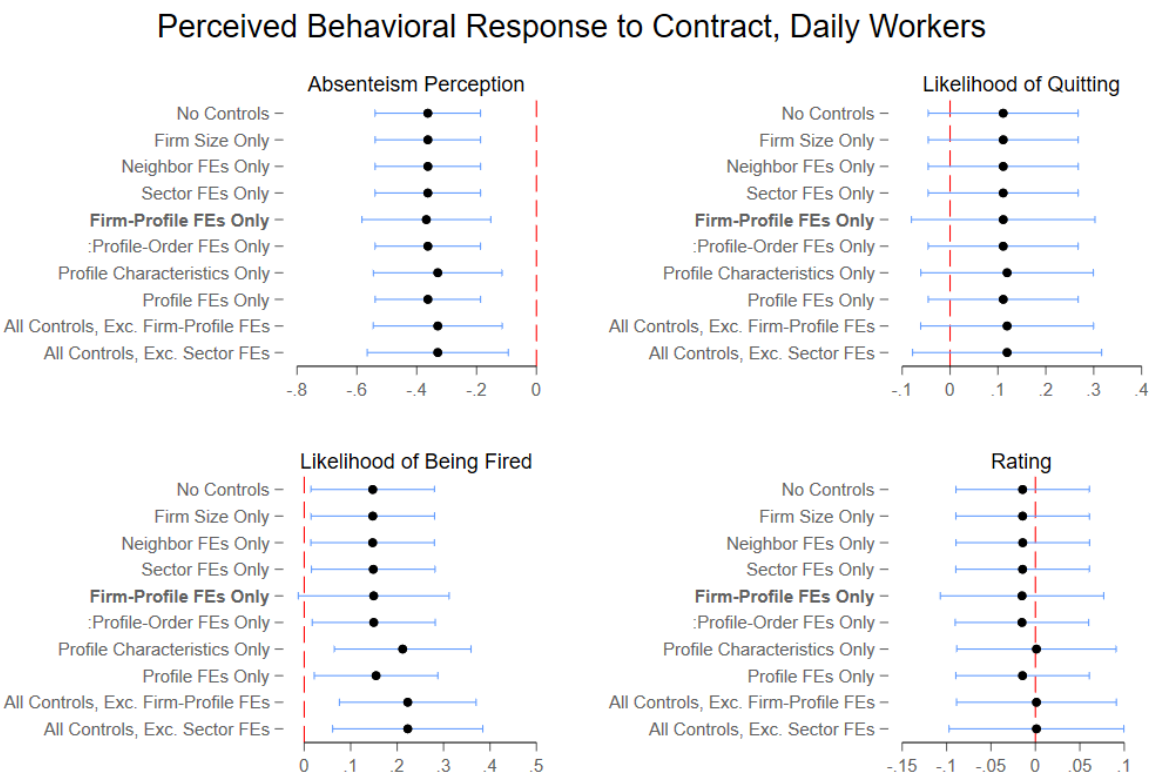


Figure A13: Robustness Check, Table 6, Regular Workers Paid Daily



E Decomposing Adjustment Costs

In Section 3.2, we show that hires involve higher adjustment costs when workers are brought in on a long-term basis. We show here in Appendix Table A2 that these results (in most cases) reflect both between- and within-role variation, and that (in all cases) a greater share of the differences are attributable to within-role variation.

First, we re-report our results from Table 5 in Panel A. Then, we re-estimate the results in Panel B, this time including role-level fixed effects (e.g., indicator variables for being a carpenter, manager, food server, etc.). In Panel B, we report the F-statistic and p-value from the test that the role coefficients are jointly equal to 0. We universally are able to reject this null hypothesis.

In Panel C, we calculate the percent reduction of the coefficient in Panel A once we add the role fixed effects. For example, we find that the coefficient on time to train is 80% in Panel B (with the inclusion of role fixed effects) as the initial version in Panel A, suggesting that a substantial share of differences between long-term and short-term staff reflect differences in investment in the same role.

Table A2: Decomposing Between- Versus Within-Role Variation in Adjustment Costs

	Panel A: Raw Comparison			
	(1)	(2)	(3)	(4)
	Weeks to Train Worker	Weeks Until Worker Was as Good as Average Worker	Weeks Worker is Expected to Remain at Firm	Weeks Needed to Replace Worker
Long-Term Worker	2.366***	2.835***	43.570***	2.224***
	(0.480)	(0.612)	(8.530)	(0.592)
	Panel B: Including Role Fixed Effects			
	(1)	(2)	(3)	(4)
	Weeks to Train Worker	Weeks Until Worker Was as Good as Average Worker	Weeks Worker is Expected to Remain at Firm	Weeks Needed to Replace Worker
Long-Term Worker	1.904***	2.465***	51.370***	2.185***
	(0.439)	(0.573)	(8.690)	(0.620)
F-Test of Joint Significance of Role Coefficients	3.076	3.921	6.283	3.059
P-Value of Joint Significance of Role Coefficients	0.00	0.00	0.00	0.00
	Panel C: Decomposition of Role Effect			
	(1)	(2)	(3)	(4)
	Weeks to Train Worker	Weeks Until Worker Was as Good as Average Worker	Weeks Worker is Expected to Remain at Firm	Weeks Needed to Replace Worker
Share of Raw Coefficient Explained by Role Effects	.2	.13	-.18	.02

Notes: Each column in Panel A and B is from a regression comparing the training and adjustment costs of hires at a firm on an indicator for whether the hire was for a regular (rather than short-term) role. There are up to two observations per firm. Standard errors are clustered at the firm level. Panel B adds indicator variables for the 32 roles we observe in our dataset; the F-test statistic is the result of a joint test that they are significant. Panel C reports the share of the coefficient in Panel A that is reduced once we add role coefficients. A negative coefficient implies that the size of the coefficient increases once role fixed effects are added.

F Model proofs

F.1 Proof of Proposition 1

We need to show that the cut-offs exist and are unique. We will appeal to two well-established results:

> **Result 1 (Intermediate-value theorem)** If f is a real-valued continuous function, and there exists a,b such that $f(a) > 0$ and $f(b) < 0$, then it must be the case that there exists c s.t. $f(c) = 0$.
>
> **Result 2 (Uniqueness of roots of a monotonic function)**: If f is a real-valued strictly monotonic function, then it has at most one root.

Define the following functions:

$$\Delta V_{12} = V_1(\theta, \eta) - V_2(\theta, \eta) \tag{13}$$

$$\Delta V_{01} = V_1(\theta, \eta) - V_0 \tag{14}$$

$$\Delta V_{02} = V_1(\theta, \eta) - V_0 \tag{15}$$

F.1.1 Proof of existence

Note that as $\eta \to \infty$, $\Delta V_{01} \to \infty$ and $\Delta V_{02} \to \infty$, similarly, as $\eta \to -\infty$ $V_{01} \to -\infty$ and $V_{02} \to \infty$. By the Intermediate Value Theorem, this suggests that ΔV_{01} and ΔV_{02} have at least one root.

We can write ΔV_{12} as follows:

$$\begin{aligned}\Delta V_{12} =&[\theta_0 + \eta_i - w - c + \frac{\beta(1-q)\sigma[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w})]}{1-\beta(1-q)} + \frac{\beta q}{1-\beta(1-q)}V_0]- \\ &[\theta_0(1+\tau) + \frac{1}{1-\beta}(\beta\mu(1+\tau) + \eta_i - w) - c]\end{aligned} \tag{16}$$

Notice that $V_2(\theta_0, \eta_i)$ is linear in η_i, whereas $V_1(\theta_0, \eta_i)$ is approximately linear in η, this is because the only non-linear terms in V_1 are $\Phi(\tilde{w})$, which is bounded between 0 and 1, and $\phi(\tilde{w})$, which is bounded between 0 and $\frac{1}{\sqrt{2\pi}}$. We can therefore re-write this as a linear function of η, along with terms that are bounded, and other terms that are constant.

$$\Delta V_{12} = \left[\frac{\beta(1-q)\sigma(1-\Phi(\tilde{w}))}{1-\beta(1-q)} - \frac{\beta}{1-\beta}\right]\eta_i + \text{Bounded terms} + \text{Constant terms} \tag{17}$$

Notice that in anything other than a knife-edge case this will be a function that tends either to be positive or negative infinity as $\eta \to \infty$ and $\eta \to -\infty$.

By the Intermediate Value Theorem, it must be the case that ΔV_{01}, ΔV_{02},ΔV_{12} cross zero at least once, which means that η_{01}, η_{02}, and η_{12} exist.

F.1.2 Proof of uniqueness

(1) Uniqueness of η_{02}

The derivative of the $\frac{\partial V_2}{\partial \eta}$ is straightforward:

$$\frac{\partial V_2}{\partial \eta} = \frac{1}{1-\beta}$$

Since it is strictly positive everywhere, ΔV_{02} is strictly monotonic and hence crosses zero exactly once. Therefore η_{02} is unique.

(2) Uniqueness of η_{01}

Consider the derivatives of $\frac{\partial V_1}{\partial \eta}$

$$\frac{\partial V_1}{\partial \eta} = \frac{\partial}{\partial \eta_i}\left[\eta + \frac{\beta(1-q)\sigma\left[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w})\right]}{1-\beta(1-q)}\right]$$

Given that $\tilde{w} = \frac{w-\eta_i-\mu}{\sigma}$, the derivative of $\tilde{w}$ with respect to η_i is:

$$\frac{\partial \tilde{w}}{\partial \eta_i} = -\frac{1}{\sigma}$$

Coming to the brackets, starting with $-\tilde{w}(1-\Phi(\tilde{w}))$, differentiating using the product rule:

$$\frac{\partial}{\partial \eta_i}\left[-\tilde{w}(1-\Phi(\tilde{w}))\right] = -\left(\frac{\partial \tilde{w}}{\partial \eta_i}(1-\Phi(\tilde{w})) + \tilde{w}\frac{\partial}{\partial \eta_i}(1-\Phi(\tilde{w}))\right).$$

Substituting $\frac{\partial \tilde{w}}{\partial \eta_i} = -\frac{1}{\sigma}$ and $\frac{\partial(1-\Phi(\tilde{w}))}{\partial \eta_i} = -\phi(\tilde{w}) \cdot \frac{1}{\sigma}$:

$$\frac{\partial}{\partial \eta_i}\left[-\tilde{w}(1-\Phi(\tilde{w}))\right] = \frac{1}{\sigma}(1-\Phi(\tilde{w})) - \frac{\tilde{w}}{\sigma}\phi(\tilde{w}).$$

Differentiating $\phi(\tilde{w})$:

$$\frac{\partial \phi(\tilde{w})}{\partial \eta_i} = \frac{\partial \phi(\tilde{w})}{\partial \tilde{w}} \cdot \frac{\partial \tilde{w}}{\partial \eta_i}.$$

Using $\frac{\partial \phi(\tilde{w})}{\partial \tilde{w}} = -\tilde{w}\phi(\tilde{w})$ and $\frac{\partial \tilde{w}}{\partial \eta_i} = -\frac{1}{\sigma}$:

$$\frac{\partial \phi(\tilde{w})}{\partial \eta_i} = \tilde{w} \cdot \frac{\phi(\tilde{w})}{\sigma}.$$

Combining the two derivatives:

$$\frac{\partial}{\partial \eta_i}\left[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w})\right] = \frac{1}{\sigma}(1-\Phi(\tilde{w})).$$

Then,

$$\frac{\partial V_1}{\partial \eta} = 1 + \frac{\beta(1-q)}{1-\beta(1-q)} \underbrace{(1-\Phi(\tilde{w}))}_{A}$$

Notice that term A is bounded, the smallest value that it attains is 0, while the highest value is 1. It follows that $\frac{\partial V_1}{\partial \eta} > 0$ everywhere and so η_{01} is unique.

(3) Uniqueness of η_{12}

The derivative of $\Delta V_{12}(\theta_0, \eta_i)$ with respect to η_i is:

$$\frac{\partial \Delta V_{12}(\theta_0, \eta_i)}{\partial \eta_i} = 1 + \frac{\beta(1-q)}{1-\beta(1-q)} (1-\Phi(\tilde{w})) - \frac{1}{1-\beta}$$

$$= \frac{\beta(1-q)}{1-\beta(1-q)} \left[\underbrace{\frac{1}{\sigma}(1-\Phi(\tilde{w}))}_{A} \right] - \frac{\beta}{1-\beta}$$

Notice that **for $\sigma > 1$ term A is bounded by 1**. Since $q > 0$, it follows that $\frac{\partial \Delta V_{12}(\theta_0, \eta_i)}{\partial \eta_i} < 0$ is negative everywhere, and therefore strictly monotonic. Since it is negative it further suggests that in the area to the left η_{12}, short-term contracts are preferred, and in the area to the right long-term contracts are preferred.

□

F.2 Proof of Prediction 1

To prove that an increase in σ will increase the share of contracts hired on short-term contracts, I will show that the threshold value to hire on a long-term contract relative to not hiring at all increase, that is, $\frac{\eta_{02}}{\sigma} > 0$ and that the threshold to hire on a long-term contract relative to a short-term contract increases, that is, $\frac{\eta_{12}}{\sigma} > 0$. Together these suggests that the share of long-term contracts will decrease.

The proof will proceed by showing (in order): (1) $\frac{\partial V_2}{\partial \sigma} = 0$ (2) $\frac{\partial V_0}{\partial \sigma} > 0$, (3) $\frac{\partial V_1}{\partial \sigma} > 0$, (4) $\frac{\eta_{02}}{\sigma} > 0$ (5) $\frac{\eta_{12}}{\sigma} > 0$ Since the proof for long-term contracts follows exactly the proof for short-term contracts, with a major simplification since V_0 does not feature in the value function of V_2, we will only prove the case for short-term contracts.

F.2.1 Proof $\frac{\partial V_2}{\partial \sigma} = 0$

This follows directly from differentiation as V_2 is not a function of σ.

F.2.2 Proof $\frac{\partial V_0}{\partial \sigma} > 0$

We will prove this in two parts, first we will derive the derivative of V_1 w.r.t to σ and then will use it to show that show that the derivative of V_0 w.r.t σ is positive.

(1) Derivation of $\frac{\partial V_1}{\partial \sigma}$

Re-write V_1 as follows:

$$V_1 = \tilde{V}_1(\eta, \theta_0, \sigma) + \frac{\beta q}{1 - \beta(1-q)} V_0 \tag{18}$$

Start by focusing on the derivative of $\tilde{V}_1$. Re-write it as follows:

$$\tilde{V}_1 = \overset{terms\ that\ do\ not}{\underset{depend\ on\ \sigma}{}} + \frac{\beta(1-q)\sigma[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w})]}{1-\beta(1-q)} \tag{19}$$

$$\frac{\beta(1-q)}{1-\beta(1-q)} \cdot \left[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w})\right] + \sigma\left[\frac{\partial}{\partial\sigma}(-\tilde{w}(1-\Phi(\tilde{w}))) + \frac{\partial\phi(\tilde{w})}{\partial\sigma}\right]\Bigg]$$

Since $\tilde{w} = \frac{w-\eta_i-\mu}{\sigma}$, we have:

$$\frac{\partial\tilde{w}}{\partial\sigma} = -\frac{\tilde{w}}{\sigma}.$$

The derivative of $-\tilde{w}(1-\Phi(\tilde{w}))$ with respect to σ:

$$\frac{\partial}{\partial\sigma}\left[-\tilde{w}(1-\Phi(\tilde{w}))\right] = \frac{\tilde{w}}{\sigma}(1-\Phi(\tilde{w})) - \frac{\tilde{w}^2\phi(\tilde{w})}{\sigma}.$$

The derivative of $\phi(\tilde{w})$ with respect to σ:

$$\frac{\partial\phi(\tilde{w})}{\partial\sigma} = \frac{\tilde{w}^2\phi(\tilde{w})}{\sigma}.$$

Combining we have:

$$\frac{\partial}{\partial\sigma}\left[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w})\right] = \frac{\tilde{w}}{\sigma}(1-\Phi(\tilde{w}))$$

Thus, the derivative of the full term with respect to σ is:

$$\frac{\partial\tilde{V}_1(\theta_0,\eta_i)}{\partial\sigma} = \frac{\beta(1-q)}{1-\beta(1-q)}\left[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w}) + \sigma\frac{\tilde{w}}{\sigma}(1-\Phi(\tilde{w}))\right]$$

After simplification, we get:

$$\frac{\partial\tilde{V}_1(\theta_0,\eta_i)}{\partial\sigma} = \underbrace{\frac{\beta(1-q)}{1-\beta(1-q)}\phi(\tilde{w})}_{>0}$$

It follows then that

$$\frac{\partial V_1(\theta_0,\eta_i)}{\partial\sigma} = \frac{\partial\tilde{V}_1(\theta_0,\eta_i)}{\partial\sigma} + \frac{\beta q}{1-\beta(1-q)}\frac{\partial V_0}{\partial\sigma}$$

(2) Derivation of $\frac{\partial V_0}{\partial\sigma}$

Recall that,

$$V_0 = \frac{\beta h}{1-\beta(1-h)} \left\{ E_{\theta,\eta} \max\{V_0, V_1, V_2\} \right\}$$

We can re-write this as:

$$V_0 = \frac{\beta h}{1-\beta(1-h)} \left\{ E_{\theta,\eta} 1\{V_0^*\} V_0 + 1\{V_1^*\} V_1(\theta,\eta) + 1\{V_2^*\} V_2(\theta,\eta) \right\}$$

where the indicator variables indicate the range of $\eta(\theta)$ values where each contract is optimal. Note that these ranges are fully characterized in terms of η_{01}, η_{02}, and η_{12}, however since we don't know the ordering of the cut-offs (i.e. whether, e.g., $\eta_{02} < \eta_{12}$), we cannot explicitly write out these regions in terms of the cutoffs. Note, however, that since these areas are fully characterized but the cutoffs, and the cut-offs are maximizers of V_0, the Envelope Theorem implies that when looking at how V_0 changes w.r.t to any variable x, we do not need to consider changes that operate through changes in the cut-off values. As such, we can write:

$$\frac{\partial V_0}{\partial \sigma} = \frac{\beta h}{1-\beta(1-h)} E_{\theta,\eta} \left\{ 1\{V_0^*\} \frac{\partial V_0}{\partial \sigma} + 1\{V_1^*\} \frac{\partial V_1}{\partial \sigma}(\theta,\eta) + 1\{V_0^*\} \frac{\partial V_2}{\partial \sigma} \right\}$$

Plugging in the expression for the partial derivative of V_1, we get that:

$$\frac{\partial V_0}{\partial \sigma} = \frac{\beta h}{1-\beta(1-h)} E_{\theta,\eta} \left\{ 1\{V_0^*\} \frac{\partial V_0}{\partial \sigma} + 1\{V_1^*\} \left[\frac{\partial \tilde{V}_1}{\partial \sigma}(\theta,\eta) + \frac{\beta q}{1-\beta(1-q)} \frac{\partial V_0}{\partial \sigma} \right] + 1\{V_0^*\} \underbrace{\frac{\partial V_2}{\partial \sigma}}_{=0 \text{ (by proof F.2.1)}} \right\}$$

Re-shuffling:

$$\frac{\partial V_0}{\partial \sigma} = \frac{\beta h}{1-\beta+\beta h\, E_{\theta,\eta}\left[1 - 1\{V_1^*\} - \frac{\beta q}{1-\beta(1-q)} 1\{V_0^*\}\right]} E_{\theta,\eta}\left[1\{V_1^*\} \frac{\partial \tilde{V}_1}{\partial \sigma}(\theta,\eta) \right]$$

Note that $E_{\theta,\eta} 1\{V_1^*\}$ is just the probability that a short-term contract is optimal. As such the following holds true:

$$E_{\theta,\eta} 1\{V_0^*\} + E_{\theta,\eta} 1\{V_1^*\} + E_{\theta,\eta} 1\{V_2^*\} = 1$$

Which in turn suggests that,

$$E_{\theta,\eta}\left[1 - 1\{V_1^*\} - \frac{\beta q}{1-\beta(1-q)} 1\{V_0^*\}\right] > 0$$

Therefore,

$$\frac{\partial V_0}{\partial \sigma} > 0$$

F.2.3 Proof $\frac{\partial V_1}{\partial \sigma} > 0$

This follows directly from proof F.2.2 and the definition of V_1.

F.2.4 Proof $\frac{\partial \eta_{02}}{\partial \sigma} > 0$

Note first, that $\frac{\partial V_2}{\partial \sigma} = 0$. Totally differentiating the definition of η_{02} w.r.t σ, we get:

$$\frac{\partial V_2(\eta_{02}, \theta_0)}{\partial \eta} \frac{\partial \eta_{02}}{\partial \sigma} = \frac{\partial V_0}{\partial \sigma}$$

Re-arranging:

$$\frac{\partial \eta_{02}}{\partial \sigma} = \frac{\partial V_0}{\partial \sigma} \cdot \left[\frac{\partial V_2(\eta_{02}, \theta_0)}{\partial \eta} \right]^{-1} \tag{20}$$

We've shown in proof F.1.2 that $\frac{\partial V_2(\eta_{02}, \theta_0)}{\partial \eta} > 0$. Likewise, we've shown in proof F.2.2 that $\frac{\partial V_0}{\partial \sigma} > 0$. As such,

$$\frac{\partial \eta_{02}}{\partial \sigma} > 0$$

F.2.5 Proof $\frac{\partial \eta_{12}}{\partial \sigma} > 0$

Totally differentiating the definition of η_{12} w.r.t σ, we get:

$$\frac{\partial V_2(\eta_{12}, \theta_0)}{\partial \eta} \frac{\partial \eta_{12}}{\partial \sigma} = \frac{\partial V_1(\eta_{12}, \theta_0)}{\partial \sigma} + \frac{\partial V_1(\eta_{12}, \theta_0)}{\partial \eta} \cdot \frac{\partial \eta_{12}}{\partial \sigma}$$

Re-arranging, we get:

$$\frac{\partial \eta_{12}}{\partial \sigma} = \frac{\partial V_1(\eta_{12}, \theta_0)}{\partial \sigma} \left[\frac{\partial V_2(\eta_{12}, \theta_0)}{\partial \eta} - \frac{\partial V_1(\eta_{12}, \theta_0)}{\partial \eta} \right]^{-1}$$

$$\frac{\partial \eta_{12}}{\partial \sigma} = \frac{\partial V_1(\eta_{12}, \theta_0)}{\partial \sigma} \left[\frac{-\partial \Delta V_{12}(\eta_{12}, \theta_0,)}{\partial \eta} \right]^{-1}$$

We've shown in proof F.1.2 that $\frac{\partial \Delta V_{12}(\theta_0, \eta_{12})}{\partial \eta} < 0$ everywhere. Moreover, we've shown in proof F.2.3 that $\frac{\partial V_1(\eta_{01} \theta_0)}{\partial \sigma} > 0$. As such,

$$\frac{\partial \eta_{12}}{\partial \sigma} > 0$$

□

F.3 Proof of Prediction 2

We begin by considering how a change in on-boarding costs for a particular contract type impacts the share of workers hired on that contract type. The proof will proceed by showing (in order):

(1) $\frac{\partial V_0}{\partial c_1} < 0$, $\frac{\partial V_2}{\partial c_1} = 0$ (2) $\frac{\partial V_1}{\partial c_1} < 0$, (3) $\frac{\partial \eta_{01}}{\partial c_1} > 0$, (4) $\frac{\partial \eta_{12}}{\partial c_1} < 0$, (5) $\frac{\partial \eta_{02}}{\partial c_1} < 0$. Since the proof for long-term contracts follows exactly the proof for short-term contracts, with a major simplification since V_0 does not feature in the value function of V_2, we will only prove the case for short-term contracts.

F.3.1 Proof $\frac{\partial V_1}{\partial c_1} < 0$, $\frac{\partial V_2}{\partial c_1} = 0$

Differentiating the expressions for V_1 and V_2 w.r.t c_1, we get that,

$$\frac{\partial V_1}{\partial c_1} = -1 + \frac{\beta q}{1-\beta(1-q)}\frac{\partial V_0}{\partial c_1}$$

$$\frac{\partial V_2}{\partial c_1} = 0$$

F.3.2 Proof $\frac{\partial V_0}{\partial c_1} < 0$

Differentiating the expression for V_0 w.r.t c_1, we get that,

$$\frac{\partial V_0}{\partial c_1} = \frac{\beta h}{1-\beta(1-h)} E_{\theta,\eta}\left\{ 1\{V_0^*\}\frac{\partial V_0}{\partial c_1} + 1\{V_1^*\}\frac{\partial V_1(\theta,\eta)}{\partial c_1} + 1\{V_2^*\} \underbrace{\frac{\partial V_2(\theta,\eta)}{\partial c_1}}_{=0 \text{ (by proof F.3.2)}} \right\}$$

$$\frac{\partial V_0}{\partial c_1} = \frac{\beta h}{1-\beta(1-h)} E_{\theta,\eta}\left\{ 1(V_0^*)\frac{\partial V_0}{\partial c_1} + 1(V_1^*)\left(-1 + \frac{\beta q}{1-\beta(1-q)}\frac{\partial V_0}{\partial c_1}\right)\right\}$$

Re-arranging we get,

$$\frac{\partial V_0}{\partial c_1} = -\underbrace{\frac{\beta h}{1-\beta+\beta h(1-E_{\theta,\eta}1(V_0^*) - \frac{\beta q}{1-\beta(1-q)}E_{\theta,\eta}1(V_1^*))}}_{\in(0,1) \text{ (following the same argument as in proof F.2.2)}} E_{\theta,\eta}1(V_1^*)$$

And so,

$$-1 < \frac{\partial V_0}{\partial c_1} < 0$$

F.3.3 Proof $\frac{\partial \eta_{01}}{\partial c_1} > 0$

Totally differentiating the expression for η_{01} w.r.t to c_1. we get that:

$$\frac{\partial V_1(\eta_{01},\theta_0)}{\partial \eta_{01}}\frac{\partial \eta_{01}}{\partial c_1} + \frac{\partial V_1(\eta_{01},\theta_0)}{\partial c_1} = \frac{\partial V_0}{\partial c_1}$$

$$\underbrace{\frac{\partial V_1(\eta_{01},\theta_0)}{\partial \eta_{01}}}_{>0 \text{ (by proof F.1.2)}} \frac{\partial \eta_{01}}{\partial c_1} = 1 + \underbrace{\left[1 - \frac{\beta q}{1-\beta(1-q)}\right]}_{\in(0,1)} \underbrace{\frac{\partial V_0}{\partial c_1}}_{\in(-1,0) \text{ (by proof F.3.2)}}$$

which implies,

$$\frac{\partial \eta_{01}}{\partial c_1} > 0$$

F.3.4 Proof $\frac{\partial \eta_{12}}{\partial c_1} < 0$

Totally differentiating the definition of η_{12} w.r.t c_1 we get:

$$\frac{\partial V_1(\eta_{12},\theta_0)}{\partial c_1} + \frac{\partial V_1(\eta_{12}\theta_0)}{\partial \eta} \cdot \frac{\partial \eta_{12}}{\partial c_1} = \frac{\partial V_2(\eta_{12},\theta_0)}{\partial \eta} \frac{\partial \eta_{12}}{\partial c_1}$$

$$\underbrace{\frac{-\partial \Delta V_{12}(\eta_{12}\theta_0)}{\partial \eta}}_{>0 \text{ (by proof F.1.2)}} \cdot \frac{\partial \eta_{12}}{\partial c_1} = \underbrace{\frac{\partial V_1(\eta_{01}\theta_0)}{\partial c_1}}_{<0 \text{ (by proof F.3.1)}}$$

which implies,

$$\frac{\partial \eta_{12}}{\partial c_1} < 0$$

F.3.5 Proof $\frac{\partial \eta_{02}}{\partial c_1} < 0$

Totally differentiating the definition of η_{02} w.r.t c_1 we get:

$$\underbrace{\frac{\partial V_2(\eta_{02}\theta_0)}{\partial \eta}}_{>0 \text{ (by proof F.1.2)}} \cdot \frac{\partial \eta_{02}}{\partial c_1} = \underbrace{\frac{\partial V_0}{\partial c_1}}_{<0 \text{ (by proof F.3.2)}}$$

$$\frac{\partial \eta_{02}}{\partial c_1} < 0$$

As such, we have shown that the threshold value to hire under a short-term contract relative to not hiring at all increases, the threshold value to hire under a long-term not to hire at all decreases, and the threshold value to hire under a long-term contract relative to a short-term contract decreases. As such the share of short-term contracts decreases.

□

F.3.6 Impact of general change in on-boarding costs

Let $c_1 = c_2 = c$. We will prove that $\frac{\partial \eta_{12}}{\partial c} < 0$, that is firms are more likely to hire on a long-term contract relative to a short-term contract as adjustment costs increase symmetrically for both.

Totally differentiating the definition of η_{12} w.r.t c we get:

$$\frac{\partial V_1(\eta_{12}\theta_0)}{\partial c} + \frac{\partial V_1(\eta_{12},\theta_0)}{\partial \eta}\cdot\frac{\partial \eta_{12}}{\partial c} = \frac{\partial V_2(\eta_{12},\theta_0)}{\partial c} + \frac{\partial V_2(\eta_{12},\theta_0)}{\partial \eta}\cdot\frac{\partial \eta_{12}}{\partial c}$$

$$-1 + \left[\frac{\beta q}{1-\beta(1-q)}\right]\frac{\partial V_0}{\partial c} + \frac{\partial V_1(\eta_{12},\theta_0)}{\partial \eta}\cdot\frac{\partial \eta_{12}}{\partial c} = -1 + \frac{\partial V_2(\eta_{12},\theta_0)}{\partial \eta}\cdot\frac{\partial \eta_{12}}{\partial c}$$

$$\frac{\partial \eta_{12}}{\partial c}\left[\frac{\partial V_1(\eta_{12},\theta_0)}{\partial \eta} - \frac{\partial V_2(\eta_{12},\theta_0)}{\partial \eta}\right]\cdot\frac{\partial \eta_{12}}{\partial c} = -\left[\frac{\beta q}{1-\beta(1-q)}\right]\frac{\partial V_0}{\partial c}$$

$$\frac{\partial \eta_{12}}{\partial c} = -\underbrace{\left(\frac{\partial \Delta V_{12}(\eta_{12},\theta_0)}{\partial \eta}\right)^{-1}}_{<0 \text{ (by proof F.1.2)}}\underbrace{\frac{\beta q}{1-\beta(1-q)}}_{>0}\underbrace{\frac{\partial V_0}{\partial c}}_{<0 \text{ (by proof F.3.2)}}$$

Which implies,

$$\frac{\partial \eta_{12}}{\partial c} < 0$$

□

F.4 Proof of Prediction 3

We begin by considering how a change in the productivity boost of long-term contracts impacts the share of workers hired on that contract type. The proof will proceed by showing (in order): (1) $\frac{\partial V_1}{\partial \gamma} = 0$, $\frac{\partial V_2}{\partial \gamma} > 0$ (2) $\frac{\partial V_0}{\partial \gamma} > 0$, (3) $\frac{\partial \eta_{01}}{\partial \gamma} > 0$, (4) $\frac{\partial \eta_{12}}{\partial \gamma} < 0$.

F.4.1 Proof $\frac{\partial V_1}{\partial \gamma} = 0$, $\frac{\partial V_2}{\partial \gamma} > 0$

Differentiating the expressions for V_1 and V_2 w.r.t γ, we get that,

$$\frac{\partial V_2}{\partial \gamma} = \theta_0 + \frac{\beta\mu}{1-\beta}$$

$$\frac{\partial V_1}{\partial \gamma} = 0$$

F.4.2 Proof $\frac{\partial V_0}{\partial \gamma} > 0$

Differentiating the expression for V_0 w.r.t γ, we get that,

$$\frac{\partial V_0}{\partial \gamma} = \frac{\beta h}{1-\beta(1-h)} E_{\theta,\eta}\left\{1\{V_0^*\}\frac{\partial V_0}{\partial \gamma} + 1\{V_1^*\}\underbrace{\frac{\partial V_1(\theta,\eta)}{\partial \gamma}}_{=0 \text{ (by proof F.4.1)}} + 1\{V_2^*\}\frac{\partial V_2(\theta,\eta)}{\partial \gamma}\right\}$$

$$\frac{\partial V_0}{\partial \gamma} = \frac{\beta h}{1-\beta(1-h)} E_{\theta,\eta}\left\{1(V_0^*)\frac{\partial V_0}{\partial \gamma} + 1(V_2^*)\frac{\partial V_2(\theta,\eta)}{\partial \gamma}\right\}$$

Re-arranging we get,

$$\frac{\partial V_0}{\partial \gamma} = \underbrace{\frac{\beta h}{1-\beta+\beta h(1-E_{\theta,\eta}1(V_0^*))}}_{\in(0,1) \text{ (following the same argument as in proof F.2.2)}} \quad E_{\theta,\eta}\underbrace{\left\{1(V_2*)\frac{\partial V_2}{\partial \gamma}\right\}}_{>0 \text{ (by proof F.4.1)}}$$

And so,

$$\frac{\partial V_0}{\partial \gamma} > 0$$

F.4.3 Proof $\frac{\partial \eta_{01}}{\partial \gamma} > 0$

Totally differentiating the expression for η_{01} w.r.t to γ. we get that:

$$\frac{\partial V_1(\eta_{01},\theta_0)}{\partial \eta_{01}}\frac{\partial \eta_{01}}{\partial \gamma} + \underbrace{\frac{\partial V_1(\eta_{01},\theta_0)}{\partial \gamma}}_{=0 \text{ (by proof F.4.1)}} = \frac{\partial V_0}{\partial \gamma}$$

$$\underbrace{\frac{\partial V_1(\eta_{01},\theta_0)}{\partial \eta_{01}}}_{>0 \text{ (by proof F.1.2)}}\frac{\partial \eta_{01}}{\partial \gamma} = \underbrace{\frac{\partial V_0}{\partial \gamma}}_{>0 \text{ (by proof F.4.2)}}$$

Which implies,

$$\frac{\partial \eta_{01}}{\partial \gamma} > 0$$

F.4.4 Proof $\frac{\partial \eta_{12}}{\partial \gamma} < 0$

Totally differentiating the definition of η_{12} w.r.t γ we get:

$$\frac{\partial V_1(\eta_{12},\theta_0)}{\partial \eta_{12}}\frac{\partial \eta_{12}}{\partial \gamma} + \underbrace{\frac{\partial V_1(\eta_{12},\theta_0)}{\partial \gamma}}_{=0 \text{ (by proof F.4.1)}} = \frac{\partial V_2(\eta_{12},\theta_0)}{\partial \eta}\frac{\partial \eta_{12}}{\partial \gamma} + \frac{\partial V_2(\eta_{12},\theta_0)}{\partial \gamma}$$

Re-arranging,

$$\underbrace{\frac{\partial \Delta V_{12}(\eta_{12},\theta_0)}{\partial \eta}}_{<0 \text{ (by proof F.1.2)}}\cdot\frac{\partial \eta_{12}}{\partial \gamma} = \underbrace{\frac{\partial V_2(\eta_{12},\theta_0)}{\partial \gamma}}_{>0 \text{ (by proof F.4.1)}}$$

Which implies,

$$\frac{\partial \eta_{12}}{\partial \gamma} < 0$$

□

G Additional Appendix Figure

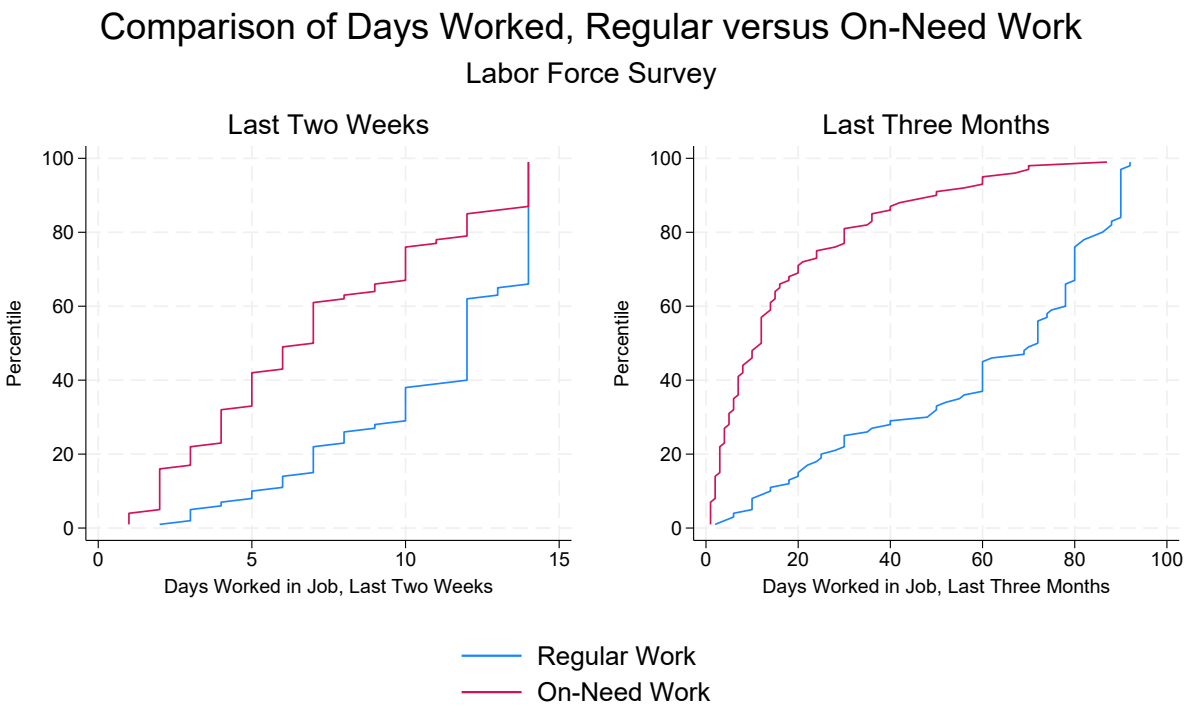


Figure A14: Comparison of Days Worked by Type of Job

H Additional Appendix Tables

Table A3: Relationship Between Margins of Adjustment in Sales and Variation in Business Demand

	Panel A: Relationship between Margins of Adjustment in Sales and Variation in Business Demand			
	(1) Uses On-Need Staff in Response to Business Variation	(2) Varies the Number of Regular Staff Across Days in Response to Business Variation	(3) Varies the Salary Regular Staff Receive in Response to Business Variation	(4) Share of Staff Hired on On-Need Basis
$\hat{\sigma}$ - Sales	0.225***	0.128**	0.138**	0.099***
	(0.062)	(0.058)	(0.055)	(0.030)
Observations	504	504	504	493
Mean of Outcome	0.45	0.28	0.32	0.16
R-Squared	0.121	0.084	0.204	0.100
SD of Outcome	0.50	0.45	0.47	0.21
	Panel B: Total Staffing			
	(1) Coefficient of Variation (CV): Total Staff	(2) CV: On-Need Staff	(3) CV: Regular Employees, Paid Daily	(4) CV: Regular Employees, Paid Monthly
$\hat{\sigma}$ - Sales	0.149***	0.304***	0.065**	0.031*
	(0.037)	(0.091)	(0.032)	(0.017)
Observations	494	504	495	503
Mean of Outcome	0.17	0.54	0.05	0.03
SD of Outcome	0.27	0.68	0.17	0.10
χ^2: Coefficient = Column (2)			6.1	9.9
p-value: Coefficient = Column (2)			0.013	0.002
	Panel C: Total Wage Bill			
	(1) CV: Total Staff Wage Bill	(2) CV: On-Need Staff Wage Bill	(3) CV: Regular Employees, Paid Daily Wage Bill	(4) CV: Regular Employees, Paid Monthly Wage Bill
$\hat{\sigma}$ - Sales	0.243***	0.303***	0.128***	0.082***
	(0.049)	(0.091)	(0.039)	(0.030)
Observations	439	501	462	483
Mean of Outcome	0.29	0.54	0.13	0.07
SD of Outcome	0.36	0.68	0.25	0.19
χ^2: Coefficient = Column (2)			3.2	6.3
p-value: Coefficient = Column (2)			0.075	0.012

Notes: Each column is from a regression of firms' use of staffing and wage bill adjustment behavior on the estimated $\hat{\sigma}$ of the firm's distribution of sales. All regressions include sector fixed effects, neighborhood fixed effects, and controls for the total number of (regular) employees. Standard errors are Huber-White heteroskedastic. The differences in sample sizes reflect don't know(s) or refusal responses from respondents; aggregates are only calculated for cases with all non-missing values. The χ^2 tests reported in columns 3 and 4 are from a seemingly unrelated regression test of the null hypothesis that the coefficients in those columns are equal to the coefficients in column 2, i.e., that firms adjust their regular daily and monthly staff (in terms of staffing and wages, respectively) to the same degree as they do their on-need staff. We present the same results with $\hat{\sigma}$ in terms of customers as our independent variable in Table 4.

Table A4: Firm Response to Shocks

	(1) Job Offer	(2) On-Need Job Offer	(3) Regular Job Offer	(4) Log Wage	(5) Wage, if Offered Job	(6) Wage, All Profiles
Next Three Weeks Busy	0.417***	0.373***	0.044**	0.066**	28.131	270.639***
	(0.024)	(0.025)	(0.021)	(0.028)	(17.551)	(17.507)
Observations	4152	4150	4150	1889	1889	4139
Mean of Outcome	0.46	0.22	0.24	6.37	657.74	300.19
R-Squared	0.511	0.376	0.421	0.790	0.783	0.515
SD of Outcome	0.50	0.41	0.43	0.51	336.22	398.67

Notes: Each column is from a regression comparing firms' willingness to hire a candidate under their current circumstances (omitted) and if the next three weeks were busy. Firm fixed effects and controls for profile covariates are included. Standard errors are clustered at the business level. Outcomes (1) to (3) are dummy variables; outcomes (5) to (6) are winsorized at the 99% and 1% level.

Table A5: Regression of Worker's Characteristics on Willingness to Offer a Job Offer

	(1) Job Offer	(2) Job Offer if any Position Was Open	(3) Rating
Any Experience	0.031	0.040	0.267***
	(0.026)	(0.029)	(0.077)
High Level of Experience	0.006	0.028	-0.052
	(0.024)	(0.027)	(0.077)
Same Sector Experience	0.073***	0.139***	0.143*
	(0.023)	(0.026)	(0.077)
High * Same Sector Experience	0.025	-0.006	0.255**
	(0.033)	(0.036)	(0.111)
Observations	3559	3556	3560
Mean of Outcome	0.40	0.67	4.99
R-Squared	0.560	0.424	0.411
SD of Outcome	0.49	0.47	1.38

Notes: Each column is from a regression of firms' responses to hypothetical workers in our vignette section. Column (1) is an indicator for whether the firm would be willing to hire a hypothetical worker, Column (2) for whether they either (a) would be willing to hire, or (b) would be willing to hire if they were currently hiring for the role, and Column (3) is a 1/7 Likert rating of the quality of the candidate. These results show the returns to different levels of (randomized) experience in firms' willingness to make an offer. They thus provide validation that results were taken seriously, and provide a benchmark against which to compare our results with respect to demand shocks (in Appendix Table A4) and of the type of contract offered to a worker (in Table 6).

Table A6: Comparison of Vignette Responses and Use of Short-Term Labor

	(1) On-Need Offers Made	(2) Share Offers On-Need
Share Staff On-Need	0.632**	0.20**
	(0.271)	(0.091)
Observations	583	394
Sample Mean	1.02	0.430
Sample SD	1.41	0.373

Notes: Each column regresses firm's choices in the vignettes on the actual share of their labor hired on a short-term basis. Column (1) measures the total number of vignettes (of six) they were willing to make an offer to, and Column (2) the share of offers made that were for an on-need contract, among the subset of firms that made any offers. Both regressions include sector fixed effects, neighborhood fixed effects, and controls for the total number of (regular) employees. Standard errors are Huber-White heteroskedastic.

80