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

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

September 2024

Abstract

We study why firms in Nairobi, Kenya offer short-term and insecure labor contracts to many of their workers. We combine detailed information on 601 firms' hiring and contracting practices with a novel hypothetical survey instrument. We find that in deciding whether to offer long-term, stable contracts to workers, firms face a trade-off between managing demand variation on the one hand and minimizing on-boarding costs and optimizing worker performance on the other. Firms face substantial variation in product demand, much of which they pass onto their workers. We observe a strong, positive correlation between firms' demand variation and their use of short-term labor. Hypothetical vignettes corroborate these results: when faced with a demand shock, firms expect to respond by hiring short-term labor. However, bringing in new labor requires paying on-boarding costs, potentially offsetting the gains from flexibility. For the median hire, firms report relatively low costs of on-boarding new workers, making short-term arrangements viable for many matches. However, we find considerable heterogeneity across matches; hires with higher on-boarding costs are more likely to be on long-term contracts. Finally, our hypothetical vignettes show that firms perceive that the same worker performs worse when hired on a short-term contract. We bring these findings together through a model of firm contracting.

*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. 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.

1 Introduction

Labor markets in developing countries are frequently characterized by short, disrupted employment spells. Rather than climbing a job ladder, these transitions are more commonly the result of workers switching between jobs, cycling in and out self-employment, and experiencing frequent spells of unemployment. These dynamics are in sharp contrast to labor markets in higher-income countries, in which worker-firm matches are considerably more stable (Donovan et al., 2023). In the context that this paper considers–poor neighborhoods in Nairobi, Kenya–nearly half of young workers had fragmented work patterns over the past two weeks. A quarter were juggling two or more jobs. These fragmented employment patterns do not appear to be driven by worker preferences; 95% of those working less than full-time report a preference to be working more than they currently are. Instead, these patterns suggest that firms' preferences or constraints plausibly play an important role in determining the prevalence of short-term work in Nairobi and in other settings with frequent turnover and employment disruptions.

We study the incentives that shape firms' decision-making regarding contracting and commitment to workers in Nairobi, Kenya. The core finding of this paper 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. 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 to performance. Every time firms bring in new labor, they have to pay the costs of searching for and evaluating workers. New hires take time to be trained and to learn the match- and sector-specific skills that make them productive. Moreover, knowing that the employment spell may quickly come to an end for reasons beyond their control, these workers exert less effort and are less reliable. We provide evidence that the decisions of firms regarding the commitment to offer their workers depend on the relative magnitude of these forces.

We provide evidence for this trade-off using a survey of 601 firms, employing a total of 5,687 workers, recruited from a census of firms in a representative sample of enumeration areas from the major business corridors of Western and Central Nairobi. Our sample firms are primarily in the retail, hospitality, beauty, and manufacturing sectors. We collect detailed information on the recruitment, hiring, and business practices of firms, and their current stock of workers. In addition to this factual information, we also design a set

of novel survey instruments to elicit firms' hypothetical responses to various hiring and contracting decisions under different scenarios in addition to their subjective beliefs about how different contracting regimes affect firm and worker performance. These instruments allow us to examine the mechanisms underlying the trade-offs we are studying in more detail and under fewer assumptions than using observational data alone would permit.

First, we administer a survey module that specifically asks firms about busy, normal, and quiet periods for their business, including these periods' frequency, the level of demand they experience in these periods, and the way that they fill their staffing needs in these periods. We find that firms experience considerable variation in their demand (we estimate that, for the median firm, the standard deviation in the volume of sales is larger than its mean), and that, moreover, firms pass much of this variation into their workers. They bring in staff on a short-term basis, adjust their contracts with regular staff, reduce staff salaries, and pay workers on commission. Moreover, we find a strong positive association between a given firm's product demand variation and its use of flexible work arrangements, including when controlling for firm size and sector and neighborhood fixed effects. Firms in the top quintile of variation in daily number of customers have a coefficient of variation in total staff across days that is over double that of the bottom quintile, and a coefficient of variation in their wage bill that is over three times as high as that of those in the bottom quintile. We corroborate this result using vignettes on hiring under different scenarios. When firms expect short-term demand shocks, they become more willing to hire candidates, and critically, almost entirely by offering short-term, flexible arrangements.

Second, we show that adjustment costs play an important role in governing the types of contracts that firms offer. Overall, firms perceive an abundance of workers with the skills needed to do the tasks at their firm, and perceive adjustment costs to be low. 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. However, firms do perceive differences in adjustment costs for different types of workers. On average, firms report that workers on long-term contracts take 70% longer to replace, 40% longer to train, and 223% longer until they are as good as a typical worker within the firm. These two facts—-that labor is perceived to be readily available, and that higher-adjustment cost workers are more likely to be employed on a long-term basis—-speak to an important, demand-side feature of the economic environment.

Third, we also show that firms believe that the same worker will be less productive if the firm is unwilling to commit to them long-term. We present hiring managers with

"profiles" of hypothetical jobseekers, designed to mimic the content and structure of typical job applications in this context– and elicit their beliefs over the quality of a candidate, how likely they would be to be absent on a given day, and how likely they would be to quit or be fired. For a subset of these profiles, we ask firms to evaluate the same candidate under three contract regimes: as a regular employee paid monthly, as a regular worker paid daily, or brought in short-term on an on-need 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. Perhaps because they believe that workers will have a foot out the door, they report that the same worker would perform worse, be absent more often, and be more likely to quit or be fired if only offered a short-term, on-need contract.

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 short-term versus long-term contracts that form the basis of this paper.

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, in a context where they face uncertainty over their future demand for labor, and uncertainty over whether workers will quit. Whenever a new worker is brought on board, firms pay an adjustment cost. Moreover, 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.

Our empirical 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. While firms in higher-income countries might be able to effectively absorb variation and offer stability to their workers, this appears less true in Nairobi, where a large share of product variation gets passed on directly to 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.[1]

Second, these speak to the importance of policies and interventions that increase the demand for specialized labor. If it is the case 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 barrier to more stable employment. Firms are likely to perceive that they can find workers who can step in at a moment's notice, limiting their need to commit to their workers. If instead, the share of jobs for which specialization and on-boarding becomes essential, and the more the consequences of low-quality work are especially acute, the more firms will be willing to offer stability to its workers.

Our paper's primary contribution is our study of labor market dynamics and contracting in low-income countries through the lens of the underlying incentives guiding firm's decision-making. Our research therefore sits at the intersection of two literatures: a substantial and growing body of work that investigates the causes of high turnover in low-income countries, much of which focuses on matching technologies and workers' search efforts; and studies—primarily from higher-income countries—that explore variation in contract type.

Donovan et al. (2023) provides evidence from 49 countries that lower-income countries have substantially higher flows into and out of employment than higher-income countries, reflecting a slippery job ladder. This phenomenon has motivated considerable research into methods through which to improve the search and matching process. Interventions in this space include improving the ease with which workers can signal their abilities (Carranza et al., 2022; Abebe et al., 2021; Bassi and Nansamba, 2022), providing workers with better information about labor market outcomes to help them target their search (Alfonsi et al., 2022; Jones and Santos, 2022), subsidizing search (Franklin, 2018; Banerjee and Sequeira, 2020), training workers (Alfonsi et al., 2020; Hardy et al., 2019), lowering the costs of hiring or contracting (De Mel et al., 2019; Cefala et al., 2023), and improving the functioning of in-person and virtual platforms (Abebe et al., 2022; Fernando et al., 2023; Groh et al., 2015). We view our work as a complement to these studies, by examining on the demand-side determinants of this turnover, and unpacking the underlying forces at the root of firms' decisions to offer unstable, low-security employment contracts that plausibly contribute to

[1]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.

high turnover in low-income countries.

Second, we add to the literature which examines the factors behind firms' choice of employment arrangements, much of which has focused on how European firms engage with temporary contracts. Nickell (1986) and Goux et al. (2001) argue that variation in labor demand may drive the use of temporary employment contracts, while Faccini (2014) suggests that temporary contracts can serve as a screening tool for firms. Closely related, Adams-Prassl et al. (2023) explore the link between labor market concentration and the use of zero-hours contracts in the United Kingdom. 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. As we show, in our setting, labor market regulations do not appear to play a key role in the use of short-term contracts, and a sizeable share of firms (59%) use short-term contracts in some capacity. These distinctions speak to a different underlying economic environment, meaning it is not ex ante clear the same mechanisms will be operant in our context.

One notable context in which the factors influencing hiring and contract choice has been studied in low-income counties is the literature on informality (e.g., La Porta and Shleifer, 2014). For example, Ulyssea (2018) studies how labor market regulation impacts firms' decisions to hire workers on- or off-the-books in Brazil. Our work is closely related to, but distinct from this literature, much of which focuses on the trade-off between the flexibility associated *with not paying taxes or benefits* against the cost of detection. While our work also emphasizes the value of flexibility as a key motive to offer worker short-term arrangements, importantly, we find that using a mix of short-term and long-term arrangements 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 simply about the role that labor market regulations play.

Finally, our work provides a proof-of-concept for the use of "worker profiles" and 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 to great effect in many contexts (see, for instance, Bertrand and Mullainathan, 2004; Kroft et al., 2016; Eriksson and Rooth, 2014), implementing such a procedure is logistically infeasible in the context that we study, and plausibly many other low-income contexts. While 75% of hires we observe involve individuals the firm manager did not know personally, much of the application process happens through individuals physically

presenting themselves and discussing their skills and qualifications in-person, or by messaging them dyadically (for example, via WhatsApp). Kessler et al. (2019) show that asking hiring managers to rate resumes (that the managers know to be fictitious) offers a tool through which to more broadly measure firm preferences. Our work extends this further, by adapting hypothetical hiring to our specific context, asking about a broader range of beliefs and preferences, and considering hypothetical scenarios (for example, different contract structures, demand shocks), albeit at the cost of being unable to directly incentivize these choices. 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.[2]

2 Context and Sampling Procedure

Our primary analysis for this paper depends on in-depth interviews with 601 firms in Western and Central Nairobi with at least three regular employees, mostly across the retail, hospitality, beauty and manufacturing sectors. We supplement this data with statistics from IPUMS-International's 1% sample of the 2009 Kenya Census, and 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 of Kenya and the second-largest city in East Africa, and is a major hub for commerce, financial services, transportation, and technology in the region. Approximately 18% of Kenya lives within either Nairobi County or the Counties adjoining it.

Nairobi has many attributes in common with major urban centers in Sub-Saharan Africa. Nairobi's population has grown substantially, driven to a meaningful degree by in-migration. Nairobi County's population was an estimated 5.3 million people in 2023, as compared to 1.4 million in 1990 and 137,500 in 1950. From the IPUMS-International 1% sample of the 2009 Census, 69% of the Nairobi County population are migrants, including 84% of the population aged 18-35. A large share of Nairobi's population has recently

[2]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.

entered the labor force. The median Nairobi County resident is age 23; and 47% percent are between the ages of 18 and 35. Most young adults in Nairobi have completed either primary (39%) or secondary school (49%) but not tertiary school (11% have post-secondary technical education; 4% have a university degree). 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 (Ruggles et al., 2024).

2.2 Labor Market in Low-Income Neighborhoods

We supplement national and firm-level data with a survey of 427 young adults in two low-income neighborhoods in Nairobi.[3] The purpose of this labor data is twofold: first, we hypothesize that many of the attributes central to understanding labor market dynamics in Nairobi–short-term employment spells, the contractual nature of work, workers' preferences, and an employment history that explicitly targets both long-term and short-term work–are either missing or under-reported in more traditional labor force surveys. Second, we aim to measure the degree to which short-term and interrupted work spells matter for an important, defined population: young adults in low-income neighborhoods.

This survey provides evidence of several important features of the economic environment for low-income youth, detailed in Table 1, and displayed in Figure 1.

First, employment is frequently fragmented and irregular. While 89% of individuals have done some work in the last two weeks, for many individuals, this does not reflect stable, long-term employment. We plot the cumulative distribution of days worked in the last two weeks by respondents in Figure 1. 39% of individuals with any work in the last two weeks report having worked less than full-time, instead having worked between one and nine days of the last 14. Even those who manage to achieve a desired work amount often do so by cobbling together employment opportunities. Of the sample who has done work in the past two weeks, 42% have worked in multiple jobs. Of our sample who has worked for others in the last two weeks, 60% have done some work on an on-need basis, rather than in a job they are called in regularly for.

This work behavior largely does not reflect a preference for part-time work. When individuals are asked what their ideal amount of work would have been in the last two

[3]We enrolled individuals in this sample in August 2024. We identified major areas two low-income neighborhoods (Kibera and Kawangware), instructed enumerators to walk in random directions, and to survey every third individual they encountered plausibly between the ages of 18-35. Individuals were eligible if they were in the age range, and not a full-time student.

weeks (that is, if work was available, but they were not required to go to work every day), 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.

Second, to a meaningful degree, these young adults are not specializing, but rather accumulating experience across a wide range of sectors and jobs. We ask individuals about their work in the last two weeks, any employment they have done for at least a month in the last three years, and the last two additional jobs (apart from any in the last two weeks) that they have done on a short-term basis. In this sample, 74% report having done work in multiple sectors; 27% have worked in multiple sectors in the last two weeks. While this sample is frequently engaged in work, and desires full-time, stable employment, a sizeable share are unable to do so, motivating our study of the underlying incentives for firms and contract choice.

2.3 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 began by identifying major business corridors within Western and Central Nairobi.[4] 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. [5]

Next, we randomly sampled 174 of these EAs, and conducted an initial census of all eligible businesses in February and March of 2024 [6]. Surveyors were instructed to approach

[4] 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.

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

[6] Given the unequal distribution of businesses across town, we surveyed in all EAs in mixed commercial-

any business with evidence of at least three employees in any of eight target sectors.[7] Upon obtaining consent, surveyors conducted a short regarding the total number of employees and their roles within the firm, and in which we collected the contact information of the individual(s) at the company responsible for hiring and managing workers.[8] Through the census, we collected the information of 1624 firms, with 703 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 601 establishments, surveyed in March through May, 2024. Collectively, our sample is meant to be representative of "firms in Central and Western Nairobi that hire multiple workers, and that would be approachable for a jobseeker on foot. [9]

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 basis.

Firm characteristics are detailed in Table 2. The majority of our sample is concentrated in four sectors: (1) retail (44%), (2) food, beverage, and hospitality (24%), (3) beauty services (15%) and (4) manufacturing (11%). On average, firms employ a total of 9 workers.[10]

residential neighbourhoods in Western Nairobi, and randomly sampled 76 EAs in CBD, and 13 in the Industrial Area

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

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

[9] 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 lacks a representative sample of firms in Nairobi, and that their data could not be shared.

[10] This is a winsorized value at the top 1%; the unwinsorized value is 13.

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 (in Purchasing Power Parity terms; mean $1,078).

Employees at these firms can broadly be classified into three categories: regular workers paid monthly, regular workers paid daily, and short-term, or "on-need" workers (generally but not exclusively paid on a daily basis). This taxonomy is overlapping with, but not exactly aligned with formal versus informal or casual workers. [11] We distinguish contract type on the basis of whether a worker is a "regular worker," i.e., someone who is expected to come into work on a regular basis (regardless of their formality); separations only take place if either party explicitly requests that the work end. In contrast, on-need work reflects that employees are explicitly called in for a short period of time; their employment might be renewed (day-to-day, or at infrequent intervals) but there is no guarantee of continued work. We find that 61% of firms have regular workers they pay on a monthly basis, 59% have regular workers they pay on a daily basis (19% have both), and 59% of firms have hire on a short-term, on-need basis. Of those using short-term employees in the last month, 31% of employees engaged with the firm are hired on a short-term basis.

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 staffing decisions of firms. Overall, firms perceive that they can quickly and easily replace workers, limiting the incentive of them to commit to employees. However, when firms do commit longer to workers, these tend to be individuals for which 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

[11] Kenya Labor Law mandates that regular workers are paid monthly, yet we find that workers receiving government benefits consist of only a modest subset of the total workers paid monthly.

flexible stock of labor.

We complement these core results with consideration of other possible mechanisms driving firm contracting decisions, including the role that labor market regulations might play, and the importance of candidate screening. 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 core set of results are that: (1) businesses experience considerable variation in demand for their goods and services, (2) they transmit much of this variation onto their workers (rather than absorb this variation themselves), (3) firms with greater variation in demand have greater variation in their staffing and wage bills, and (4) short-term labor arrangements are a key mechanism through which firms manage 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 begin by collecting information that allows us to estimate each business' distribution of sales and customers. 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 presenting respondents with a laminated sheet with the three types of days (busy, normal, quiet) and 20 coins to allocate 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, μ and σ for each.[12] We discuss this estimation procedure, and our use of a log-normal distribution, in more detail in Appendix B.

In addition to estimating the distribution of customers and sales, we also ask firms about the different strategies they engage in in response to this variation. Specifically, 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

[12] That is, given our assumption of log-normality, the mean and standard deviation of the underlying normal distribution.

information about the number of regular workers, the number of short-term workers, and the total salaries paid to each during normal, busy and quiet periods.

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 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 (that again, might be excluded from annual surveys) as well as very short-run high-frequency panels (for whom the bulk of days might all fit into a normal bin).

In Table 3, we provide evidence both that firms experience meaningful variation, and that they engage in many variance-reduction strategies. From our estimated distributions, we calculate that the median firm has a coefficient of variation of sales in excess of one, i.e., that their standard deviation is greater than their mean. These results are not concentrated in a single sector or neighborhood. 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. Additionally, we find that 46% of businesses use short-term staff to manage variation, 29% adjust their staff salaries between high and low periods, and 32% adjust their supply of regular workers between periods. The share of firms who engage in none of these strategies is 37%.

In Table 3, we also examine the relationship between our estimated business variation and these behaviors on the extensive margin. Specifically, we estimate:

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

where y_{ins} is an outcome for firm i in neighborhood n, engaged in sector s, $\hat{\sigma_{ins}}$ is the estimated σ 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.

Our results show a consistent pattern: firms that experience greater variation are more likely to engage in strategies that pass this variation onto their workers. A one σ increase in customers is associated with firms being 29 percentage points more likely to use short-term workers, 24 percentage points more likely to adjust staff salaries, and 14 percentage points more likely to adjust the number of regular staff who come in. These differences are also associated with an 8 percentage point increase in the share of roles at the firm filled via

short-term workers.

These results are also present when examining differences in the firms' variations in total staffing and the total wage bill. In Figure 2 we show the association between firms' σ in their customers and sales, and the coefficient of variation of their staff and wage bill in Panels A-D. We present a binned scatter plot, binned at the vigintile level. Throughout the distribution, there is a positive association between business variation and variation in staffing and wages. In Table 4 we decompose this variation by type of worker (regular workers paid monthly, regular workers paid daily, and short-term workers). In both cases, the effects are strongly concentrated among short-term workers, suggesting this is a key margin through which businesses manage their variation.

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 none-the-less supplement our results with a vignette to examine how firms anticipate that they would respond to shocks in demand.

We administer a module of hypothetical vignettes related to hiring. For each hiring manager, we present six different hypothetical candidates, meant to be similar to the types of candidates who work for or apply for their jobs. For these candidates, we ask firms about their willingness to hire a given candidate under their current needs today, and if so, the type of employment arrangement they would offer to the worker (regular employee paid monthly, regular employee paid daily, temporary worker, other). Then, 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 rate whether they would be willing to hire a worker, and if so, the type of employment arrangement they would offer them.

In Figure 3 and in Appendix Table A4, we show how firms believe that this demand shock would affect their willingness to hire, and the type of hiring they would use. Firms are 42 percentage points more likely to hire a given worker; the vast majority of this hiring operates through the channel of offering them short-term contracts (there is just a 4 percentage point increase in the share who report they would hire a worker on a long-term basis). Thus, in addition to firms' actual behavior being consistent with the idea that disrupted, short-term contracts being a key way to manage business variations, firms also report that in hypothetical scenarios, this is a mechanism they would use. These results

collectively suggest that a key benefit of short-term work is the flexibility that it affords firms, particularly those for whom demand variation is especially acute.

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 dictating 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 on-boarding 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 summary statistics from 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:

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

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. [13]

We estimate that firms perceive there to be higher adjustment costs for staff hired on a long-term role for several of the dimensions we measure. Firms report that long-term hires will take 2.3 weeks longer to train on average, and 2.6 weeks longer until they become as good as the typical employee. They believe that these employees will remain with the firm longer, and would take longer to replace.

Our results are consistent both with the possibility that roles require no training, but also with the possibility that some training is required, but that the pool of workers available broadly already possess these skills. For this question, we rely on our vignettes of hiring, and find at least some evidence in favor of the second interpretation.

Specifically, in our vignettes on hiring, we experimentally manipulate the work experiences of applicants, into one of five categories: (i) no experience, (ii) low experience, different sector, (iii) high experience, different sector, (iv) low experience, same sector, and (v) high experience, same sector. In Appendix Table A2, we report the returns to these experience levels with respect to firm willingness to hire, and on rating of the candidate quality. We do find evidence that firms value workers with experience in the same sector (and are 8 percentage points more likely to say they would offer the candidate a job). However, we are unable to reject the null that more extensive experience is valued relative to any experience. We view this evidence as suggesting that at least some of the low perceived adjustment costs simply reflect that there is a large supply of labor with the requisite skills needed to do a given job; the relatively small point estimates in turn suggest that the relatively modest task complexity also plays an important role.

3.3 Behavioral Response to 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

[13] We winsorize all time outcomes at the top 1%.

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 [behavior table]. We estimate regressions of the form:

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

Where $y_{ipc}$is the rating for firm i, evaluating profile p, under contract-regime c. daily and on need 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,[14] 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 example, this structure avoids any issues with beliefs that the types of workers most often hired on a particular sort of contract are the ones most likely to quit.

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

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

(relative to no experience, in Appendix Table A6) 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, expending less effort (and potentially searching for a better opportunity) when only offered short-term employment.

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 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 both that firms have a challenge observing candidate quality, and that improving the ability of candidates to signal their quality can affect both jobseeker and firm behavior. A relevant question to consider therefore is whether the prevalence of short-term contracts reflects that firms are using this employment as a screening mechanism.

The evidence we collect leads us both to believe that screening has some relevance in the economy, but also that it is not the key driver of variation in contract choice. 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 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—more likely to not show up on a given day, will perform worse when on the job, and the firm is more likely to choose to let them go. 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). And yet, 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 employee's must receive social security and health care contributions, have tax contributions deducted, and offers sick leave and maternity leave. It 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. Thus, if these laws have actual bite in the economy, work-term stints of less than one month might be a way to circumvent

these laws.

We measure firm compliance with labor law indirectly, through our section on hiring. We ask managers about the benefits they offer to their recent hires, including both informal benefits (loans and salary advances, transportation allowances, meals) and benefits mandated by labor law (Social Health Insurance Fund / National Health Insurance Fund, and National Social Security Fund contributions). We also ask employees what information is needed from applicants when they apply to work, including education credentials, a letter of reference, and a Kenya Revenue Authority PIN (i.e., a number needed to pay taxes and a required document for formal employment). We find that just 14% of firms both pay benefits and require a KRA pin; when we use the less stringent criteria that they pay benefits, 20% of firms pay benefits.

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; if anything, firms not compliant with regulations appear slightly more 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, something 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 required benefits). However, we find that 71% of employees are re-used, and 87% 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

To provide intuition on how the stylized facts we have outlined impact the use of short-term contracts, we now turn to a simple conceptual framework. Our focus is on understanding how three factors – (1) increasing product demand volatility, (2) decreasing on-boarding costs and search frictions, and (3) increasing the "incentive effects" of long-term contracts — affect the likelihood that a worker is hired on a short-term contract. In order to focus on the impact of these on the contract choice margin, and to keep the model tractable, we do not allow for adjustments through wages, and abstract away from general equilibrium considerations.

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 from a normal distribution with mean μ_j and variance σ_j. Firms hold correct beliefs about the parameters characterising 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 enough workers to meet demand or have idle workers.[15]

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 that is less than one. The frequency at which job-seekers find vacancies is an exogenous reduced form parameter that reflects the intensity of job search. After a vacancy is filled, firms have to pay a one-time on-boarding cost, which may be heterogeneous across contract type. For simplicity, we do not allow offers to be rejected or bargaining to take place.

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 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 match has been

[15] 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.

destroyed, which occurs with probability less than one. 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.

There are two sources of heterogeneity in the model. First, firms face different productivity distributions, resulting from differences in mean productivity, the spread of productivity, or both. Second, when a worker and firm match, a worker independently draws her quality from a quality distribution. 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.[16]

In this set-up, hiring on a short-term contract presents firms with a trade-off. Short-term contracts allow firms to manage variable product demand and hence reduce slack at the firm. At the same time, workers on short-term contracts are expected to be less productive, which directly reduces the firm's profits. Lastly, since a match is more likely to be destroyed under a short-term contract, firms are exposed to the risk of not having a worker at the firm when they need one. In addition, this creates an added cost since each time a firm hires a worker they must pay an on-boarding cost.

4.1 Solution

4.1.1 State space

We can simplify the model to three states: V_0, the value of the state in which the firm does not have a worker; V_1, the value of hiring a worker on a short-term contract; and V_2, the value of hiring a worker on a long-term contract.[17]

4.1.2 Value functions

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

[16] 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.

[17] This is done to make the model stationary, in reality there are 5 states: hiring a worker (on a STC or LTC) when the firm doesn't have a worker, which involves paying c, the continuation value of having a worker on a STC or LTC at the firm, the value of not having a worker at the firm.

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

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:

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

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

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

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

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 workers productivity is given by $\theta(1+\tau)$, where $\tau > 0$ captures the 'productivity boost' associated with a long-term contract. The second term corresponds to the discounted expected future value 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:

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

Finally, the value of being without a worker is given as follows:

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

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, λ, 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 meatching with a worker today.

4.1.3 Optimality conditions

The firm's optimality conditions are characterised 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.

> **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 C

4.2 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 increases, decreasing the the cut-off value to hire on a long-term contract, η_{12}, and increasing and the share of short-term contracts.

> **Prediction 1 (product demand variability):** A firm which faces more variable product demand, will hire workers of higher quality into long-term contracts, decreasing the share of long-term employment.

Proof in appendix C

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, increase 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 C

Finally, an increase the impact of long-term contracts on the productivity of workers, directly increases the value of long-term contracts, which increases 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 C

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 documented 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

Our results are also consistent with a number of facts about labor market flows in low-income countries. Donovan et al. (2023) document high levels of turnover in low-income countries, concentrated early in workers' tenure and in lower-wage jobs. We provide a framework through which to rationalize these patters. 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 labour market–a high use of short-term contracts facilitates frequent separations. This interpretation is broadly consistent with the authors' interpretation that low-income countries may have a higher share of low-value, and high-turnover jobs.

Furthermore, We show that demand variation, together with low adjustment and training costs drive the use of these short-term contracts. Our results suggest that one hypothesis that may be consistent with the cross-country patterns documented by Donovan et al. (2023) is that firms in low-income countries face greater demand variation (or a greater tendency to pass demand variation onto their workers) and lower adjustment and training costs than firms in high-income countries, which results in greater employment risk for workers, shorter employment spells, and more frequent turnover.

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 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 none-the-less 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-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 also consistent with, albeit not dispositive of, the possibility of a "productivity trap" of sorts, in which firms and workers are stuck in a low commitment equilibrium. We document that firms perceive adjustment and training costs for their workers to be low, suggesting that firms make limited investments in the skills of their hires. Similarly, firms 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. These results suggest low commitment on both sides of the market– firms make minimal investments in, and offer limited stability to their workers, while firms perceive that short-term workers have a foot out of the door, and may not be reliable. These results are thus consistent with the hypothesis that firms and workers are stuck in a low-commitment equilibrium, posing a challenge to increased specialization and productivity growth. This pattern is 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 this is especially likely to be true in the context of farming, where there are specific stages in the agricultural cycle of high labor demand (and moreover, a time-sensitive nature to these periods), coupled with periods of lower demands. Our broader set of results regarding customers and sales variation suggest that farming is not unique in this regard. Instead, many sectors in low-income countries are likely to experience this variation (and/or a worsened ability to manage this variation, relative to higher-income countries). Our work thus suggests that the phenomenon they document may be broadly

prevalent in low-income countries, where demand variation appears to play a key role in guiding firm behavior.

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, and adjustment 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, either because they cannot readily absorb the variation themselves, or that this does not limit their ability to hire workers, there are likely to be substantial challenges in ensuring worker stability. Similarly, in labor markets where a substantial 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. So long as demand variation is prevalent, and 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.

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*.

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.

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

Banerjee, A. V. and S. Sequeira (2020). Spatial mismatches and imperfect information in the job search.

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

Bertrand, M. and S. Mullainathan (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review 94*(4), 991–1013.

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.

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

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.

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

Eriksson, S. and D.-O. Rooth (2014). Do employers use unemployment as a sorting criterion when hiring? evidence from a field experiment. *American economic review 104*(3), 1014–1039.

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.

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., 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.

Hardy, M., I. M. Mbiti, J. L. Mccasland, and I. Salcher (2019). The apprenticeship-to-work transition: experimental evidence from ghana. *World bank policy research working paper* (8851).

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

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.

Kroft, K., F. Lange, M. J. Notowidigdo, and L. F. Katz (2016). Long-term unemployment and the great recession: the role of composition, duration dependence, and nonparticipation. *Journal of Labor Economics 34*(S1), S7–S54.

La Porta, R. and A. Shleifer (2014). Informality and development. *Journal of economic perspectives 28*(3), 109–126.

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

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

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

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.

7 Figures

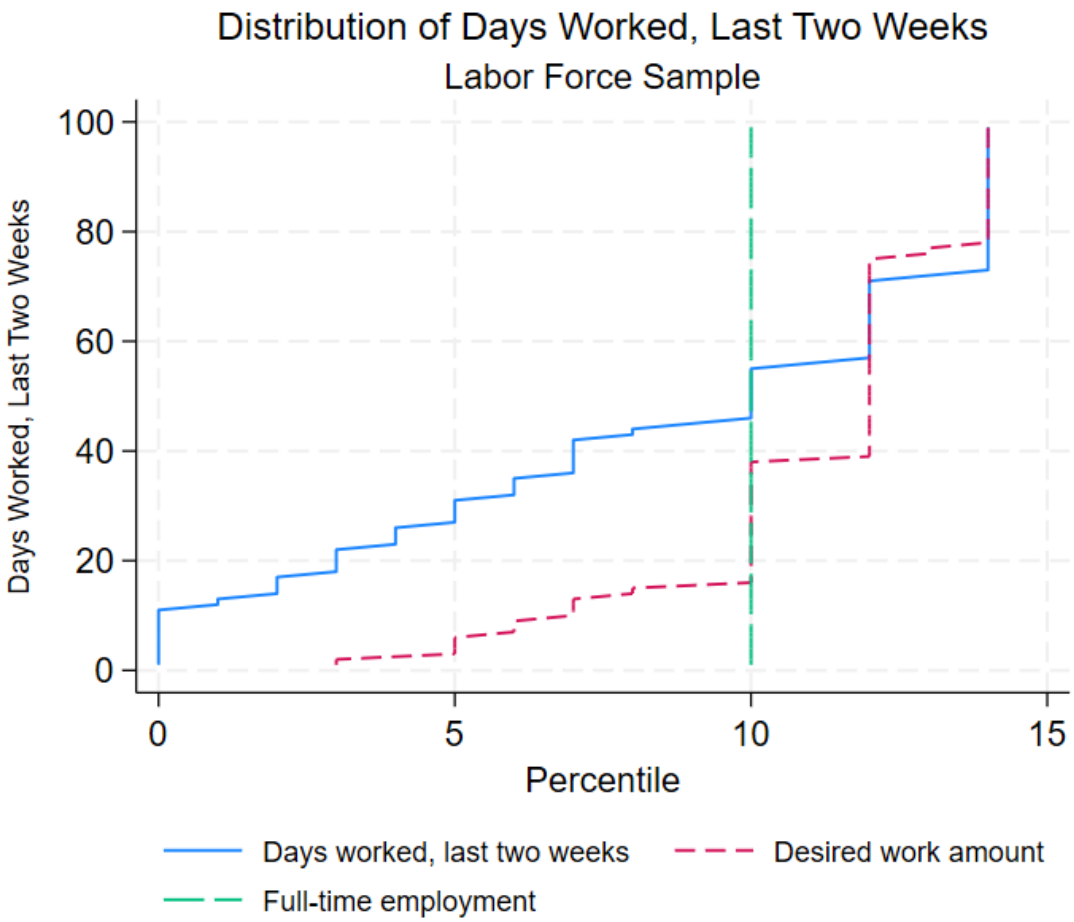


Figure 1: Distribution of Days Worked, Young Adult Sample

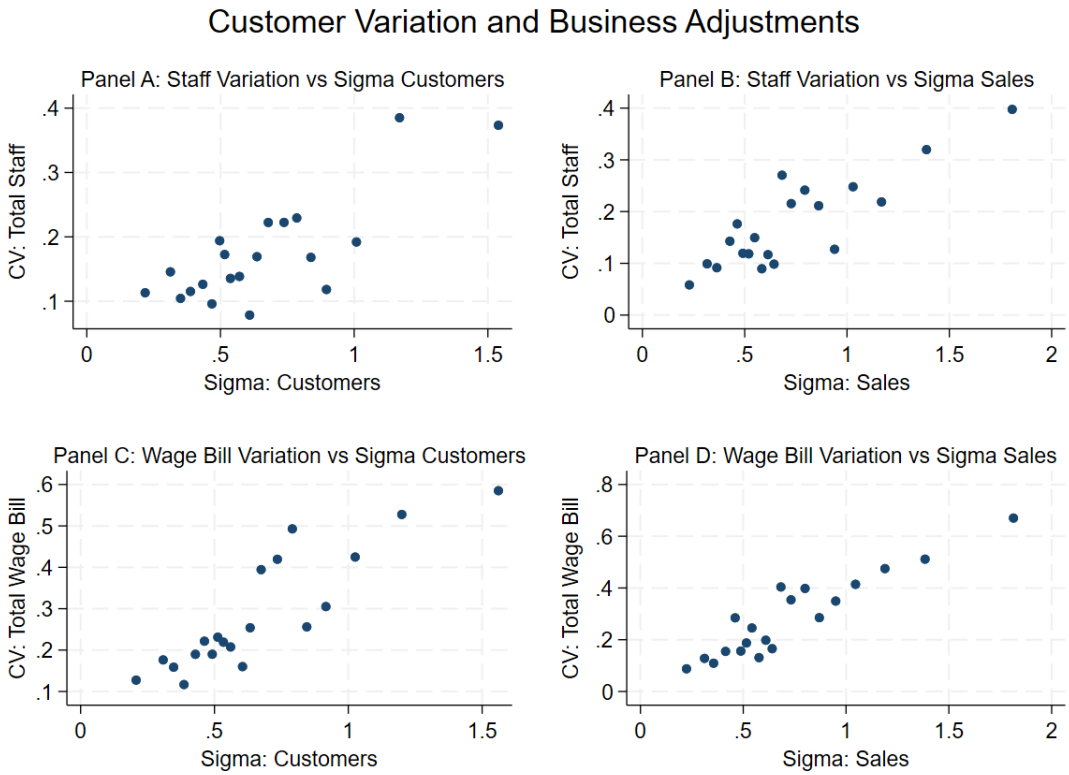


Figure 2: Relationship between customer variation and firm staffing variation

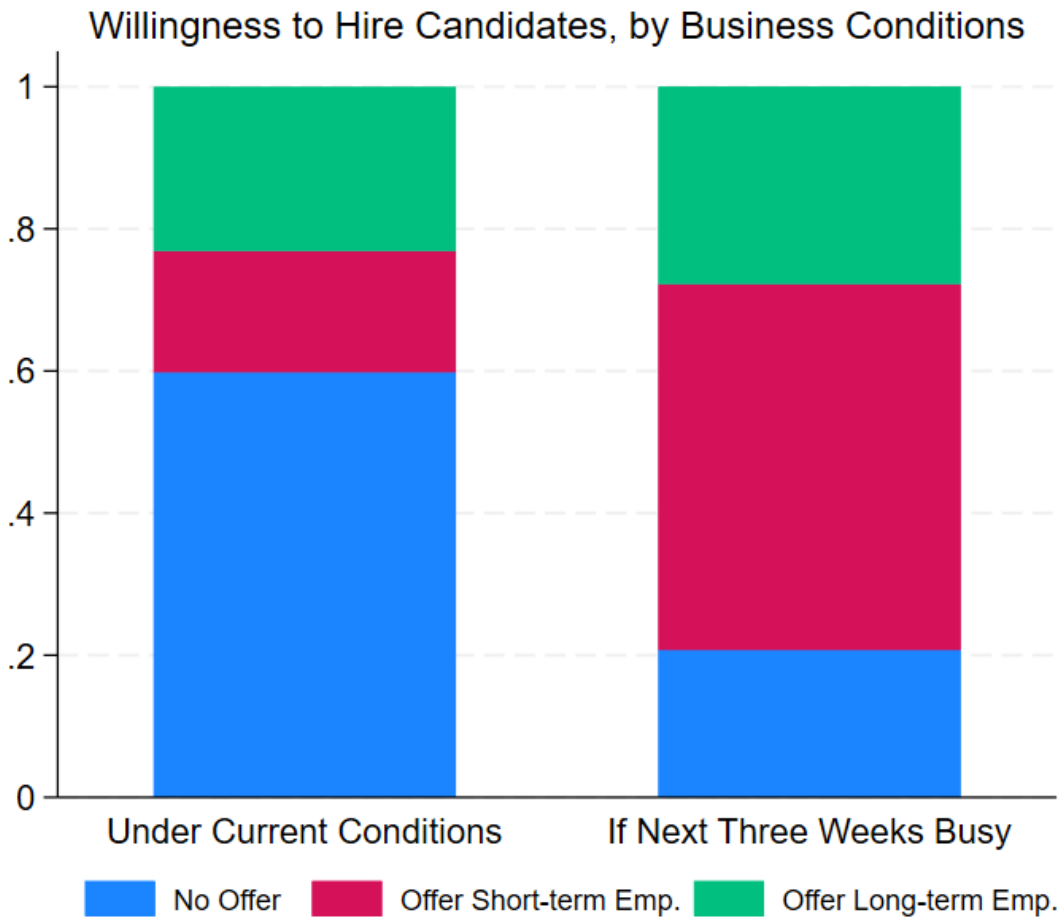


Figure 3: Firm Responses to Demand Shocks: Hypothetical Hiring

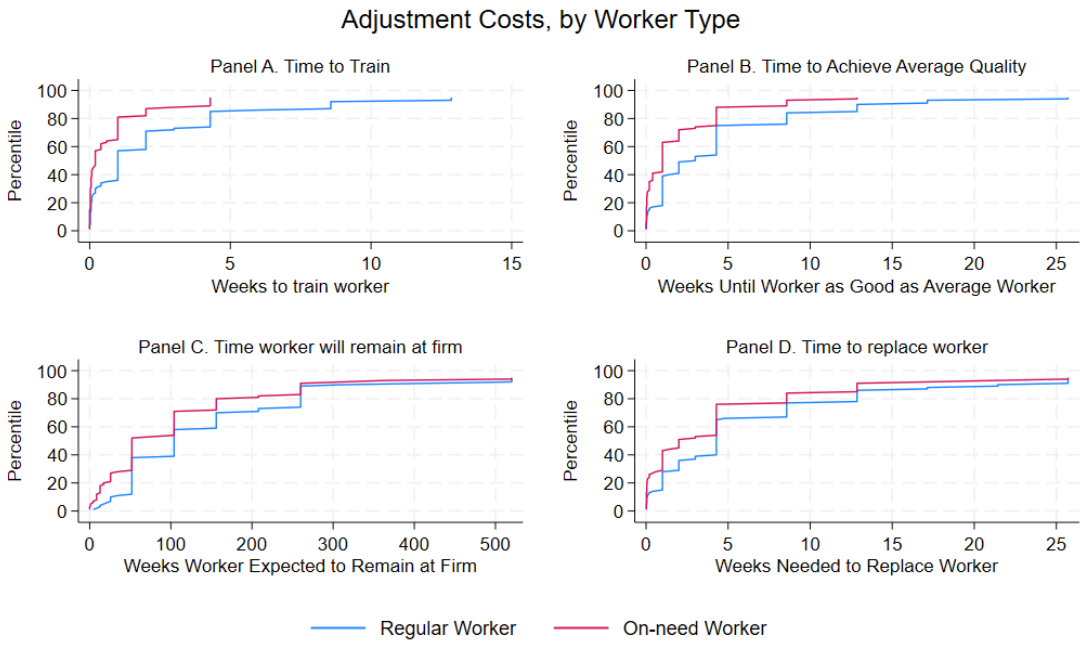


Figure 4: Adjustment Costs by Worker Type

8 Tables

Table 1: 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 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 2: 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 3: 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
σ - Customers	0.290***	0.144**	0.211***	0.075**
	(0.068)	(0.068)	(0.065)	(0.033)
Observations	536	536	536	531
Mean of Outcome	0.46	0.29	0.32	0.16
SD of Outcome	0.499	0.453	0.465	0.214

Notes: Each column is from a regression of firm behavior on the estimated $\hat{\sigma}$ of the firm's distribution of customers. Standard errors are Huber-White heteroskedastic.

Table 4: Variation of Staffing and Wage Bill in Response to Customer Variation

	Panel A: 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
σ - Customers	0.157***	0.342***	0.013	0.042**
	(0.040)	(0.099)	(0.031)	(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
	Panel B: 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
σ - Customers	0.262***	0.347***	0.100**	0.088***
	(0.053)	(0.099)	(0.042)	(0.028)
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

Notes: Each column is from a regression of a coefficient of variation (of staffing or wage bill) on the estimated $\hat{\sigma}$ of the firm's distribution of customers. Standard errors are Huber-White heteroskedastic. The differences in sample sizes reflect don't know or refusal responses from respondents; aggregates are only calculated for cases with all non-missing values.

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	2.278***	2.717***	46.751***	2.302***
	(0.435)	(0.554)	(7.955)	(0.548)
Observations	923	933	873	938
Mean of Outcome	2.58	4.90	140.04	7.09
R-squared	0.054	0.070	0.066	0.070
SD of Outcome	6.72	9.29	145.03	10.88

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 regular (rather than short-term) role. 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
Daily	-0.367***	0.110	0.149**	-0.015
	(0.099)	(0.088)	(0.075)	(0.042)
On-Demand	0.739***	0.973***	0.313***	-0.223***
	(0.138)	(0.101)	(0.082)	(0.042)
Observations	3559	3561	3415	3555
Mean of Outcome	2.67	4.14	4.05	4.89
R-Squared	0.418	0.378	0.514	0.510
SD of Outcome	2.96	2.43	2.38	1.39
F-Test (Daily = On-Need)	68.657	76.192	4.815	21.988
P-Value (Daily = On-Need)	0.00	0.00	0.03	0.00

Notes: Standard errors are clustered at the business level and included in parentheses. Frequency-order, firm and profile fixed effects are included. The omitted category of the dependent variable "frequency" is "regular employee, paid monthly". Outcome (1) corresponds to the number of days the respondent perceived the worker would miss in a month. 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".

A Sampling

Figure A1 presents the enumeration areas we sampled, and which form the basis of our study.

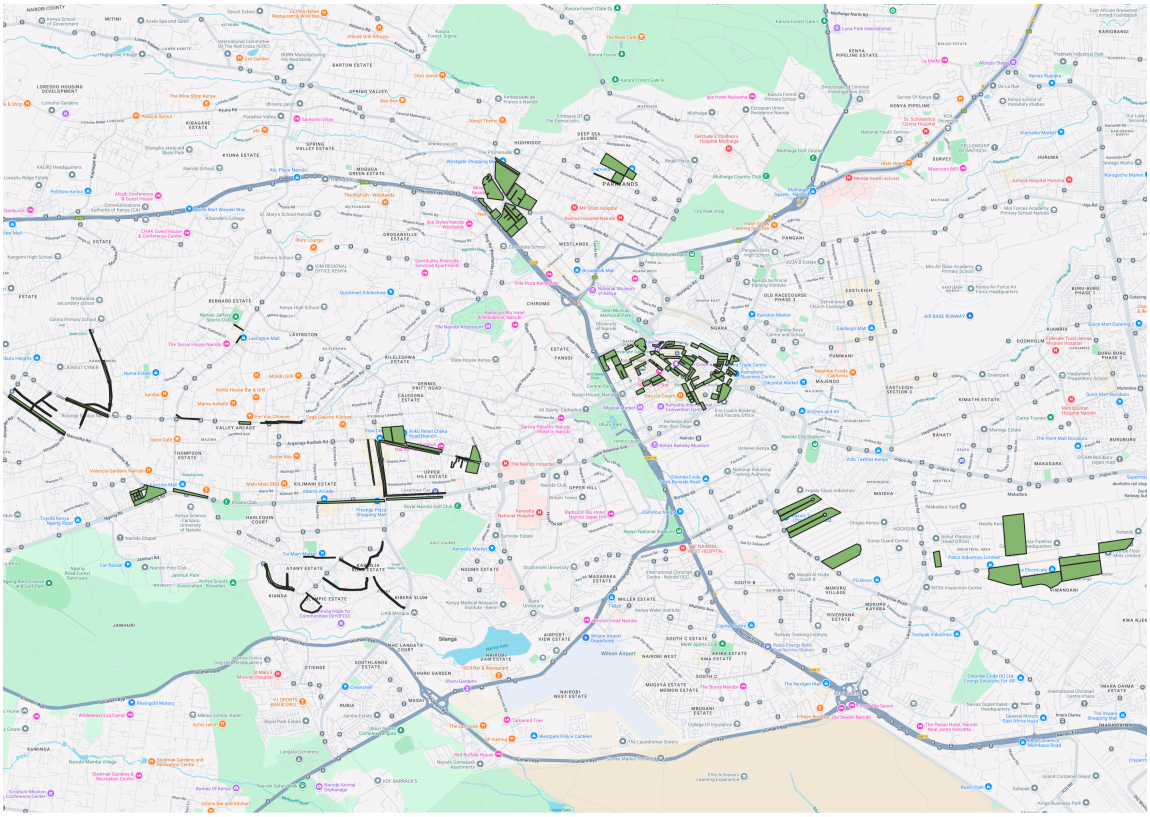


Figure A1: Enumeration areas

B Distribution of Customers and Sales

B.1 Estimation procedure

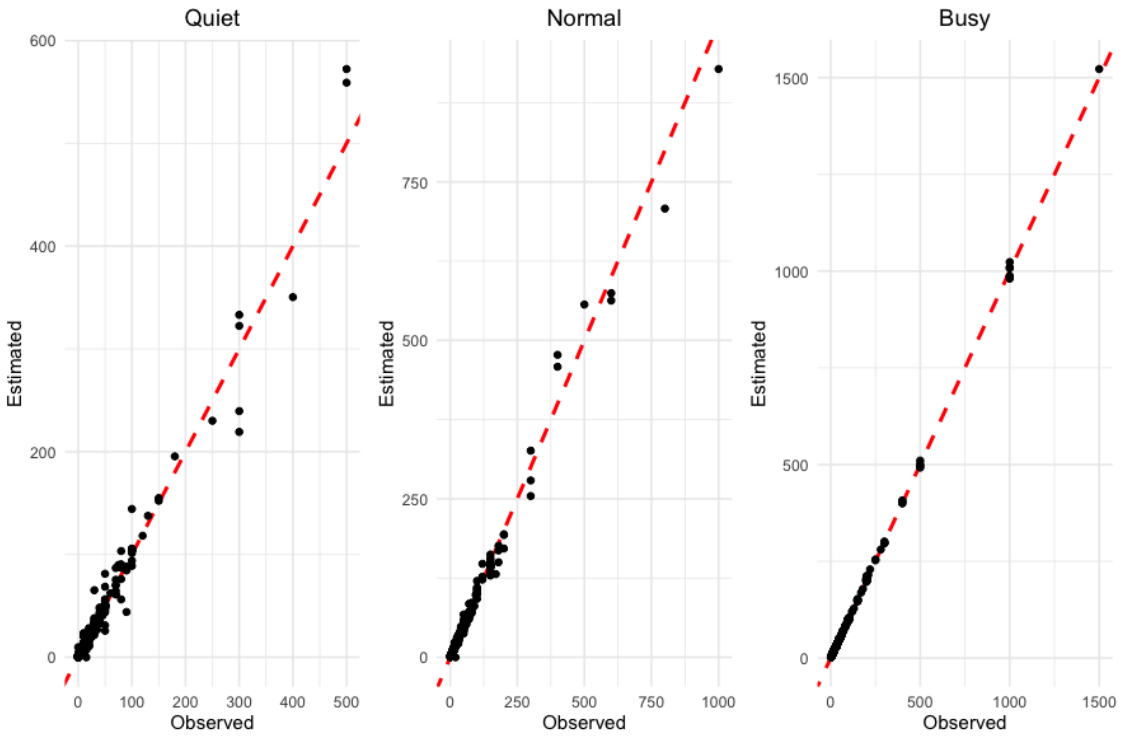
For each firm we obtain the share of days that are quiet, normal, and busy, as well as the corresponding volume of sales and costumers on those days. We convert the shares into percentiles, P, and represent the volume of sales and costumers as the associated quantiles, Q(P).

For a given distribution, F, parameterized by θ, we recover θ through a "method of

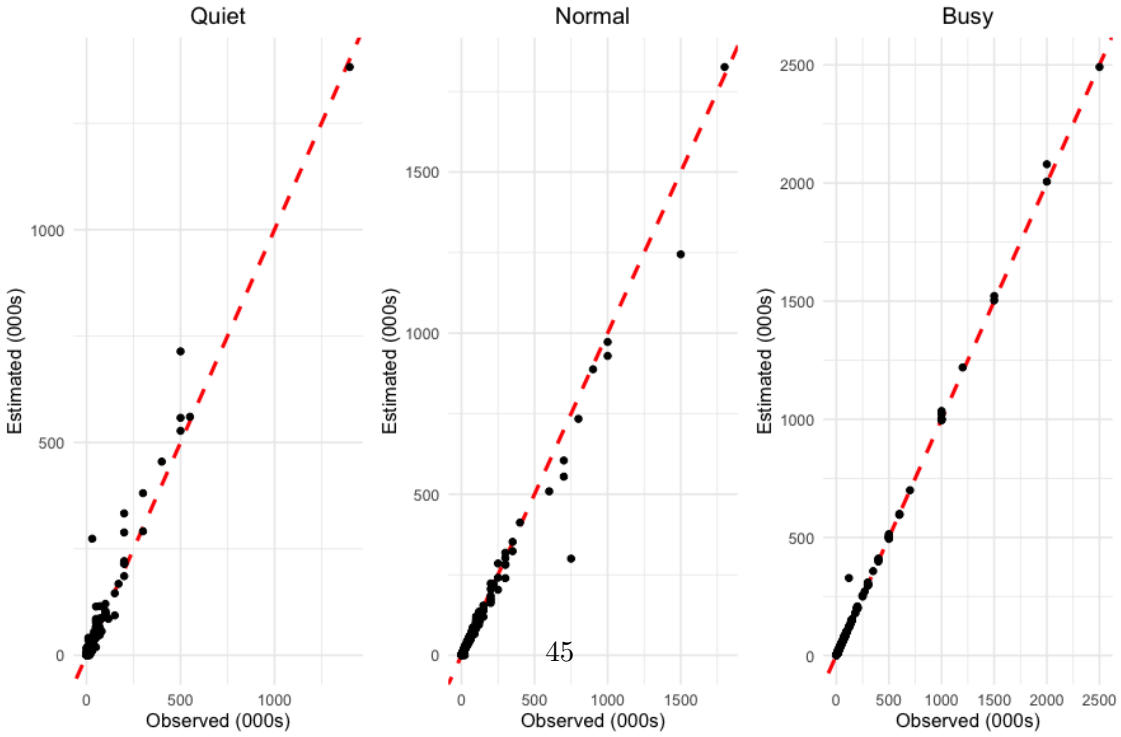
moments" procedure. Formally, we find θ such that:

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

Figure A2 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 A2: Fit of log-normal distribution for customers and sales

B.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. To test the fit of each, we begin by fitting each to two of our three quantiles, and using the last one as a 'hold-out' or test observation. Table A1 summarises how well each distribution fits our data using three measures of fit: (1) mean sum of square residuals (Mean SSE), (2) mean mean of squared residuals (Mean MSE), (3) mean percent deviation (Mean MD). These are averaged across firms for a given hold-out group, and across different hold-out groups. That is, values correspond to the mean measure across observations when using the first two observations to fit a distribution, the second and the third, and the first and the third. The results highlight that both the truncated normal and the log-normal distribution fit the data considerably better than other distributions.

Distribution	Mean MSE (Customers)	Mean MD (Customers)	Mean MSE (Sales)	Mean MD (Sales)
Exponential	4675.38	0.44	15581472269.04	0.54
Log-normal	1280.30	0.12	79655338356.15	0.20
Negative binomial	2199.50	0.33	10938018199.60	1.12
Poisson	2662.61	0.48	10941992335.80	1.12
Truncated Normal	418.35	0.09	2521136320.39	0.13

Table A1: Summary of distribution fit

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

Appendix figure A3 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.

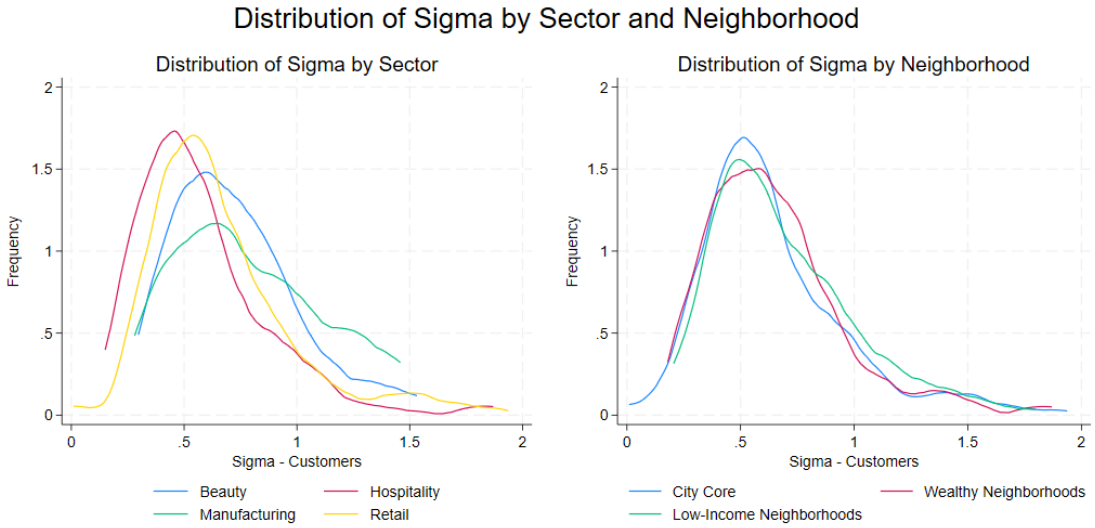


Figure A3: Distribution of $\hat{\sigma}$

C Model proofs

C.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:

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

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

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

C.1.1 A. Proof of existence

Note that as $\eta \to \infty$, $V_{01} \to \infty$ and $V_{02} \to \infty$, 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 function that tends either to 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.

C.1.2 B. 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,

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

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:

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

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

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 the area to the left η_{12}, short-term contracts are preferred, and to the area to the right long-term contracts are preferred.

C.2 Proof of prediction 1

The proof of prediction 1 will proceed by showing that $\frac{\eta_{02}}{\sigma} > 0$ and $\frac{\eta_{12}}{\sigma} > 0$.

C.2.1 A. 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:

$$\frac{\partial \eta_{02}}{\partial \sigma} = \frac{\partial V_0}{\partial \sigma} \cdot \left[\frac{\partial V_0}{\partial \eta} \right]^{-1} \tag{18}$$

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

We've shown above that $\frac{\partial V_2(\eta_{02}, \theta_0)}{\partial \eta} > 0$, we show below in lemma 2.1, that $\frac{\partial V_0}{\partial \sigma} > 0$.

C.2.2 B. Proof $\frac{\partial \eta_{12}}{\partial \sigma} > 0$

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

$$\frac{\partial V_2(\eta_{12},\theta_0)}{\partial \eta}\frac{\partial \eta_{02}}{\partial \sigma} = \frac{\partial V_1(\eta_{01}\theta_0)}{\partial \sigma} + \frac{\partial V_1(\eta_{01}\theta_0)}{\partial \eta}\cdot\frac{\partial \eta_{02}}{\partial \sigma}$$

Re-arranging, we get:

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

$$\frac{\partial \eta_{12}}{\partial \sigma} = \frac{\partial V_1(\eta_{01}\theta_0)}{\partial \sigma}\left[\frac{-\Delta V_{12}(\theta_0,\eta_{12})}{\partial \eta}\right]^{-1}$$

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

We've shown above that $\frac{\Delta V_{12}(\theta_0,\eta_{12})}{\partial \eta} < 0$ everywhere. We show in lemma 2.2, that $\frac{\partial V_1(\eta_{01}\theta_0)}{\partial \sigma} > 0$.

C.2.3 C. Proof of Lemma 2.1 ($\frac{\partial V_0}{\partial \sigma} > 0$)

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:

$$V_1 = \tilde{V}_1(\eta,\theta_0,\sigma) + \frac{\beta q}{1-\beta(1-q)}V_0 \tag{19}$$

Start by focusing on the derivative of $\tilde{V}_1$.

$$\tilde{V}_1 = something + \frac{\beta(1-q)\sigma[-\tilde{w}(1-\Phi(\tilde{w})) + \phi(\tilde{w})]}{1-\beta(1-q)} \tag{20}$$

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

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:

$$\frac{\partial V_1(\theta_0,\eta_i)}{\partial\sigma}=\frac{\beta(1-q)}{1-\beta(1-q)}\left[-\tilde{w}(1-\Phi(\tilde{w}))+\phi(\tilde{w})+\sigma\frac{\tilde{w}}{\sigma}(1-\Phi(\tilde{w}))\right]$$

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:

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

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:

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

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

$$\frac{\partial V_0}{\partial \sigma} = \frac{\beta h}{1-\beta(1-h)} E_{\theta,\eta}\left\{1_{\{V_0 \geq V_1, V_0 \geq V_2\}} \frac{\partial V_0}{\partial \sigma} + 1_{\{V_1 \geq V_0, V_1 \geq V_2\}} \left[\frac{\partial \tilde{V}_1}{\partial \sigma}(\theta,\eta) + \frac{\partial V_0}{\partial \sigma}\right] + 1_{\{V_0 \geq V_1, V_0 \geq V_2\}} \frac{\partial V_2}{\partial \sigma}\right\}$$

Re-shuffling:

$$\frac{\partial V_0}{\partial \sigma} = \frac{\beta h}{(1-\beta(1-h))\ E_{\theta,\eta}\left[1_{\{V_1 \geq V_0, V_1 \geq V_2\}} + 1_{\{V_0 \geq V_1, V_0 \geq V_2\}}\right]} E_{\theta,\eta}\left[1_{\{\tilde{V}_1 \geq V_0, V_1 \geq V_2\}} \frac{\partial \tilde{V}_1}{\partial \sigma}(\theta,\eta)\right]$$

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

C.2.4 D. Lemma 2.2 $\frac{\partial V_1}{\partial \sigma} > 0$

This follows directly from Lemma 2.1 and the definition of V_0.

D Additional Appendix Tables

Table A2: 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
σ - 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	0.498	0.451	0.466	0.214

Notes: Each column is from a regression of firm behavior on the estimated $\hat{\sigma}$ of the firm's distribution of sales. Standard errors are Huber-White heteroskedastic.

Table A3: Variation of Staffing and Wage Bill in Response to Customer Variation

	Panel A: 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
σ - Customers	0.157***	0.342***	0.013	0.042**
	(0.040)	(0.099)	(0.031)	(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
	Panel B: 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
σ - Customers	0.262***	0.347***	0.100**	0.088***
	(0.053)	(0.099)	(0.042)	(0.028)
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

Notes: Each column is from a regression of a coefficient of variation (of staffing or wage bill) on the estimated σ of the firm's distribution of customers. Standard errors are Huber-White heteroskedastic. The differences in sample sizes reflect don't know or refusal responses from respondents; aggregates are only calculated for cases with all non-missing values.

Table A4: Firm Response to Shocks

	(1) Job Offer	(2) On-Need Job Offer	(3) Regular Job Offer	(4) Log Wage	(5) Wage Winsorized	(6) Wage Winsorized, if Offer Was Made
Shock	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: Standard errors are clustered at the business level and included in parentheses. Firm and profiles fixed effects are included. Outcomes (1) to (3) are dummy variables corresponding to 0 if the respondent does not offer the job, and 1 if they do. Outcome (5) is winsorized at the 99% and 1% level.

Table A5: 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	2.278***	2.717***	46.751***	2.302***
	(0.435)	(0.554)	(7.955)	(0.548)
	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	.16	.09	-.1	.05

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.

Table A6: 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 Experience	0.006	0.028	-0.052
	(0.024)	(0.027)	(0.077)
Sector Experience	0.073***	0.139***	0.143*
	(0.023)	(0.026)	(0.077)
High 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 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.