CAUSES AND CONSEQUENCES OF LABOR UTILIZATION: EVIDENCE FROM UGANDAN FIRMS*

VITTORIO BASSI IMRAN RASUL OTTAVIA A.VEROUX ANNA VITALI

JULY 2026

Abstract

Labor idleness is pervasive in small firms in developing countries, yet its causes and consequences remain poorly understood. Using a novel five-year panel of over 1,000 firms in Uganda, we document that firms face substantial and hard-to-predict demand, that capital and labor adjustment costs are large, and that labor utilization is low. Firms with more uncertain demand exhibit lower utilization and invest less in on-the-job training. Guided by these facts, we develop a model of firm input and training choices under demand uncertainty. A key mechanism is that training is a fixed investment paid upfront, while workers are paid piece-rates only when utilized, so higher demand uncertainty lowers expected utilization and reduces the returns to training. We estimate the model, identifying key parameters via a field experiment offering firms a wage subsidy to recruit and train an additional worker. Counterfactual simulations show that reducing demand uncertainty raises hiring and training nearly as much as lowering training costs directly, and that the effectiveness of typical development policies to stimulate firm size, productivity and growth – such as capital grants, management training, and market access interventions – all depend critically on the uncertainty of demand firms face. Our findings highlight a novel demand-side mechanism limiting human capital accumulation, firm size, and the effectiveness of firm-targeted policies in low-income settings.
JEL: D22, J24, O15.

*We gratefully acknowledge financial support from the Mastercard Foundation, PEDL, IGC, J-PAL, the ESRC CPP at IFS (ES/H021221/1), James M. and Cathleen D.Stone Centre on Wealth Concentration, and an anonymous donor. We thank Oriana Bandiera, Stefano Caria, Dennis Egger, Dean Karlan, Franck Portier, Tommaso Porzio, Todd Schoellman, Garima Sharma, Chad Syverson, Chris Udry, Gabriel Ulyssea and numerous seminar participants for comments. IRB approval is from UCL (5115/003, 007). The study is registered (AEARCTR-0000698). All errors are our own. Bassi: USC, vbassi@usc.edu; Rasul: UCL and IFS, i.rasul@ucl.ac.uk; Veroux: IFS, ottavia.veroux@ifs.org.uk; Vitali: NYU, av3906@nyu.edu.

1 Introduction

Nothing is more anemic to economic growth than productive resources being left idle. Labor idleness — the phenomenon of workers being present at the firm but not engaged in productive work — is a pervasive feature of small firms in developing countries. A growing empirical literature documents that utilization rates are strikingly low in Sub-Saharan Africa [Bassi *et al.* 2025, Walker *et al.* 2025]. This stands in sharp contrast to firms operating in high-income countries, where capacity utilization fluctuates over the business cycle but remains high on average [Burnside *et al.* 1993, Burnside and Eichenbaum 1996].

A foundational insight from the macroeconomics literature is that variable utilization of inputs is a key margin of adjustment through which firms respond to shocks: firms hoard factors in downturns, keeping workers and machines partially idle rather than incurring adjustment costs, and then respond to positive shocks by intensifying input use [Greenwood *et al.* 1988, Burnside and Eichenbaum 1996, Basu *et al.* 2006]. In settings where firms face large demand uncertainty and costly input adjustment, input underutilization may arise as a second-best optimum where firms hedge against uncertainty by keeping resources idle [Cooley *et al.* 1995, Fagnart *et al.* 1999, Boehm and Pandalai-Nayar 2022, Amberg *et al.* 2026].

We build on this intuition to study the causes and consequences of labor utilization in small and medium-sized firms in Uganda. We begin by documenting and formalizing a mechanism not studied before in developing country contexts: demand uncertainty combined with adjustment costs leads firms to keep resources idle. We provide direct evidence on this mechanism using high-frequency data from a five-year panel of firms in Uganda.

We then go beyond documenting low labor utilization to trace out its downstream consequences, both for on-the-job training and the effectiveness of typical development interventions to stimulate firm growth. Across much of Sub Saharan Africa, on-the-job training is the primary way in which individuals gain skills post-labor market entry, and it is key to the growth of firms and the economy as a whole [Filmer and Fox 2014, ILO 2023, Alfonsi *et al.* 2020]. To do so, we develop and estimate a model of firm input and on-the-job training choices under uncertainty. To identify the model we exploit a field experiment in which firms were randomly offered wage subsidies to hire and train a new worker. A key mechanism is that training is a fixed investment paid upfront, while workers are paid piece-rates only when utilized, so higher demand uncertainty lowers expected utilization and reduces the returns to training. Our counterfactuals show that reducing demand uncertainty raises hiring and training by nearly as much as lowering training costs directly, and that capital grants, improved market access, and management training are all substantially less effective when demand uncertainty is high.

Our firm panel is based on a representative sample of $1,025$ firms with between 1 and 15 employees (excluding the firm owner) operating in eight manufacturing and services sectors that are central to employment generation in Uganda: motor-mechanics, electrical wiring, welding,

plumbing, construction, hairdressing, tailoring and catering. Our baseline sample covers such enterprises in 15 urban areas in 2012, and then firms were tracked near annually over four follow-ups until 2017.

We collected survey data enabling us to measure and link demand uncertainty, adjustment costs, labor utilization, and firm-sponsored training. We gathered high-frequency information on both realized sales and firm owners' expectations over the distribution of future sales, which we use to construct measures of actual and expected demand volatility. We collected rosters of machines and workers in each firm, to examine how quickly inputs adjust in response to demand shocks. Importantly, we construct a measure of labor utilization using time-on-tasks data. More precisely, for each worker we collected information on a typical product they work on, the time required to complete that product, and the piece-rate payment received for their work on the product. Dividing total monthly earnings by the per-product piece-rate yields an estimate of the number of products completed per month. Multiplying this by the average time spent per product allows us to compute the number of hours effectively worked per month: our key measure of labor utilization. Finally, we collect detailed information on the on-the-job training received by employees.

The firm panel was collected as part of a field experiment where firms were offered various labor market interventions. Some of these treatments were analyzed in previous work [Alfonsi *et al.* 2020, Bandiera *et al.* 2025, Alfonsi *et al.* 2026]. For this paper, we exploit an additional treatment arm in the original design, where firms were randomly offered a six-month wage subsidy to hire and train a new worker of their choice. We use this experimental variation to identify and validate the returns to training in our model of firm behavior.

We begin the analysis by establishing novel stylized facts on the demand uncertainty faced by firms, input adjustment costs, and their implications for labor utilization and training provision.

First, we document that firms face substantial variability in both actual and expected demand — either measured by revenues or customer numbers. Running fixed effects regressions on realized monthly revenues, the majority of variability in revenues cannot be explained by firm, sector $\times$ time or location $\times$ time fixed effects. Our evidence points to demand being not only highly variable but also idiosyncratic and difficult for firms to predict. In contrast, capital and labor inputs display lower and more predictable variability.[1]

Second, we show that adjustment costs in both capital and labor are substantial. More precisely, we exploit our asset and worker rosters to conduct event studies showing firms adjust inputs slowly over time following the hiring or firing of a worker. In addition, we show that in response to demand shocks (proxied by instrumenting revenues with by the mean revenues of other firms in the same location) firms do not adjust capital or labor inputs, and instead adjust rates at which existing labor is utilized within the firm.

Third, we show that labor idleness is pervasive. While the average worker spends 269 hours per

[1]Our survey shows that rental markets are not well developed in our sectors of interest, in contrast with the findings of Bassi *et al.* [2022] for the carpentry sector in the same context.

month at their firm, only 84 of these hours are spent engaged in production, implying a utilization rate in production of 30% for the average firm. Allowing for a utilization rate in non-production tasks of 20% [Bassi *et al.* 2025] brings the estimated overall utilization rate to about 50%, a figure in line with Walker *et al.* [2025] for micro-enterprises in Kenya.

Fourth, training is an essential way in which firms access skilled workers, and training costs are substantial. Half of the employees in our data have received on-the-job training at their current employer. Training takes about nine months and is usually conducted by the firm owner: conditional on having trainees, the average owner spends over 100 hours per month training workers, implying large time costs associated with training provision.

Finally, we bring the analysis together to study the interplay between demand uncertainty, labor utilization and training. Consistent with the mechanism proposed, we find that firms with more variable demand have lower utilization and invest less time training workers. In short, when faced with demand uncertainty, firms keep resources idle, lowering the expected returns to paying the fixed costs of hiring and training workers, thereby reducing investment in training.

Building on these stylized facts and relationships, we develop a two-period model that links the causes of utilization — demand uncertainty and input adjustment costs — to its consequences for underinvestment in labor inputs, and in particular training. The model is in partial equilibrium, and features monopolistically competitive firms making input choices under uncertainty. The model incorporates standard mechanisms that limit on-the-job training — high training costs and the possibility of workers leaving the firm after being trained — along with the novel utilization channel. In the first period of the model, before the realization of an idiosyncratic demand shock, firms set prices and choose capacity by investing in capital, hiring workers, and deciding how much training to provide. In the second period, the demand shock is realized and firms adjust utilization, choosing how intensively to deploy existing inputs, but cannot adjust capital stocks, total employment (firm size), or training due to the presence of adjustment costs.

A key feature of the environment is that capital and training investments are fixed and paid up-front, while workers are compensated on a piece-rate per unit of output and are therefore paid only when utilized. As a result, higher demand uncertainty reduces expected utilization and lowers the returns to fixed training investments. This creates a novel demand-side channel through which uncertainty discourages firms from investing in workers' skills.

We estimate the model using simulated method of moments. We target cross-sectional moments from the baseline and control firm data, exploiting our detailed measures of labor utilization and firm-provided training to identify demand uncertainty and training costs. We use experimental variation from the wage subsidy intervention to identify the returns to training and workers' labor supply responses. To validate the framework, we simulate the wage subsidy intervention in the estimated model and compare the implied treatment effects to those observed in the data. The simulated impacts on employment, training, utilization, and revenues closely track the experimental estimates, providing additional support for the model's ability to capture the key mechanisms

linking demand uncertainty, utilization, and investment in training.

We then use the model to conduct three sets of counterfactuals.

First, we run a horse race between demand uncertainty and three other frictions commonly proposed as drivers of low levels of on-the-job training in low-income countries: training costs, high turnover, and low average levels of demand. To do so, we change the parameters corresponding to these different frictions so as to achieve a 10% increase in sales, and study the corresponding impacts on firm choices of labor input and training. The key result is that reducing demand uncertainty yields an increase in trained workers that is almost as large as that achieved by reducing training costs or separations — typical channels that the earlier literature has focused on.[2]

Second, we show that demand uncertainty shapes the returns to typical policies aimed at fostering firm size, productivity and growth. We simulate the effects of capital grants (reductions in the cost of capital), improved market access (an increase in the level of demand) or management training (an increase in firm TFP if we think of management as a technology [Bloom *et al.* 2017]) across environments with different levels of demand uncertainty, and show that all these policies yield significantly higher returns for firms when demand uncertainty is lower, with labor utilization being an important mediating factor. This strong complementarity between demand stability and policy effectiveness suggests the returns to conventional firm-targeted interventions are substantially mediated by the demand environment in which they are implemented — a dimension largely overlooked in the earlier literature.

The final set of counterfactual exercises come full circle to consider the origins of demand uncertainty — and policies that could feasibly reduce it. We take as our starting point descriptive evidence from our setting that uncertainty is negatively correlated with the number of customers, for example. We use this to consider a counterfactual in which firms could share customers with each other to smooth demand shocks. Such within-sector-location aggregation could be implemented through centralized platforms where customers place orders and these are reallocated across firms based on available capacity [Bergquist *et al.* 2025]. Such arrangements allow firms to effectively share demand while maintaining decentralized production. To quantify the scope for such sharing, we compare the coefficient of variation of revenues at the firm level to that of aggregate sales within sector-location clusters. We show that such demand aggregation leads to a substantial decline in demand uncertainty, consistent with the earlier evidence indicating a large share of firm-level demand fluctuations reflect idiosyncratic shocks.

We then use the model to conduct a counterfactual in which each firm faces the level of demand uncertainty implied by sharing customers in the cluster, while holding constant the mean level of demand and all other parameters. We document this leads to large increases in firm outcomes:

[2]In a similar vein, we show that the effectiveness of wage subsidies to hire and train new workers depend on the level of demand uncertainty. There is a strong complementarity: wage subsidy interventions are more effective when demand is more stable.

firm size increases by 71%, while capital, effective labor, and profits increase by more than 80%. As in most low-income settings, the fundamental friction that limits such demand sharing practices being used is contractual. Reallocating customers across firms requires agreements over quality, compensation, and future customer relationships, which may be difficult to monitor and enforce in settings characterized by high levels of informality and weak institutions. As a result, we quantify the extent to which demand remains fragmented across many small firms, each exposed to substantial idiosyncratic risk — and the consequent impacts on the provision of training and firm outcomes.

Our work contributes to three strands of the literature. First, we contribute to the literature on capacity utilization and its macroeconomic implications. Walker *et al.* [2025] establish that slack is pervasive in low-income economies and has large aggregate consequences for fiscal multipliers and aggregate supply elasticity. We complement their analysis by identifying demand volatility and adjustment costs — distinct from input indivisibilities — as an additional source of underutilization, and by studying its consequences for firm-level investment, firm size and human capital accumulation.

The macro literature on utilization and shock propagation [Greenwood *et al.* 1988, Burnside and Eichenbaum 1996, Basu *et al.* 2006] provides the theoretical foundations for our firm-level framework: the same logic that makes utilization a key margin of aggregate adjustment in rich-country business cycle models operates at the firm level in our setting, but with much larger steady-state slack and qualitatively different consequences. Closest to our paper are two studies examining the relationship between demand uncertainty and utilization in developed-country contexts. Butters [2020] uses data on hotels and airlines to show that demand volatility and adjustment costs generate optimal underutilization and bias standard productivity measures. Amberg et al. [2026] use quarterly survey data on Swedish manufacturing firms to show that firm-level demand uncertainty lowers expected capacity utilization, with a larger effect among high-markup firms, consistent with a newsvendor model of optimal capacity choice. We complement and extend this work in two main ways. First, we study a low-income setting where utilization rates are markedly lower and adjustment costs likely more binding. Second, we trace the consequences of labor slack for on-the-job training investment and for the effectiveness of canonical firm-targeted development policies. In doing so, we also contribute to the broader literature on uncertainty and firm investment [Bloom 2009, Bloom *et al.* 2018] by identifying a distinct utilization-based channel operating in developing-country firms.

Second, we contribute to the literature on policies to improve firm performance and stimulate SME growth in low-income countries. A large experimental literature has evaluated capital grants [De Mel *et al.* 2008, Blattman *et al.* 2014, De Mel *et al.* 2019], management training [Bloom *et al.* 2013, Bloom *et al.* 2014, McKenzie and Woodruff 2014, 2017], wage or training subsidies [De Mel *et al.* 2019, Alfonsi *et al.* 2020, Hardy and McCasland 2023, Brown *et al* 2024], and market access interventions [Atkin *et al.* 2017, Ali *et al.* 2025], often finding that impacts are

positive but modest, and sometimes short-lived. We contribute by documenting a new mechanism that can limit the effectiveness of such interventions: when firms face high demand uncertainty they optimally keep resources underutilized, which reduces the returns to any intervention that increases input availability or quality.

Third, we contribute to the literature on the determinants of firm-sponsored training and, more broadly, human capital accumulation on the job in low-income countries [Acemoglu and Pischke 1998, 1999, Lagakos *et al.* 2018, Caicedo *et al.* 2022, Ma *et al.* 2024, Cefala *et al.*, 2025a, 2025b]. This literature has emphasized training costs, worker turnover, and poaching externalities as the primary frictions limiting on-the-job training. We identify a novel demand-side channel: demand uncertainty reduces expected utilization and therefore lowers the returns to the fixed costs of training, providing a new explanation for why firm-sponsored training is so limited in developing countries and why it responds so little to conventional supply-side interventions. As on-the-job training is partially transferable across firms [Alfonsi *et al.* 2020], the reduced incentives to invest in training due to demand uncertainty have implications that extend well beyond the training firm and can ripple through the labor market.

The paper is organized as follows. Section 2 describes our firm data and the key measurements enabling us to study the causes and consequences of labor utilization. Section 3 establishes stylized facts on the interplay between demand uncertainty, input adjustment costs, labor utilization and firm-provided training. In Section 4 we develop a model of firm input choices that reflects the key stylized facts established and uses this to draw implications for firm-sponsored training of workers and firm outcomes. Section 5 estimates the model, and uses the experiment to validate it. Section 6 presents counterfactual analysis to assess the importance of demand uncertainty for training, the overall effectiveness of firm-targeted interventions in low-income countries, and considers the impacts of policies aiming to reduce demand uncertainty. Section 7 concludes. Additional data description and proofs are in the Appendix.

2 Data and Descriptives

2.1 Sample

Our data was collected as part of a two-sided labor market field experiment studying firms and workers in Uganda [Alfonsi *et al.* 2020]. We drew our sample in 2012, from a census of firms operating in 15 urban labor markets across Uganda and selecting those: (i) having between one and 15 employees (plus an owner); (ii) operating in one of eight sectors, covering both manufacturing and service sectors: motor-mechanics, electrical wiring, welding, plumbing, construction, hairdressing, tailoring and catering. The first restriction excludes micro-entrepreneurs and ensures we focus on small and medium sized firms that have meaningful choices to make over capital and labor inputs. The second restriction ensures we cover the near universe of firms in the same sector-city location

— what we refer to as a cluster of firms. This is valuable in some of the later counterfactual analysis when we explore mechanisms to reduce the demand uncertainty faced by firms.

Our baseline sample comprises 1,025 firms, with the average firm size being three (plus a firm owner). These firms offer good jobs: earnings are higher in these sectors than many others in Uganda, and they collectively employ about 30% of individuals aged 20-30 working outside of agriculture. These firms are thus central to employment generation and the provision of human capital to workers post-labor market entry.

Timeline and Attrition Panel A of Figure A1 shows the study timeline. The baseline survey took place between December 2012 and June 2013. Firms were tracked annually over four follow-ups until between March and July 2017. We consider firms to have attrited if they cannot be found by enumerator teams or are known to have shut down. By the fourth follow-up, 32% of firms interviewed at baseline had attrited. This is in line with an annual death rate of 8.2% documented by McKenzie and Paffhausen [2019] for similar SMEs across 12 developing countries. Table A1 shows correlates of attrition at each follow-up: less profitable, smaller and younger firms are more likely to attrit. Owner characteristics do not predict attrition.

2.2 Data

Four features of our data allow us to study the interplay between labor utilization, demand uncertainty, adjustment costs, and firm-provided training. Panel B of Figure A1 details the timing of data collection related to these features.

First, at the worker level, we record the number of hours each spends on production tasks at the firm through a time-on-task approach. We use this to construct a measure of labor utilization: the number of hours an employee spends on production tasks divided by the number of hours she spends at the firm premises.

Second, at the firm level, we collect high-frequency data on outcomes and inputs. We record profits and revenues for each of three months preceding every survey wave. We combine this with firm owners' expectations of profits and revenues, to provide evidence of (perceived) demand uncertainty *within* the firm over time.

Third, we collect information to construct monthly rosters of assets and employees in each firm. For the asset rosters, we asked owners about the quantity and unit price of pre-specified, sector-specific assets/machines. We record information on all assets firms have access to, regardless of whether they owned, rented, or borrowed the asset. Figure A2 illustrates the asset roster module in four of our study sectors. For the worker rosters, we collected data on each employee's demographics, recruitment process, productivity, hours and earnings. To build a monthly firm panel of assets and employees, in each follow-up survey, we asked owners whether the firm still employed the assets/employees listed in the previous survey, and recorded information on asset

sales/purchases and worker hires/layoffs. This enables us to document input changes over time, and examine whether firms adjust input choices in response to short-term changes in demand. We use this to document the importance of adjustment costs for firms.

Finally, for each worker, we collect detailed information on firm-provided training, including the length of training on-the-job, the monetary cost of training and materials/tools used during training, and who provided the training (the firm owner or other employees). We use this to evidence the costliness of firm-provided training, and how the provision of on-the-job training responds to demand uncertainty.

2.3 Experimental Design

Our data was collected as part of a two-sided field experiment studying urban labor markets in Uganda implemented in partnership with the NGO BRAC. Firms in the study were randomly assigned to either: (i) be matched with vocationally trained or untrained workers, (ii) be given a wage subsidy to hire and train a new worker, or (iii) be held as controls. As indicated on the study timeline in Figure A1, the interventions took place just after the baseline firm survey was completed, between July 2013 and February 2014. For the wage subsidy intervention, 256 firms were randomly chosen to be offered a wage subsidy of UGX120,000 ($50) per month to hire and train a new worker of their choice. This subsidy corresponds to the 75th percentile of the distribution of monthly earnings for unskilled workers, and owners and workers were free to negotiate how to split the subsidy. The subsidy lasted for six months from the time the new employee was hired, conditional on the worker being employed and present on a regular basis (as verified by randomly timed monthly spot checks). The contract between BRAC and firm owners who took up the subsidy specified firms were obliged to train workers, it did not specify a curriculum to follow. This ensures firms provide typical kinds of training to new hires.

Our earlier work using data from this project has focused on the worker side of the experiment to study the impacts of vocational training and matching interventions [Alfonsi *et al.* 2020, Bandiera *et al.* 2025, Alfonsi *et al.* 2026]. In this paper, we exploit the firm-side data and the wage subsidy intervention to target key empirical moments and validate the model developed that links labor utilization, demand uncertainty, adjustment costs, and firm-provided training.

2.4 Context

Firms Table 1 describes our firms at baseline. Panel A focuses on firm characteristics: firms are six years old and therefore closer to their steady state rather than in a start-up phase. Average monthly profits are $218, enabling them to partially self-finance inputs and employee training. Firms are split across manufacturing and service sectors. Panel B focuses on firm owners — they are aged 34, and are evenly split by gender, in part because some service sectors in our sample

are female-dominated. On average firm owners report spending 267 hours at the firm each month. Table A2 presents more information on owners' time use. Panel A shows: (i) 85% of owners report being at the premises every day; (ii) they spend just under 11 hours working on site in a typical day. Panel B details how they spend their time: they allocate time to training, assisting or searching for workers, or searching for machines and loans. Although we did not directly elicit information on finding customers, the typical scenario in our study context is one in which firm owners (not workers) are responsible for finding customers.

Panel C describes capital and labor inputs. The average value of the capital stock accessible to firms is $1256, five times the value of monthly profits and implying a capital stock per worker of over $400.[3]

Panel D provides information on the ties firms have to other firms. Such network ties are defined as firms with whom owners report regularly interacting with or sharing resources with — such as inputs, information, credit, or advice. These networks can help firms cope with demand uncertainty, for example through sharing customer orders. Networks are widespread: 79% of firms report at least one network connection, and are on average tied to one other firm. However, only 37% of these ties are purely professional, while the remainder involve ties through family or friends of the owner. Conditional on having a network, machines and credit are the most commonly shared resources (23% and 20% of firms, respectively), whereas only 9% and 10% of firms share workers or information about demand. This information on network ties among all firms in the same sector-city location cluster is valuable in the later counterfactual analysis when we explore mechanisms to reduce the demand uncertainty faced by firms.[4]

Workers Table 2 uses worker rosters to describe employees at firms. We do so for control firms at first follow-up in order to present data on training that is not impacted by our experimental intervention. Panel A shows that nearly all employees are full-time. The modal payment scheme is for workers to be paid a piece-rate per unit of output (not per hour at work), with 49% of workers being solely paid according to piece rates and 67% having compensation schemes that include some fixed payment or payment-in-kind.

Panel B details tenure and hours. Workers enjoy long tenure within firms — 19 months on average across our study sectors. This highlights these are long-term and meaningful employment relationships. Workers spend 11 hours at the firm premises each day, for 25 days a month, so spend 280 hours at the firm premises each month, closely matching the 291 hours owners report spending at the firm premises. This emphasizes workers have limited time to work for other employers.

[3]To demonstrate the representative of our sample, Table A3 compares our firms to those in the 2010 Global Entrepreneurship Monitor (GEM) Survey for Uganda from 2010. The two samples are balanced on most of the comparable characteristics of owners and firms, although our sample is likely to be slightly positively selected given they had to have at least one employee to be sampled.

[4]The limited sharing of productive assets across firms contrasts with evidence from Bassi *et al.* [2022], who document a large rental market among carpentry firms in Uganda — a sector not included in our study.

Training As in most labor markets across sub-Saharan Africa, there are two main ways in which workers acquire human capital post labor market entry: through attending a formal vocational training institute, or being trained on-the-job. Panel C shows 50% of workers employed by firms in our sample have received on-the-job training, while 27% have received some form of vocational training from a vocational training institute before joining the firm. Hence, in this setting, firm-provided training is key to human capital accumulation post labor market entry. Moreover, these skills are not entirely firm-specific but transferable across firms: 63% of firm owners report that the on-the-job training they provide is useful both in this firm as well as in other firms. This emphasizes the provision of training is key to the wider post-labor market entry accumulation of human capital in the economy as a whole.

Panel D highlights the time costs of training: on-the-job training lasts almost ten months on average. Training is nearly always delivered by firm owners (sometimes in combination with co-workers). Firm owners spend around 113 hours on training — so around 40% of their time at the firm premises. Firm owners therefore face an opportunity cost of time from training employees. Panel E evidences the monetary costs of training and who pays these costs: the majority of trainees are unpaid, 29% of trainees pay owners for their training, and 20% of trainees are paid. There are also training costs related to the purchase of tools or materials for trainees: Panel F shows these are relevant in over 80% of trainings, and paid for by owners 56% of the time.

All this evidence points to on-the-job training constituting a large fixed investment by firms into their employees. Central to our analysis is the link between the *ex ante* incentives to train workers to the rate at which workers are expected to be utilized in the firm.

Table 3 digs deeper on the nature of firm-provided training, splitting workers into those that have and have not been trained within the firm. Those with firm-training have longer tenure ($p = .000$), and higher earnings ($p = .000$). We have previously documented the high returns to on-the-job training in these sectors [Alfonsi *et al.* 2020]. This earnings premium is not due to trained workers working more hours ($p = .486$). Rather, as Panel B shows, they are more productive: owners assess firm-trained workers to be more productive than untrained workers, and more likely to be able to perform a typical sector-specific task. That this is due to them being trained in the firm is underlined by the fact that when they were originally hired, firm owners perceived these workers to have lower productivity and be less able to conduct a spector-specific task than other workers.

Taken together, this evidence indicates workers require firm-provided skills to be productive. Training enables workers to perform more complex tasks rather than simply becoming faster at basic ones. Consistent with this, Panel C shows that trained workers spend significantly more time producing the typical product/service they undertake, relative to the time spent by untrained workers on their typical product/service. As a result, they produce far fewer items than untrained workers each month (because they work the same number of total hours). The products/services that firm-trained workers produce are higher valued items than those produced by non-trained

workers, as reflected by the higher prices and value of materials used to produce them.

3 Stylized Facts

We establish novel stylized facts related to labor utilization, demand volatility, adjustment costs, and firm-provided training. Our objective is to make precise important features of our setting and use these to develop a framework to understand the causes and consequences of labor utilization for the provision of on-the-job training and firm outcomes.

3.1 Labor Utilization

We first document labor utilization across firms. We implement a time-on-task measurement approach. To do so, we exploit the worker rosters to define the labor utilization of worker i employed at firm j in the last month as:

$$\text{Labor utilization}_{ij} = \frac{\text{number of hours spent by } i \text{ engaged in production tasks at firm } j}{\text{number of hours spent by } i \text{ at firm } j \text{ premises}}. \quad (1)$$

The numerator is constructed using information on: (i) the piece-rate the worker received to complete a pre-identified sector-specific product/service; (ii) the time it typically takes them to complete that sector-specific product/service; (iii) the workers' total monthly earnings in the month prior to the survey. We calculate the number of products/services the worker completed in the previous month by dividing total earnings by the piece-rate. We then derive the number of monthly hours spent in production as the number of products/services the worker produced times the typical product/service-specific time taken by them to complete the task. We thus measure labor utilization for workers that: (i) are paid a piece rate, and, (ii) engage in the pre-identified sector-specific product/service. This information is consistently available from the second follow-up onwards. The denominator is constructed using information on: (i) the number of hours the worker usually spends on the firm premises per day, (ii) the number of days the worker spends on the premises per week. The total number of hours spent on the premises in a given month is then calculated as the product of these two values, multiplied by four.[5]

[5]For each worker paid piece-rates (67% of workers), we ask them to identify the sector-specific product or service the worker performs most often. We then asked the following questions regarding the product or service identified by the worker: (i) *How much time (in minutes) would <name> normally take right now to complete his/her work on this product/service alone to a standard of quality considered satisfactory by the average customer of this business?*; (ii) *How much of the price charged to customers for this product/service does <name> normally receive?* The implicit assumption behind the construction of the labor utilization measure is that the piece-rate on other tasks is proportional to the amount of time the worker spends on those other tasks. We can also construct the labor utilization measure at first follow-up but with slightly more imprecision. Specifically, at first follow-up, we pre-specified a task when we asked firms about the time it would take a worker to complete it, but not when asking about the piece-rate paid to workers for the task. This implies that the tasks to which these two questions refer might be different. Starting from the second follow-up, we asked both questions about the same task.

For control firms, Panel A of Figure 1 plots the distributions of the number of hours that workers spend engaged in production tasks in the month prior to baseline (light blue), and the hours spent at the firm premises over the same month (dark blue). Panel B then shows the distribution of the labor utilization rate. The underutilization or *idleness* of labor is pervasive: while the average worker spends 269 hours per month at their firm, only 84 of these hours are spent engaged in production, a labor utilization rate of 31%. The distribution is right skewed: only 1% of workers have a utilization rate above .9, and 22% have utilization rates below .1.

Figure A3 plots labor utilization rates for workers by sector, demonstrating the phenomena is not driven by specific sectors: labor utilization rates vary from .16 in plumbing to .57 in catering. Two points are of note. First, there is substantial variation in utilization rates across workers in the same sector, who conduct similar tasks — emphasizing the idiosyncratic nature of labor utilization. Second, the highest utilization rates are in catering — where 5% of workers actually have a utilization rate of one. This aligns with an intuition driving our analysis, that firms facing lower demand uncertainty optimally utilize workers at a higher rate.[6]

Interpretation and Validation Four points are of note. First, our time-on-task approach focuses respondents on a single concrete task and its associated piece-rate and time, eliciting information workers know precisely from their daily work. An obvious drawback, however, is that the approach does not capture time spent on non-production tasks, such as interacting with customers or procuring inputs. Using data from the same setting and similar sectors, Bassi *et al.* [2025] document that employees spend 20% of their time on such non-production tasks. Assuming such tasks are conducted to the same extent in our study sectors, the average labor utilization rate rises to 51%. This would be in line with the level of utilization documented by: (i) the Uganda Manufacturing Association and the Uganda Small Scale Industry Association [UMA 2023], which finds an average utilization rate of 54% in a random sample of 205 manufacturing firms; (ii) Walker *et al.* [2025] in micro-enterprises in Kenya.[7]

Second, given our measure of labor utilization is constructed at the worker level, we can examine how individual characteristics correlate to utilization rates. These results are in Table

[6]Table A4 details the firms and workers among controls that our measure of labor utilization can be constructed for. Column 1 considers characteristics of firms in which the labor utilization measure can be constructed for at least one worker, conditional on sector. We can construct our labor utilization measure for 60% of firm-survey observations. As expected, firms with more unskilled workers are less likely to be used to construct this outcome, because unskilled workers are less likely to be paid piece rates. The next two Columns repeat the analysis based on worker characteristics. Column 2 shows that we are can measure labor utilization for 48% of workers, and are more likely to do so for skilled or higher productivity workers. Column 3 shows this conclusion is unchanged conditional on firm fixed effects.

[7]However, our measure of the time workers spend on production tasks is substantially lower than the 60% reported by Bassi *et al.* [2025] for the carpentry sector. To reconcile this difference, we note that our sample includes firms across eight sectors spanning both manufacturing and service activities. On average, our data shows that utilization rates are almost 8 percentage points higher in manufacturing than service sectors. In addition, UMA [2023] shows that furniture, which is the focus of Bassi *et al.* [2025] and is not part of our study, is the manufacturing sector with the highest utilization rate (89%) among the 19 sectors from which they collect data.

A4. In Column 4 we control for worker characteristics, and fixed effects for survey waves, location and sector. Worker characteristics are not predictive of utilization rates — the only statistically significant predictor is tenure within the firm, but the magnitude of the partial correlation is small. In Column 5 we control for survey wave and firm fixed effects to narrow in on within-firm variation in labor utilization: we see a very similar pattern of results. Specifically we do not see that within a firm, skilled, more educated or longer tenure workers have higher utilization rates. Nor do we find evidence of family members employed by firm owners being utilized less. This suggests the idleness of labor does not reflect using the offer of work as a form of redistribution whereby family members are employed but left unused [Macchi and Stalder 2024, Swanson 2025].

Third, we find limited evidence that when not utilized, workers spend time at other firms. As described earlier, only 9% of firms report sharing workers with other firms, and 95% of workers are employed full-time. The fact that workers remain idle rather than re-employed in other firms is natural given the sporadic flow of customer arrivals within a day, implying workers could be called upon at unpredictable times. In contrast, we saw earlier that firm owners report little idle time at work (Table A2).

3.2 Demand Uncertainty

If firms are unable to easily adjust inputs in response to demand shocks, underutilization or idleness may arise as a second best optimum [Cooley *et al.* 1995, Fagnart *et al.* 1999, Boehm and Pandalai-Nayar 2022]. To gauge the demand uncertainty firms face, we exploit the fact that at first follow-up, we asked owners' their expectations about future revenues and customer numbers. Panel A of Table 4 describes these expectations of revenues for the next six months, and expectations of customer numbers in the next week. We ask: (i) the *minimum* expectation for each outcome; (ii) the *maximum* expectation for each outcome; (iii) the number of months they expect the outcome to be above the midpoint of (i) and (ii). Assuming a triangular distribution of expectations, we then calculate the expected number, standard deviation and coefficient of variation of expectations for each outcome.

Owners face uncertainty along both margins of revenue and customer numbers. On revenue over the next six months, the lowest expectation is 30% of the highest, with a coefficient of variation of .24. On customer numbers over the next week, the lowest expectation is 33% of the highest. Panel B shows that these figures closely match the variation in realized sales firms experience in the three months prior to the survey.[8]

While this evidence is consistent with owners facing considerable demand uncertainty, part of the variation could be driven by predictable trends that firms are able to anticipate, such as seasonality. To isolate the residual uncertainty, we exploit our panel to regress actual monthly

[8]This degree of uncertainty is not specific to our setting: Walker *et al.* [2025] find that among Kenyan micro-enterprises, sales in their worst week are 37% of those in their best week.

revenues on location-by-month, sector-by-month, and firm fixed effects. Figure 2 shows the share of the variance of firm monthly revenues that is explained by each of these fixed effects. First, we normalize the total variance to 1 (gray bar). Second, we regress revenues on a specific set of fixed effects and plot the remaining variance (net of these fixed effects) as a share of the initial variance. This indicates how much of the variance is unexplained by various fixed effects.

The first set of bars show that location-by-month and sector by-month-fixed effects — which might capture location- and sector-specific seasonality — only explain 15% and 32% of the variance in revenues respectively. Firm fixed effects play a large role, explaining 52% of the total variance. However, after including all fixed effects in the same regression, 42% of the total variance remains unexplained, suggesting firms face substantial idiosyncratic demand uncertainty.

The second and third sets of bars in the same figure present analogous results for realized capital and labor inputs. For the value of the capital stock, the vast majority of variation is explained by time-invariant heterogeneity across firms (97%) and less than 3% remains unexplained. Similarly, variation in employee numbers is largely explained by time-invariant heterogeneity across firms (71%), and only 25% of the variation remains unexplained.

Taken together this evidence underpins the notion that demand is a key source of uncertainty for firms. It also suggests inputs are more stable and predictable than demand, consistent with the presence of input adjustment costs as we probe in the next section.

3.2.1 Origins of Demand Uncertainty

To shed light on potential drivers of demand uncertainty, we examine correlates of the coefficient of variation of firms' monthly revenues. We do so in two ways: Panel A of Figure 3 exploits variation *across* firms in a given survey wave, while Panel B uses variation *within* firms over time. We focus on three sets of correlates: (i) firm and owner characteristics, (ii) characteristics of firms' business networks; (iii) perceived constraints on the firm.

Focusing first on the variation across firms in Panel A, firm size and the number of customers emerge as key predictors of demand uncertainty. A 10% increase in firm size and in the number of customers is associated with a reduction in the coefficient of variation of monthly sales of .002 and .004, respectively (.65% and 1.3% of the mean). The negative relationship with firm size may reflect the inability of smaller firms to access large-scale customers (e.g. through procurement or export markets), which typically provides smoother, more predictable demand over time. In addition, firms with a limited customer base are likely to be more exposed to idiosyncratic shocks affecting individual customers, such as shocks to household income, health or mobility. This may be particularly important in our setting, where 40% of customers are drawn from the same neighborhood and 83% from the same city, making firms vulnerable to correlated local shocks across households.

Managerial ability of owners is also negatively correlated with demand volatility: more skilled

owners can be more effective at forecasting and planning for demand fluctuations, as well as acquiring and retaining customers.[9]

Firms operating in Kampala, relative to other urban areas, experience less volatility, likely reflecting higher and more stable foot traffic. Firms that report borrowing do not experience lower volatility, suggesting access to credit is not a key driver and may partly reflect reverse causality — if firms facing more uncertain demand are more likely to borrow.

Turning to business networks, the number of network links is uncorrelated with uncertainty. However, breaking these into personal and professional ties we see that only the latter appear to smooth revenues, perhaps because shocks are less correlated between business-ties than within family-tied firms. Among network ties, worker sharing is most strongly associated with lower uncertainty, followed by sharing information about demand — although as reported in Table 2, these practices are not very common. In contrast, access to machines or credit through the network do not lower uncertainty.[10]

Finally, firms that report greater difficulty in finding customers exhibit higher demand uncertainty, in line with the results on customer numbers. Constraints related to upstream suppliers and managerial constraints are uncorrelated with uncertainty.

Panel B repeats the analysis within firms over time. As such, covariates that do not have variation over time become less precisely estimated. It remains the case that firms with higher customer numbers, a greater share of firm links that are professional ties, and where owners report sharing information about demand with their network, exhibit lower revenue volatility.

Overall, over and above household shocks to customer incomes or mobility, these findings point to limited access to customers and incomplete risk-sharing through networks as potentially key origins of demand uncertainty among the kinds of small and medium sized firms we study. We come back to consider these ideas later in the analysis when we examine feasible counterfactual policies that might enable firms to reduce the demand uncertainty they face.

3.3 Adjustment Costs

For demand uncertainty to lead to lower labor utilization, it must be costly for firms to adjust inputs in response to demand shocks. We provide evidence of the nature of adjustment costs.

[9]Managerial ability is measured using an index of managerial practices constructed from a battery of questions on record keeping, marketing, stock control, business performance review, and financial planning.

[10]The negative coefficients on worker sharing and demand-information sharing in Figure 3 are notable because reverse causality would push these relationships in the opposite direction. For example, if firms borrow workers to absorb realized demand fluctuations, this would generate a *positive* correlation between demand volatility and worker sharing. The observed negative correlations are therefore conservative, and suggest that network interactions actively reduce the demand uncertainty firms face in the first place.

3.3.1 Sources of Adjustment Cost

Capital To understand firm's ability to adjust capital, it is useful to consider the kinds of machine used in our study sectors. Panel A of Figure 4 shows the distribution of machine costs across firms. To compare across sectors, we normalize machine costs by sectoral averages. The value of most machines lies within a standard deviation of the sectoral average. An asset like lift jacks are a relatively low unit cost machine in the motor mechanics sector; a designing machine is a relatively high unit cost machine in the welding sector.

Panel B details characteristics by quartile of normalized machine costs: (i) while many firms own assets in the bottom quartile of machine values, fewer than 10% own machines in the top quartile; (ii) access to the most expensive machines (through ownership or rental) is more strongly associated with firm profits, perhaps because they allow mechanization/scale-up of activities [Bassi *et al.* 2022]; (iii) over 90% of assets are financed through savings, regardless of their value.

Panel C shows firm owners' reports at baseline of perceived constraints to expanding capital, the most important being difficulties in finding reliable machines (72%), and high borrowing costs (65%). Both constraints likely bind because of the wide variation in machine values. Time costs to firm owners are also relevant: (i) 38% of owners search for machines, spending an average of 23 hours on each search; (ii) 6% of owners search for sources of funding to purchase machines, spending an average of 8 hours on each search.

Labor Panel D repeats the analysis for perceived constraints to expanding the number of workers: finding workers with the right skills (68%), trustworthiness (57%), and being able to assess the quality of applicants (51%) are all key constraints. Fewer owners perceive access to unskilled labor and managerial constraints as key hiring frictions (29% and 21% respectively). Overall, high search and screening costs constrain firms adjusting labor in response to demand shocks [Cooper *et al.* 2006]. They thus face the alternative of training and/or utilizing workers in-house in response to demand shocks.

3.3.2 Existence of Adjustment Costs

Move beyond perceptions of owners, we use two strategies to dig deeper into the existence of adjustment costs based on actual input choices. Our first approach considers how firms' capital, labor, and labor utilization respond to exogenous changes in demand measured through firm revenues. We do so using a Hausman-style instrument, so instrumenting for firm revenues with the average revenues of firm's in the same sector-location-month, excluding the firm's own revenues [Hausman *et al.* 1994, Hausman 1996]. The identifying assumption is that there is sectoral-spacial correlation in demand shocks across firms, but that a given firm does not adjust its inputs in response to the demand shocks faced by other firms. Some potential sources of these correlated demand shocks were highlighted in Figure 3. Using our firm panel, we regress inputs on instrumented

firm revenue, conditional on firm fixed effects — so removing the time-invariant heterogeneity that explains much of the variation in revenues across firms, and only focusing on the within-firm variation in labor utilization across workers.

These IV results are in Table 5. In line with the presence of adjustment costs, the elasticity of labor and capital to demand shocks is very low: .004 (Column 1) and −.003 (Column 2) respectively, and neither being statistically different from zero.

To segway towards understanding whether demand shocks translate into changes in labor utilization, we repeat the analysis for firms in which the labor utilization measure is constructed. Column 3 suggests these firms may face lower capital adjustment costs, as their assets positively respond to demand shocks, with an elasticity of .15. Column 4 confirms that labor inputs do not adjust to revenues.

Column 5 then examines whether the average labor utilization across workers in the firm responds to exogenous shifts in demand, using the average labor utilization within firms because utilization rates do not vary across workers in the same firm (Table A4). In sharp contrast to assets or employee numbers, firms significantly increase labor utilization rates in response to positive demand shocks, with an estimated elasticity of .830, that we cannot reject is different from one ($p = .643$). In short, firms respond to demand shocks not by changing input levels, but by changing labor utilization. This is confirmed in Columns 6 and 7 that show labor utilization responds to demand shocks through changes in time workers spend in production (the numerator in (1), not the time workers spend at the premises.[11]

Our second approach uses an event study design to show how labor, capital, and the value of capital per worker adjust in response to: (i) the hiring/separation of a worker; (ii) the purchase/sale of a machine. We use the monthly panel data on the number of workers and the value of assets utilized by firms i in month t and survey wave s to estimate the following event study specification:

$$y_{it} = \alpha_i + \gamma_t + \sum_{m=-10, m\neq 0}^{10} \beta_m z_{i,t+m} + \varepsilon_{it}. \tag{2}$$

For each survey wave we identify the event as the hiring/separation of a worker, or the purchase/sale of a machine. z_{it} is an indicator equal to one for firms that change inputs in periods after the first input change. α_i and γ_t are firm and month fixed effects respectively.

Figure 5 shows the event studies related to changes in labor. Panel A shows how inputs change around a worker separation: (i) firms do not hire other workers to quickly replace separating workers; (ii) the value of the firm's capital stock changes smoothly around the event — suggesting

[11] We can also construct average labor utilization rates separately for skilled and unskilled workers. Doing so we find that positive demand shocks only translate into increased labor utilization of skilled workers: the estimated elasticity from the IV specification is 1.01 for skilled workers ($p = .016$) and −.314 for unskilled workers ($p = .699$). However, we need to caveat this by reiterating that there are few unskilled workers for whom labor utilization rates can be constructed because most workers are skilled, and unskilled workers are more likely to be unpaid or be paid a fixed wage rather than a piece-rate.

the separation is unrelated to market conditions limiting firm's capital accumulation; (iii) as a result, the capital per worker jumps upwards around the event. Panel B shows symmetric impacts of a new worker hire: capital adjusts slowly and the value of capital per worker never recovers to its pre-hire level. These findings do not vary with the reason for separation (Figure A4). We also document similar dynamics around the time of changes in capital stocks (Figure A5).[12]

The key takeaway from these event studies is that it takes considerable time for firms to adjust a given input around the event of a sudden change in the level of the other input, even when that change can be anticipated. Along with the evidence in Table 5, this all points to firms facing substantial input adjustment costs.

3.4 Bringing the Key Elements Together

We now bring the key elements of our analysis together to examine how labor utilization and on-the-job training correlate with demand uncertainty.

3.4.1 Labor Utilization and Demand Uncertainty

In Column 1 of Table 6 we regress labor utilization on our measure of demand uncertainty: the coefficient of variation of firms' *expected* monthly sales as measured at first follow-up, controlling for sector and location fixed effects. We find a precisely estimated negative partial correlation between demand uncertainty and labor utilization: those firms facing greater demand uncertainty utilize their workers at lower rates.

We next exploit the panel data to regress labor utilization on the coefficient of variation of firms' *current* monthly revenues, controlling for firm fixed effects. To establish a causal relationship, we again use a Hausman-style instrument, so instrumenting the coefficient of variation of firms' current monthly revenues with the average coefficient of variation of monthly revenues among firms in the same sector, location and survey wave, excluding the firm itself. The IV result in Column 2 shows a negative relationship between demand volatility and labor utilization. The magnitude of the coefficient implies that a one standard deviation increase in the coefficient of variation of monthly revenues leads to an 18% decrease in labor utilization.

In short, the negative relationship between demand volatility and labor utilization holds both across (Column 1) and within firms (Column 2).

[12]Figure A4 shows results from event studies where the worker left voluntarily or was fired. The slight difference is that in the latter case we see more evidence of firms replacing the worker around seven months after their departure. This might be because firms start searching for a new employee before firing an old one, in anticipation of the time required to find a replacement. Figure A5 shows event studies related to changes in capital. Panel A shows the event of a sale of a machine. Such events lead only to a small fall in the value of the capital stock, suggesting most firms only sell small-valued machinery — that they are more likely to own and that is easier to replace. Panel B repeats the analysis for the purchase of a machine. For this event we see a gradual increase in labor around the event but again because the value of these purchases is small, the value of machines per worker falls and does not recover back to its pre-event level.

We next break down labor utilization by the number of hours spent in production (Columns 3 and 4) and the number of hours spent at the firm premises (Columns 5 and 6). Both across and within firms, the effects on utilization are driven by changes in the hours workers spend on production tasks rather than by adjustments in hours spent at the premises.

Overall, the results are consistent with firms facing higher demand uncertainty maintaining idle resources to be able to absorb positive demand shocks [Cooley *et al.* 1995, Fagnart *et al.* 1999, Butters 2020]. In contrast, firms facing more stable demand are better able to predict their input needs and have a lower option value of keeping inputs underutilized.

3.4.2 On-the-job Training and Demand Uncertainty

The final batch of results explore impacts on a key firm outcome: the provision of on-the-job training, as measured by the number of hours the owner spends training workers in the month prior to the survey. Column 1 of Table 7 shows that, at baseline, firm owners facing higher uncertainty about future revenues spend significantly fewer hours training workers. Column 2 focuses on within-firm variation, using the same IV for revenue volatility described above. It shows that, when faced with an exogenous increase in uncertainty, owners decrease the time spent training workers. The magnitude of the coefficient implies that a one standard deviation increase in the coefficient of variation of monthly revenues leads to a 12% decrease in the number of hours the owner spends training. Columns 3 and 4 show this result is not driven by changes in the time owners spend at the premises.[13]

The results are consistent with the returns to training being lower when firms face greater demand uncertainty. This is because firms pay a fixed cost to train employees — a cost that is independent of their subsequent utilization. However, as shown earlier, workers are also utilized less intensively when uncertainty rises, reducing the returns from on-the-job training.

4 Model

We now develop a model of firm behavior incorporating the established stylized facts: that demand uncertainty and adjustment costs can impact labor utilization rates. We focus on the consequences these features have for the provision of on-the-job training and other firm outcomes.

4.1 Set-up

Agents The economy features a monopolistically competitive sector, with a continuum of homogeneous risk-neutral firms $j \in J$ producing differentiated varieties, and homogeneous risk-neutral workers $i \in I$.

[13] If demand uncertainty also led to an increase in the owner's idle time, we would expect a positive within-firm relationship between demand uncertainty and training, as owners would have more time available to train workers.

Demand The model is in partial equilibrium, so we do not explicitly model final consumption. Instead, we assume each firm j faces (exogenous) demand $p_j^{-\theta} v_j$, where v_j is an idiosyncratic demand shock drawn from $F(v)$ with (known) support $v \in [\underline{v}, \bar{v}]$, and $0 < \underline{v} < 1 < \bar{v}$, and θ is the absolute value of the price elasticity of demand. Let μ_v and σ_v denote the mean and standard deviation of v_j, respectively.

Timing The model has two periods:

Period 1: Both firms and workers face uncertainty about demand. Firms choose the number of workers b_j (firm size), capital k_j, training τ_j, and output price p_j. Workers choose the amount of time to spend at the firm's premises h_j.

Period 2: The idiosyncratic demand shock v_j the firm faces is realized. k_j, b_j, τ_j and h_j are fixed, but firms can adjust utilization (and therefore output) by choosing the share of workers' time allocated to production l_j.

This two-period setup provides a simple way to embed the type of adjustment costs observed empirically, suggesting it takes time for firms to find and train new workers and to purchase new capital assets. The model effectively assumes infinite adjustment costs between period 1 and 2.

Productive Capacity In period 1, firms invest in capacity by choosing the number of workers, capital, training, and their output price. Inputs become productive only at the beginning of period 2 (because it takes time for capital to be installed and workers to be trained). The firm's productive capacity, q_j, is given by a constant-returns-to-scale Cobb-Douglas technology:

$$q_j = A_j k_j^{\alpha} (h_j b_j \phi(\tau_j))^{1-\alpha}, \tag{3}$$

where the labor input is the composite of three elements: b_j — the number of workers hired by the firm; h_j — the share of time firm j's workers spend on the premises; and $\phi(\tau)$ – the training-augmentation factor, assumed to be increasing and concave ($\phi' > 0$, $\phi'' < 0$).

Utilization Rate In period 2, the idiosyncratic demand shock v_j is realized. At this stage inputs are fixed, but the firm can adjust utilization by choosing utilized labor $l_j \in [0, b_j]$. We can interpret $\frac{l_j}{b_j}$ as the share of scheduled time a worker actually spends in production, namely the *utilization rate* for labor. Following Fagnart *et al.* [1999], we assume that output with utilization l_j is proportional to the utilization rate, namely:

$$y_j = q_j \frac{l_j}{b_j} = A_j k_j^{\alpha} (h_j \phi(\tau_j))^{1-\alpha} b_j^{-\alpha} l_j \tag{4}$$

Three observations follow. First, this functional form implies that all inputs are utilized in the same proportion. In other words, there is no separate utilization choice for capital and labor.

Second, the marginal product of utilized labor, $\partial y_j/\partial l_j = q_j/b_j$, is increasing in capital (k_j) and training (τ_j), reflecting complementarity in the technology, as well as in workers' time on the premises (h_j). Third, $\partial y_j/\partial l_j$ decreases in the number of workers b_j, which implies that although adding workers raises capacity, it dilutes the payoff from additional utilization. This could reflect coordination costs that increase with firm size, whereby larger firms require greater communication and agreement across workers [Becker and Murphy 1992, Bassi *et al.* 2024]. These costs are likely to be larger in low-demand states, when workers must decide not only how to allocate tasks but also who works and who does not, particularly in an environment where workers are paid only when performing tasks.

Realized Output Given a realization of the demand shock v_j in period 2, the firm can adjust labor utilization to meet its production needs. This is because, as shown below, profit maximizing prices exceed marginal cost, so it is optimal for the firm to produce as much as possible subject to demand and capacity. The realized output of the firm in period 2 is therefore:

$$y_j(p_j, v_j) = \begin{cases} p_j^{-\theta} v_j & \text{if } v_j \leq \tilde{v}_j \\ q_j & \text{otherwise.} \end{cases} \tag{5}$$

There exists a cutoff demand shock, $\tilde{v}_j = q_j p_j^{\theta}$, at which the firm's demand equals its capacity. Firms meet demand at low v_j (by adjusting utilization), but are capacity constrained for $v_j \geq \tilde{v}_j$.

4.2 Solving the Model

We solve the model by backward induction.

4.2.1 Period 2: Labor Utilization

Given the realized period 2 demand shock v_j, the firm chooses utilized labor to meet its production needs. Using equation (4), optimal utilized labor is therefore:

$$l_j = \frac{y_j}{A_j k_j^{\alpha} (h_j \phi(\tau_j))^{1-\alpha} b_j^{-\alpha}}, \qquad 0 \leq l_j \leq b_j. \tag{6}$$

4.2.2 Period 1: Prices

Firms choose prices in period 1 before observing the demand shock v_j. They do so by maximizing expected variable profits:

$$\max_{p_j} E_v \Big[(p_j - MC_j) y_j(p_j, v_j) \Big] \tag{7}$$

where $MC_j = w h_j / (A_j k_j^{\alpha} (h_j \phi(\tau_j))^{1-\alpha} b_j^{-\alpha})$ denotes the firm's marginal cost. This is equal to the cost of one unit of utilized labor ($w h_j$) divided by its marginal product ($\partial y_j/\partial l_j$). Two

points are worth noting. First, firms only pay workers when they are utilized, consistent with the majority of employees in our setting being paid piece-rates. That is, wages are a variable cost. Second, capital and training costs do not enter marginal costs, as they are fixed costs incurred in the first period.

Using the capacity cutoff $\tilde{v}_j = q_j p_j^{\theta}$, expected output can be written as:

$$E_v[y_j(p_j, v_j)] = q_j \underbrace{\left[\frac{1}{\tilde{v}_j}\int_{\underline{v}}^{\tilde{v}_j} v\, dF(v) + 1 - F(\tilde{v}_j)\right]}_{\text{expected utilization}} \tag{8}$$

This expression shows that expected output is equal to the product of capacity (q_j) and expected utilization. Specifically, the first term in brackets captures the contribution of slack states to expected utilization. In these states, $v_j < \tilde{v}_j$, output is demand-determined and utilization is $v_j/\tilde{v}_j$. The second term is the probability that the firm is capacity constrained, in which case utilization equals one.

The first-order condition for prices yields:

$$p_j = \underbrace{\frac{\theta}{\theta - 1 - \frac{1-F(\tilde{v}_j)}{\frac{1}{\tilde{v}_j}\int_{\underline{v}}^{\tilde{v}_j} v\, dF(v)}}}_{\text{markup } \mu(\tilde{v}_j)} \times \underbrace{\frac{w}{A_j k_j^{\alpha} \phi(\tau_j)^{1-\alpha}(h_j b_j)^{-\alpha}}}_{\text{marginal cost } MC_j}. \tag{9}$$

The optimal price is a markup over marginal cost. Notice that the markup depends on the two terms that affect expected utilization. Specifically, this expression shows that markups are increasing in the relative importance of capacity-constrained to slack states, $(1 - F(\bar{v}_j))/\frac{1}{\tilde{v}_j}\int_{\underline{v}}^{\tilde{v}_j} v dF(v)$. Intuitively, when low-demand states are more likely (relative to capacity-constrained states), lowering prices expands expected sales in the region where the firm is not capacity constrained, making price cuts more valuable and leading to a lower optimal markup. Conversely, when the firm expects to be frequently capacity constrained, lowering prices has a weaker effect on expected output because the firm is already constrained in high-demand states. Instead, the firm can sell capacity output at a higher price, resulting in a higher optimal markup.

4.2.3 Period 1: Labor, Training, and Capital

Firms choose labor b_j, training τ_j, and capital k_j in period 1 before observing v_j, to maximize expected profits:

$$\max_{k_j, b_j, \tau_j} E_v\Big[(p_j - MC_j) y_j(p_j, v_j)\Big] - \chi_j(\tau_j) b_j - r k_j \tag{10}$$

where $\chi_j(\cdot)$ is the per-worker training cost, assumed increasing and convex ($\chi_j' > 0$, $\chi_j'' > 0$) and r is the capital rental cost. A key distinction is that, while wages are variable costs paid only when workers are utilized (through l_j and therefore MC_j), training and capital are fixed costs, incurred

in period 1 regardless of the subsequent utilization of inputs.

Let $E_v(u) = \frac{1}{\tilde{v}_j} \int_{\underline{v}}^{\tilde{v}_j} v dF(v) + \big(1 - F(\tilde{v}_j)\big)$ be the firm's expected utilization. Using the expression for expected output derived above, the input problem can be written as:

$$\max_{k_j, b_j, \tau_j} \Big(p_j - MC_j\Big) q_j E_v(u) - \chi_j(\tau_j) b_j - r k_j \tag{11}$$

This expression makes clear why utilization matters for fixed investments: training and capital are paid for before demand is realized, but their returns depend on the extent to which the resulting productive capacity is used.

Training Assuming all workers receive the same amount of training within the firm, the first order condition with respect to training τ_j is:

$$(1-\alpha)\frac{\phi'(\tau_j)}{\phi(\tau_j)} MC_j q_j \Big(E_v(u) + (\mu(\tilde{v}) - 1)(1 - F(\tilde{v}))\Big) = \chi_j'(\tau_j) b_j \tag{12}$$

The left-hand side is the marginal benefit from training, which is determined by three components. The term outside brackets captures the productivity gain from a marginal increase in training evaluated at full capacity. A higher training-augmentation factor ($\phi(\tau_j)$) raises the marginal productivity of utilized labor, lowering marginal costs in all states, and additionally increasing capacity in capacity-constrained states.

The bracketed term determines the expected profit return to this productivity gain. The first term in brackets is expected utilization: the more often the firm expects its labor to be utilized, the larger the cost savings from training. The second term captures the additional markup premium from capacity-constrained states. In these states, training allows the firm to sell extra units, and the profit gain from those units depends on the price-cost margin ($\mu(\tilde{v}_j) - 1$). The higher the markup and the more often the firm is capacity constrained, the larger the additional profit return to the additional capacity generated by training.[14] The right-hand side is the marginal cost of training all employees ($\chi'(\tau_j) b_j$), which is a fixed cost per worker paid regardless of subsequent utilization.

Capital The first order condition with respect to capital k_j is:

$$\frac{\alpha}{k_j} MC_j q_j \Big(E_v(u) + (\mu(\tilde{v}) - 1)(1 - F(\tilde{v}))\Big) = r \tag{13}$$

[14]Note that optimal markups also increase in response to higher training. Training lowers marginal cost and, holding the markup fixed, would therefore reduce prices. Because demand is elastic ($\theta > 1$), this price decline expands demand by more than training expands capacity, lowering the cutoff ($\tilde{v}_j = q_j p_j^{\theta}$). As a result, $1 - F(\tilde{v}_j)$, the probability that the firm is capacity constrained, increases. By Equation (9), this raises the optimal markup. However, the increase in markups only partially offsets the decline in marginal cost, so the overall effect of training is to lower prices.

This expression is analogous to the FOC for training. The term outside brackets captures the productivity gain from a marginal increase in capital evaluated at full capacity. The bracketed term weights this gain by expected utilization and the profit premium from additional capacity in constrained states. The right-hand side is the marginal cost of capital, the rental rate (r), paid up-front and independently of utilization.

Labor Finally, the first order condition with respect to labor b_j is:

$$\frac{1}{b_j} MC_j q_j \Big(-\alpha E_v(u) + (1-\alpha)(\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}_j)) \Big) = \chi(\tau_j) \tag{14}$$

The FOC for labor highlights the trade-off firms face when hiring an additional worker. The term outside bracket $\frac{q_j}{b_j}$ is output per utilized worker, which multiplied by the marginal cost yields the full-utilization value (in cost-equivalent terms) of an additional fully utilized worker.

The bracketed terms separates two effects. The first term (negative) captures the marginal cost effect of hiring an additional worker. Adding workers lowers output per utilized worker through the coordination/dilution effect in production, therefore increasing marginal costs. The higher expected utilization is, the larger the productivity loss from this channel. This is the key element distinguishing labor from training and capital. The second term (positive) captures the capacity expansion effect of hiring. The gain from these additional units is proportional to the price-cost margin and is weighted by the probability that the firm is capacity constrained.

The firm hires workers up to the point where the expected net marginal benefit equals the per-worker training cost ($\chi_j(\tau_j)$). Notably, this is the *only* cost incurred unconditionally when hiring an additional worker: because wages are paid as piece-rates and only when workers are utilized, they enter marginal costs rather than affecting hiring decisions.

4.2.4 Period 1: Worker Labor Supply

Workers choose their labor supply to firm j, h_j, before observing the firm's realized demand shock v_j. They do so to maximize their expected earnings net of the disutility from time at the firm premises. This disutility captures the value of leisure or the worker's outside options, such as leaving to work for another employer. Given our setting, for each unit of time spent at the firm, we assume workers are only paid for the fraction of time they are utilized. Given the earlier results, we take this as constant across workers and equal to the firm's utilization rate. Workers thus solve the following maximization problem:

$$\max_{h_j} E_v \left[w h_j \frac{l_j}{b_j} - \gamma(h_j) \right], \tag{15}$$

where $\gamma(h_j)$, with $\gamma'(h_j) > 0$, capture the disutility from time spent at the firm. The first-order condition with respect to h_j yields:

$$wE_v(u) = \gamma'(h_j) \tag{16}$$

As is intuitive, at the optimum, h_j increases with the wage w, decreases with the marginal disutility of time spent at the firm γ, and increases in expected utilization because it raises the marginal value of time spent at the firm relative to the worker's outside option.

4.3 Comparative Statics wrt Demand Uncertainty σ_v

We present comparative statics for expected utilization, inputs choices, training and labor supply with respect to demand uncertainty (σ_v).

Proposition 1: Utilization. *Hold inputs (k, b, τ, h) and the expected value of the demand shock $E(v)$ constant.*

1. ***Fixed prices.*** *Holding prices fixed, expected utilization is weakly decreasing in a mean-preserving increase in demand uncertainty: $\frac{dE_v(u)}{d\sigma_v} < 0$.*
2. ***Re-optimized prices.*** *If firms re-optimize prices, expected utilization is weakly decreasing in a mean-preserving increase in demand uncertainty if and only if markups do not fall too fast: $\frac{dE_v(u)}{d\sigma_v} < 0$ if and only if $\frac{d\log\mu}{d\sigma_v} \geq \frac{1}{\theta}\frac{\partial E_v(u)/\partial\sigma_v}{\frac{1}{\tilde{v}}\int_{\underline{v}}^{\tilde{v}} v dF(v)}$.*

All proofs are in the Appendix. The intuition for the result is as follows and also illustrated in Panel A of Figure A6. Holding inputs and prices fixed, the firm's input choices yield a capacity cut-off $\tilde{v}$. For demand realizations below the cutoff, the firm meets demand and utilization is $u(v) = v/\tilde{v}$. At the cutoff, the firm reaches capacity, and utilization is capped at one. Therefore, utilization is an increasing but concave function of the demand shock: $u(v) = \min\left\{\frac{v}{\tilde{v}}, 1\right\}$. A mean-preserving increase in demand uncertainty σ_v shifts probability mass toward both low- and high-demand states. The low-demand realizations reduce utilization because the firm is further below capacity. The high-demand realizations, however, cannot raise utilization symmetrically, because utilization is capped at one once the firm is capacity constrained. Therefore, by concavity, expected utilization falls.

When markups — and therefore prices — are allowed to adjust, the cutoff itself may shift. In particular, a decrease in prices following a rise in volatility may move the cutoff $\tilde{v}$ to the left, generating an increase in utilization that offsets the decrease generated by the change in the distribution of demand shocks alone. For the direct distributional effect to dominate, markups cannot fall too much in response to higher volatility. Notice that this effect is mediated by θ: the higher the demand elasticity, the more the cutoff will fall in response to a decrease in prices, and the larger the countervailing effect on utilization.[15]

[15] An increase in volatility can cause prices to rise or fall depending on how it affects the relative importance of

Proposition 2: Training and Inputs. *Let the firm re-optimize inputs and prices* (k, b, τ, p) *when* σ_v *changes, holding labor supply* h *and the expected value of the demand shock* $E[v]$ *fixed. Then inputs are weakly decreasing in a mean-preserving increase in demand uncertainty if the following conditions hold:*

1. ***Training***: $\frac{d}{d\sigma_v}\Big(E_v(u) + (\mu(\tilde{v}) - 1)(1 - F(\tilde{v}))\Big) \leq 0$;
2. ***Capital***: $\frac{d}{d\sigma_v}\Big[b_j\Big(E_v(u) + (\mu(\tilde{v}) - 1)(1 - F(\tilde{v}))\Big)\Big] \leq 0$;
3. ***Number of workers***: $\frac{d}{d\sigma_v}\Big[- \alpha E_v(u) + (1 - \alpha)(\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}))\Big] \leq \frac{1}{wh_j}\frac{d}{d\sigma_v}\chi_j(\tau_j)$.

The intuition is as follows. The training FOC shows that training raises profits through two channels: it lowers the marginal cost of utilized labor and it expands sellable capacity in constrained states. The relative importance of these two channels are summarized by the term in brackets. A mean-preserving increase in demand uncertainty decreases expected utilization $E_v(u)$ under the conditions stated in Proposition 1, so the derivative of the first term with respect to σ_v is negative. Intuitively, when higher demand uncertainty lowers expected utilization, the productivity gains from training apply to fewer units of output, reducing the return to paying the fixed cost of training.

At the same time, an increase in demand uncertainty changes the probability mass above the capacity cutoff. This affects the second term in brackets through two margins: the probability that the firm is capacity constrained and the markup earned on additional units sold. A sufficient condition for training to fall is that the derivative of this term with respect to σ_v is also negative.

In the Appendix, we discuss the conditions under which this sufficient condition holds assuming fixed prices and v distributed as a log-normal. The intuition is shown in Panel A of Figure A6. We show that, as long as the cutoff $\tilde{v}$ is not too high, the tail probability $(1 - F(\tilde{v}))$ — and therefore training training — falls with a mean-preserving increase in demand uncertainty. To assess of whether the $\tilde{v}$ falls below the relevant threshold in our empirical setting, we back out the share of firms that should be constrained if $\tilde{v}$ were at the threshold above which training may increase in demand uncertainty. In our estimation, $\sigma_v = 3.4$, implying that the threshold for $\tilde{v}$ corresponds to the 99.97th percentile. This implies that firms fall in the relevant region as long as they are not so unconstrained that capacity binds in fewer than 0.034% of demand states. Consistent with this threshold condition, Panel B of Figure 1 shows that almost 6% of firms are operating at capacity at follow-up, which is substantially higher than 0.034%. This suggests that our firms are plausibly located in the region where a mean-preserving increase in demand uncertainty reduces the return to training.

slack and capacity-constrained states, as shown in Equation (9). In particular, markups increase when volatility raises the capacity-constrained term $(1 - F(\tilde{v}, \sigma_v))$ relative to the slack-state utilization term $(\frac{1}{\tilde{v}}\int_{\underline{v}}^{\tilde{v}} v(\tilde{v}, \sigma_v))$, and decrease when the opposite occurs. A sufficient condition for expected utilization to fall is that markups weakly increase in response to volatility: $\frac{d \log \mu}{d\sigma_v} \geq 0$.

The intuition for capital is analogous to the one for training, but with one important difference. As with training, the marginal benefit of capital depends on the two terms in the capital FOC: the utilization term and the markup-weighted capacity-state term. However, while the cost of training is paid per worker and therefore scales with b, the cost of capital does not. As a result, changes in the number of workers increase the return to capital relative to its cost. For capital to fall in response to higher demand volatility, the decline in the combined term $(E_v(u)+(\mu(\tilde{v})-1)(1-F(\tilde{v})))$ must be sufficiently large to offset any increase in the workforce.

The comparative static for the number of workers (firm size) is even more ambiguous because an increase in b has two opposing effects. On the one hand, adding workers lowers productivity per utilized worker and raises marginal costs. This effect is captured by the negative term $-\alpha E_v(u)$. The key difference with respect to training and capital is that lower utilization weakens this negative force: when higher demand uncertainty lowers expected utilization, additional workers are less likely to be fully utilized, so the marginal-cost increase generated by each worker is smaller. This force can push firm size upward.

On the other hand, adding workers also expands capacity. This benefit arises only in capacity-constrained states and is captured by the second term in brackets. If higher demand uncertainty lowers the probability of firms being capacity constrained, this capacity-expansion benefit falls, pushing firm size downward. The right-hand side is a function of the per worker training cost. Since training is re-optimized when uncertainty changes, this cost can also change with σ_v: if firms reduce training in response to higher uncertainty, the cost of hiring an additional worker falls, which can partly offset the decline in the marginal benefit of labor.

Thus, higher demand uncertainty has an ambiguous effect on employment: lower utilization can reduce the marginal-cost penalty from hiring, a lower probability of being capacity constrained can reduce the benefit of additional capacity, and changes in training costs affect the fixed cost of adding workers.

Proposition 3: Worker Labor Supply. *Suppose* $\frac{d\log\mu}{d\sigma_v} \geq \frac{1}{\theta}\frac{\partial E_v(u)/\partial\sigma_v}{\frac{1}{\tilde{v}}\int_{\underline{v}}^{\tilde{v}} v dF(v)}$. *Then* $\frac{d}{d\sigma_v}h_j < 0$.

To see the intuition for this, equation (16) shows that the time workers spend at the firm premises increases in utilization. By Proposition 1, under endogenous prices expected utilization decreases with demand volatility when markups do not decrease too much in response to higher demand volatility. Hence, worker's labor supply at the firm decreases with demand volatility under the same conditions.

5 Model Estimation

5.1 Method

We jointly estimate the model parameters using simulated method of moments and indirect inference. The returns and the cost of training are parameterized as $\phi \log(1+\tau)$ and $\chi\tau^2$, respectively. We assume the demand shock follows a log-normal distribution and normalize μ_v to 0. We then estimate the following nine parameters, $\{\alpha, \theta, \chi, \phi, \sigma_v, \gamma_1, \gamma_2, A, w\}$, by matching the nine data moments listed in Table 8. While there is no one-to-one mapping between an individual moment and a specific parameter, Table 8 indicates which moments are more closely related to the identification of each parameter targeted in the estimation. Below, we discuss several of these moments, for which the unique features of our data are particularly important.

Cross-section Moments The first moment is the average utilization rate. Specifically, we target the observed share of time workers spend in production in our data, combined with the share of time devoted to non-production tasks in data from Bassi *et al.* [2025]. As discussed in Proposition 1, holding inputs and mean demand fixed, utilization is pinned down by demand volatility. This moment therefore plays a central role in identifying, σ_v.

The second moment is the share of workers receiving on-the-job training, which helps identify the average costs of training. The rental rate of capital is calibrated externally using estimates from Bassi *et al.* [2022], while wages and the output elasticity of labor in the production function (α) are estimated by targeting the labor share of total costs and average firm size. Although the estimated wage may appear high relative to the rental rate of capital, it is paid only when workers are utilized. Expressed on a per-hour basis over the total time workers spend at the firm, the implied wage would therefore be substantially lower.

The average number of machines per firm helps identify total factor productivity (A), through the effect of productivity on firm scale. Finally, we calculate markups as the ratio of monthly profits to total costs, where total costs are measured as revenues minus profits. This moment is particularly informative about the demand elasticity parameter θ.

Experimental Moments As described in Section 2, our data was collected as part of a field experiment in which firms were randomly assigned the offer a monthly wage subsidy for six months to hire and train a new worker of their choice. We use the field experiment to target two additional parameters. The first (ϕ) captures the returns to training. Because understanding how underutilization affects firms' incentives to provide on-the-job training is central to our analysis, accurately estimating the returns to training is particularly important. The wage subsidy experiment generates an exogenous shock to the cost of hiring and training a new worker. We estimate the returns to training by targeting the change in sales induced by hiring a new treated worker,

where the hiring of such a worker is instrumented by treatment assignment.[16] Similarly, to estimate non-linearities in workers' labor supply responses to utilization (γ_2), we target the change in the time incumbent workers spend on the firm premises in response to the wage subsidy.

Table 8 summarizes the estimated parameter values and reports both the observed and simulated moments. Overall, we achieve a good fit between the model and data moments.

5.2 Using the Experiment to Validate the Model

We validate the model using the field experiment. More precisely, we simulate the wage subsidy intervention in the model and compare the implied treatment effects on training provision and other firm outcomes with those observed in the data. The simulation follows the timing structure of the model while imposing the constraints of the experimental design. Consistent with the intervention, the wage subsidy is modeled as a cash transfer provided to the firm conditional on hiring an *additional* unskilled worker.

In the first period (before the realization of the demand shock v_j), firms choose all inputs that are predetermined at the time of production. In particular, they choose capital k_j and training for the new worker τ_j^{treat}. As we did not design a specific training curriculum, we assume that firms can freely choose how much training to provide to the new worker. However, they cannot change other labor inputs — namely, the number of incumbent workers b_j, which was a condition for receiving the cash transfers, and the amount of training τ_j provided to them, which we assume is sunk. We also assume that prices p_j remain unchanged, consistent with the experimental evidence, which shows little adjustment in prices. Workers then choose labor supply h_j in this period, in response to the intervention.

For capital, we consider two versions of the simulation: one in which it is fixed and one in which capital can be adjusted, reflecting the possibility that capital adjustment costs limit short-run adjustments, as discussed earlier.

In the second period (after the realization of v_j), firms adjust only utilization, as in the baseline model. Given $(k_j, b_j, \tau_j^{treat}, h_j)$, firms choose the optimal utilization rate of labor.

Therefore, in the second period the firm's maximization problem follows the same steps outlined above, with the key difference that we the production is modified to reflect the design of the wage subsidy intervention:

$$y_j = A k_j^{\alpha} \Big((\phi(\tau_j) b_j + \phi(\tau_j^{treat})) h_j \Big)^{1-\alpha}. \tag{17}$$

As tasks are allocated equally across workers and do not depend on training, all workers, including the new trainee, optimally choose the same labor supply to the firm h_j. In the first period, firms

[16] Without an experiment, we would likely target the correlation between the firm's share of trained workers and its sales to measure returns to training. However, this correlation would be subject to endogeneity, capturing either heterogeneity in time-invariant firm characteristics, or, when using within firm variation, contemporaneous supply or demand shocks.

choose τ^{treat} and k to maximize expected profits:

$$\max_{k_j,\tau_j^{treat}} E_v[p_j y_j - w h_j l_j - \chi_j(\tau_j) b_j - \chi_j(\tau_j^{treat}) - r k_j + WS], \tag{18}$$

where WS denotes the wage subsidy, which, as in the data, is set equal to 65% of the wage.

Table 9 shows the reduced form experimental impacts of the intervention (Panel A), and those generated by the simulated model (Panel B). Within Panel A we report LATE estimates from an instrumental variable specification in which the hiring of a new worker is instrumented by treatment assignment. All effects are expressed in levels and as percentage changes relative to control firms. Panel B reports analogous changes in firm behavior between the baseline model and the model in which firms receive a subsidy to hire a new worker, for the two versions of the simulation: one in which capital is fixed and one in which capital can adjust.[17]

Treatment Effects We begin by describing the effects of the wage subsidy intervention on firm outcomes, focusing on the LATE estimates because of the high take-up rate of the wage subsidy (71%) and because, in the model, firms are homogeneous and therefore do not exhibit heterogeneity in take-up: given the value of the wage subsidy, it is optimal for all firms to hire an additional worker. These estimates are measured at first follow-up, immediately after the end of the six-month wage subsidy period.

Column 1 in Panel A shows that the wage subsidy treatment increased employment by 1.2 workers among firms that took up the offer, corresponding to a 54% increase relative to control firms, and implying that firms expanded their workforce slightly beyond the hiring of the subsidized worker. As Column 2 shows, this effect is primarily driven by an increase in the number of skilled workers, indicating that the additional workers hired by treated firms were provided with on-the-job training.

In line with firms facing costs of adjusting other inputs, we find no treatment effect on the value of assets owned or rented by the firm (Column 3). Utilization — the number of hours employees spend in production divided by the number of hours spent on the firm premises — remains unchanged (Column 4). However, total hours spent by workers on the firm premises decline by 11% (Column 5), suggesting that the wage subsidy has displacement effects on the intensive margin (fewer hours on the firm premises) rather than the extensive margin (fewer employees).

[17] In Panel A, we estimate LATE effects using the following second-stage ANCOVA regression,

$$Y_{ibst} = \beta \mathrm{WS}_i + \alpha_1 Y_{i0} + \alpha_2 \mathbf{X}_{i0} + \gamma_l + \gamma_s + \gamma_t + \varepsilon_{ibst},$$

where WS_i is a dummy equal to one if the firm *took up* the wage subsidy and hired an additional worker. It is instrumented with a dummy equal to one for firms assigned to receive the wage subsidy offer. Y_{i0} is the value of the outcome at baseline, and $\mathbf{X}_{i0}$ is a vector of additional controls measured at baseline (the firm owner's gender and education level, firm age, and firm age squared). γ_l,γ_s and γ_t are location, sector, and survey wave fixed effects, respectively.

Overall, taking up the wage subsidy is effective in that it increases sales by 32% (Column 6), with a corresponding, though statistically insignificant, 10% increase in profits (Column 7).[18]

Data vs. Simulated Impacts Panel B shows the impacts simulated in the model. We focus on the case where capital is fixed and firms are assumed to face capital adjustment costs. The results for outcomes other than assets are similar when firms can re-optimize capital.

Overall, the magnitude of the simulated effects aligns closely with the experimental estimates. The two outcomes for which the model deviates from the data are hours spent at the firm premises and profits. The divergence in hours arises because, in the model, labor supply is tightly linked to expected utilization, as shown in equation (16). Both in the data and in the model, utilization declines only modestly in response to the intervention, implying a limited adjustment in labor supply. This muted response helps explain why the model slightly overpredicts the profit impact of the wage subsidy: higher hours supplied in the model imply lower labor costs relative to the data. In contrast, if firms had to raise wages to induce existing workers to supply time at the firm (a mechanism not captured in our partial equilibrium framework), profits would increase by less than in the model.

6 Counterfactual Analysis

We use the model to conduct three types of counterfactual analysis: (i) a quantitative horse race between frictions proposed as drivers of low levels of on-the-job training in low-income economies; (ii) to simulate the impact of well-studied supply-side policies designed to improve firm performance across environments characterized by different levels of demand uncertainty; (iii) to simulate the impact of policies explicitly designed to reduce demand uncertainty.

6.1 Training Frictions

Our first set of counterfactuals run a quantitative horse race between demand uncertainty (σ_v) and three frictions that have been proposed as drivers of the low levels of on-the-job training in low-income countries: training costs (χ), high turnover (γ) and low levels of demand (μ_v). Training costs and worker separations have received extensive attention in the literature [Acemoglu and Pischke 1998, 1999, Caicedo *et al.* 2023], which emphasizes the role of up-front investment costs and limited worker retention in shaping firms' incentives to train. In sharp contrast, demand uncertainty has been largely overlooked as a determinant of training, despite its potential to distort utilization and reduce the returns to fixed training investments. Our first exercise therefore asks whether reducing demand uncertainty, holding mean level of demand fixed, can generate

[18] While the hours worked outcome in Column 5 is averaged across incumbent and new workers, the reduction in the number of hours on the firm premises is similar for both types of employees.

improvements in training and effective labor that are comparable to relaxing canonical frictions studied in the literature thus far.

To do so we construct counterfactual scenarios where we change the magnitude of each of the model parameters of interest (χ, γ, μ_v and σ_v) to induce a 10% increase in firm revenues. We then compare the resulting impacts on the number of workers (firm size), utilization, time spent by workers on the firm premises, training, and a composite measure of effective labor that incorporates all of these margins. The results are in Figure 6.

Panel A shows the relative changes in the model parameters necessary to induce a 10% rise in revenues. The same revenue increase can be achieved by lowering training costs or workers' outside option by approximately 5%, increasing the mean of the demand shock by 8.7%, or reducing demand uncertainty by .7%. This points to the quantitative importance of demand uncertainty: a small reduction in uncertainty generates the same revenue gains as substantially lowering canonical frictions related to training costs, worker's outside options or the level of demand.

Panel B shows the effects of the changes in model parameters from Panel A on various components of labor. The first is the number of workers employed — firm size. All parameter changes induce firms to hire more workers, as they either decrease the cost of, or increase the returns from, additional employees. Reductions in training costs, increases in demand levels and reductions in demand uncertainty have the largest effects on employment: 10.5%, 9.4% and 8.9% respectively. In contrast, reducing workers' outside option only leads to a 4.6% increase in employment. This is because, in response to a lower outside option, workers spend 5.5% more time at the firm premises, while the other parameters have negligible effects on this margin. As a result, a smaller increase in firm size is sufficient to generate comparable increases in total labor supply.

The third set of bars in Panel B show that utilization is also largely unaffected by the parameter changes. While this is what we would expect for training costs and the outside option, the result is more surprising for the demand-uncertainty counterfactual. The reason we observe only a small increase in utilization is that σ_v needs to fall by just .7% to generate a 10% increase in firm revenues. Due to its limited effect on utilization, lower demand uncertainty increases the share of trained workers by only .5%. Lower training costs and a lower outside option have larger effects, leading to a 2.4% increase in the share of trained workers.

Panel C shows results on the overall number of trained workers and effective labor, calculated as the product of the number of workers, utilization, time spent at the premises, and returns to training $\phi(\tau)$ (with the results being similar when using the share of trained workers τ).

On the number of trained workers, a reduction in training costs has the largest effect (13.1%), followed by an increase in demand levels and a decrease in demand uncertainty (9.7% and 9.4%, respectively). In other words, reducing demand uncertainty (while holding the *mean* fixed) raises the number of trained workers by almost as much as lowering training costs, a channel that has received far more attention. Decreasing workers' outside option has the smallest impact on the number of trained workers (7.1%), as it incentivizes firms to provide training to a smaller number

of workers who supply more time to the firm, rather than expanding firm size.

On effective labor, reductions in training costs and outside options generate the largest increases (12.9%). However, the effects induced by changes in demand uncertainty and demand levels are only slightly lower (10.3% and 10.6%, respectively). Importantly, this implies that stabilizing demand — without increasing its average level — can deliver gains in effective labor that are nearly as large as those obtained by directly reducing training costs or turnover frictions. In other words, demand uncertainty itself is a quantitatively important determinant of training and effective labor, alongside canonical mechanisms previously studied. Demand stabilization policies can therefore be as powerful as traditional supply-side interventions aimed at stimulating skill investment.

6.2 The Effectiveness of SME Policies and Demand Uncertainty

Small and medium-sized enterprises are the heart of the economies of many developing countries — given skewed firm size distributions, they constitute the majority of firms in such contexts, and the majority of workers are employed in such enterprises. As a result, a large literature has developed to examine policies to improve firm performance and stimulate SME growth in low-income countries. A large experimental literature has evaluated capital grants [De Mel *et al.* 2008, Blattman *et al.* 2014, De Mel *et al.* 2019], management training [Bloom *et al.* 2013, Bloom *et al.* 2014, McKenzie and Woodruff 2014, 2017], wage or training subsidies [De Mel *et al.* 2019, Alfonsi *et al.* 2020, Hardy and McCasland 2023, Brown *et al.* 2024], and market access interventions [Atkin *et al.* 2017, Ali *et al.* 2025]. This work finds impacts of such interventions that are often positive, but modest. We contribute by documenting a new mediating mechanism that can limit the effectiveness of such interventions: when firms face high demand uncertainty, they optimally keep resources underutilized, which reduces the returns to any intervention that increases input availability or quality.

To explore this intuition, we simulate the impact of supply-side policies across environments characterized by different levels of demand uncertainty, to assess how policy effectiveness varies with this key source of uncertainty faced by firms. This allows us to evaluate whether demand stabilization can amplify policy effectiveness by increasing the returns of investing in inputs that require up-front investments. We consider a range of outcomes — firm size, capital, profits and effective labor — where the interventions considered typically aim to increase the profitability of firms and at least one of these other margins.

Figure 7 reports results from model simulations of four counterfactual policies: reducing training costs (Panel A), reducing capital rental rates (Panel B), improved market access (Panel C), and management training (Panel D). These are simulated through changes in χ, r, μ_v, and A respectively, that generate a 10% increase in sales revenue in the baseline scenario. We simulate these policies under three environments characterized by differing levels of demand uncertainty:

baseline (in yellow), 5% lower volatility relative to baseline (in blue), and 5% higher volatility relative to baseline (in green). For each environment, the graphs show changes in the number of workers (firm size), capital, profits and effective labor induced by the intervention.

The effectiveness of this broad set of intervention types is always mediated by the demand uncertainty firms face. A 5% decrease in volatility more than doubles the impact of all policies on labor, capital, effective labor, and profits. Conversely, a 5% increase in volatility leads to effects that are half the size of those in the baseline scenario. This pattern is consistent across interventions and reflects a common mechanism: lower uncertainty raises expected utilization, thereby increasing the returns to expanding capacity through hiring and capital accumulation. In contrast, when uncertainty is high, firms anticipate inputs will remain idle for a larger share of time, dampening responses to policies operating through these margins.

Two further points are of note.

First, Figure 7 highlights that the magnitude of this amplification effect is similar across different interventions. Policies that act directly on costs (such as reductions in χ or r) and those that increase demand or productivity (through μ_v or A) all exhibit substantially larger effects in low-uncertainty environments. This suggests demand stabilization not only shifts outcomes directly, as highlighted in the first counterfactual, but also indirectly increases the responsiveness of firms to a broad class of supply-side policies because it is a mediator for these other policy instruments. The returns to these interventions are not fixed structural parameters but are mediated by the demand environment in which firms operate.

Second, the results on how the impact of interventions on profits is mediated through demand uncertainty has implications for firm's compliance with the offer of such interventions. If firms face costs of compliance, then even limited up front costs of compliance can offset profit impacts in environments characterized by high demand uncertainty.

6.3 Reducing Demand Uncertainty

The final set of counterfactual exercises come full circle to consider the origins of demand uncertainty — and policies that could feasibly reduce it. We take as our starting point the earlier descriptive evidence on the origins of demand uncertainty from Figure 3, that suggested uncertainty is negatively correlated with: (i) the number of customers, (ii) the share of professional ties in business networks. These patterns suggest that firms with access to a more professional network and a broader, more interconnected customer base experience lower uncertainty, consistent with a mechanism in which idiosyncratic customer-level shocks are imperfectly diversified at the firm level. Policies that reduce the correlation of demand shocks across firms — for example, by encouraging professional rather than personal business relationships — or that facilitate the sharing of customers across firms could therefore help smooth fluctuations in demand. In Figure 8, we investigate the quantitative effect of these two channels — professional business networks

and demand sharing — on demand volatility and firm outcomes.

We begin by examining the impact of firms only engaging in professional networks (as opposed to also family-based networks). This could be implemented by incentivizing business owners to meet with other firm owners outside their existing network [Brooks *et al.* 2018, Cai and Szeild 2018]. To do so, we use the coefficient estimated in Figure 3 to infer the counterfactual coefficient of variation of sales that would prevail if all the existing network connections were strictly professional (among those firms that have a network). As previously discussed, this could reduce demand uncertainty faced by firms by lowering the correlation between the shocks they are exposed to relative to a scenario where firms are linked primarily through personal ties. Panel A shows that despite the large increase in the share of firms with professional networks (from 35% to 62% across all follow-ups), this only reduces average (median) demand uncertainty from .30 to .28 (from .25 to .23). As a result, the impact on firm outcomes, shown in Panel B, is modest: the number of workers, capital, profits, and effective labor increase by between 9.1% and 10.9%.

The second policy we consider to reduce demand uncertainty is for firms to share customers. Such within-sector-location aggregation could be implemented through cooperatives, shared showrooms, or centralized platforms where customers place orders and these are reallocated across firms based on available capacity [Bergquist *et al.* 2025]. These arrangements allow firms to effectively share demand while maintaining decentralized production. To quantify the scope for such smoothing, we compare the coefficient of variation of revenues at the firm level to that of aggregate sales within sector-location clusters. We are able to quantify this accurately given we survey nearly all firms within the same cluster.

If firm-level shocks were perfectly correlated, aggregation would not reduce uncertainty and the two coefficient of variations would be identical. Conversely, if shocks are largely idiosyncratic, aggregation should reduce uncertainty. The result in Panel A shows this latter scenario to be more relevant. Aggregation leads to a substantial decline in demand uncertainty: the mean (median) CV falls from .30 to .13 (from .25 to .15). Consistent with the earlier evidence (Figure 2), this indicates a large share of firm-level demand fluctuations reflects diversifiable, idiosyncratic shocks.

We then use the model to conduct a counterfactual in which we lower σ_v so that each firm faces the level of demand uncertainty implied by sharing customers in the cluster, while holding mean demand and all other parameters constant.[19] The result in Panel B shows that lower demand uncertainty leads to large increases in firm outcomes: firm size increases by 71%, while capital, effective labor, and profits increase by more than 80%.

Taken together, these results suggest that demand sharing is currently limited, even among firms connected through professional networks. While such networks are associated with slightly lower demand uncertainty, the effect is small relative to the gains that could be achieved through effective pooling of customer demand. This wedge suggests contractual frictions may limit actual

[19]The model is in partial equilibrium and so this counterfactual assumes that the returns to investing in training or capital do not diminish as more firms in the cluster invest as a result of the lower demand uncertainty.

sharing of demand within the cluster. Reallocating customers across firms requires agreements over quality, compensation, and future customer relationships, which may be difficult to monitor and enforce in settings characterized by high levels of informality and weak institutions. As a result, demand remains fragmented across many small firms, each exposed to substantial idiosyncratic risk.

The results therefore highlight how two central features of low-income economies — small firm size and contractual frictions — increase demand uncertainty and, through this channel, constrain firm growth. Small firms are less able to diversify idiosyncratic demand shocks, while contractual frictions prevent the aggregation of demand across firms. As a result, firms face lower expected utilization, reducing incentives to invest in capital and on-the-job training. This has consequent impacts on aggregate growth and paths of development.

7 Conclusions

We study the causes and consequences of labor utilization in small and medium-sized firms in a low-income setting. While our data come from Uganda, the four stylized facts we establish are likely present in other similar settings. First, firms face substantial and difficult-to-predict demand volatility. Second, capital and labor adjustment costs are significant, limiting firms' ability to respond to shocks along extensive margins. Third, the idleness of labor is pervasive. Fourth, firms with more uncertain demand exhibit lower labor utilization and invest less in on-the-job training. Guided by these facts, we develop and estimate a model of firm input and training choices under demand uncertainty and adjustment costs. A key mechanism is that training requires up-front fixed investment, while workers are largely compensated on a piece-rate basis for output produced. Higher demand volatility lowers expected utilization and therefore reduces the returns to fixed training investments, generating a novel demand-side channel through which uncertainty depresses firm-sponsored human capital accumulation. We validate the model by showing that simulated treatment effects of a wage subsidy intervention closely track the experimental estimates.

Our findings highlight the importance of demand-side frictions in shaping firm dynamics and human capital accumulation in low-income settings [Goldberg and Reed 2023]. Policies that reduce demand uncertainty through improved market integration, demand stabilization, or better market information may therefore be a crucial complement to conventional supply-side interventions, and their absence may help explain why the latter have so often disappointed.

Our analysis suggests two broad directions for future work.

First, there is much more to be understood about the fundamental drivers of demand uncertainty. Our model allows us to take a first step in understanding how aggregating demand within sector-location clusters can push forward firm size, capital accumulation, profits and effective labor in the economy. However, the current level of demand sharing within clusters is limited. In similar

contexts, well-functioning rental markets allow firms to share *capital* across the cluster in sectors where firms need access to many large machines [Bassi *et al.* 2022]: a machine is something that one firm can clearly rent to another for a specified period and fee, with little ambiguity about what was traded. Pooling *demand* is contractually harder. It requires agreements over which firm fulfills which order, the quality standards each must meet, how revenue is allocated, and how future customer relationships are governed — none of which is naturally separable from the firm itself. Policy intervention may therefore be needed to stimulate demand pooling, for example through cooperatives, shared showrooms, or centralized platforms where customers place orders that are reallocated across firms based on available capacity [Bergquist *et al.* 2025]. Testing the effectiveness of such interventions in reducing demand uncertainty and stimulating firm growth within clusters remains a priority for future research.

Other fundamental drivers of demand uncertainty might stem from income volatility of consumers. Indeed, in our context the low utilization of labor stemming from demand uncertainty implies worker earnings are lower for those employed in firms exposed to greater uncertainty, starting a cycle of volatile worker earnings that then impact the volatility of demand faced by other firms. Finally, another fundamental source of demand uncertainty might be the fact that consumers are able to search over many firms in geographically concentrated sectors. This raises a delicate issue — that although reducing frictions to consumer search raises consumer welfare [Vitali 2025], if this then leads to greater demand uncertainty faced by firms, this can reduce producer welfare.

Second, our analysis has implications for the measurement and causes of productivity differences across firms [Hsieh and Klenow 2009]. That focuses on static across-firm misallocation, typically inferring distortions from dispersion in marginal revenue products across firms. Recent work has shown that inference is complicated by endogenous capital utilization: when firms adjust utilization flexibly in response to shocks, dispersion in average revenue products of capital can reflect efficient reallocation of capital services (i.e., utilized labor) rather than misallocation [Kabir *et al.* 2025]. Our paper identifies a further and distinct problem: demand uncertainty generates within-firm underutilization that not only makes firms appear less productive [Butters 2020] but actively suppresses investment in training that would raise long-run productivity. This dynamic channel compounds over time in a way that static misallocation frameworks do not capture.

A Appendix

A.1 Proofs

Proof of Proposition 1 Let $U = E_v(u)$ be the firm's expected utilization, which can be written as

$$U(\tilde{v}, \sigma_v) = \frac{M_1(\tilde{v}, \sigma_v)}{\tilde{v}} + \bar{F}(\tilde{v}, \sigma_v)$$

where $M_1(\tilde{v}, \sigma_v) = \int_{\underline{v}}^{\tilde{v}} v dF(v)$ and $\bar{F}(\tilde{v}, \sigma_v) = 1 - F(\tilde{v}, \sigma_v)$. We first prove part (1) of the proposition. Holding prices and inputs fixed also holds the cutoff $\tilde{v}$ fixed. The function

$$u(v) = \min\left\{\frac{v}{\tilde{v}}, 1\right\}$$

is increasing and concave in v: it is linear for $v < \tilde{v}$ and flat for $v \geq \tilde{v}$. Therefore, by second-order stochastic dominance, any mean-preserving increase in the dispersion of v weakly lowers expected utilization:

$$\frac{\partial U(\tilde{v}, \sigma_v)}{\partial \sigma_v} \leq 0$$

The inequality is strict whenever the mean-preserving spread places positive probability mass on both sides of the kink at $\tilde{v}$. This proves part (1).

We now prove part (2). When firms re-optimize prices, inputs remain fixed, so q and marginal cost are fixed. However, the cutoff changes with the price: $\tilde{v} = qp^{\theta}$. Taking logs, $\log \tilde{v} = \log q + \theta \log p$. Since q is fixed,

$$\frac{d \log \tilde{v}}{d\sigma_v} = \theta \frac{d \log p}{d\sigma_v}$$

With fixed inputs marginal cost is fixed and optimal prices can be written as: $p = \mu mc$. Therefore, $\frac{d \log p}{d\sigma_v} = \frac{d \log \mu}{d\sigma_v}$ and $\frac{d \log \tilde{v}}{d\sigma_v} = \theta \frac{d \log \mu}{d\sigma_v}$. To derive the total effect of volatility on utilization, we first compute the effect of the cutoff on utilization. Notice that:

$$\frac{\partial M_1(\tilde{v}, \sigma_v)}{\partial \tilde{v}} = \tilde{v} f(\tilde{v}, \sigma_v), \qquad \frac{\partial \bar{F}(\tilde{v}, \sigma_v)}{\partial \tilde{v}} = -f(\tilde{v}, \sigma_v)$$

Therefore,

$$\frac{\partial U}{\partial \tilde{v}} = \frac{\partial}{\partial \tilde{v}}\left(\frac{M_1}{\tilde{v}}\right) + \frac{\partial \bar{F}}{\partial \tilde{v}} = \frac{\tilde{v} f(\tilde{v}) \tilde{v} - M_1}{\tilde{v}^2} - f(\tilde{v}) = -\frac{M_1}{\tilde{v}^2}$$

Taking the total derivative of utilization with respect to volatility gives:

$$\frac{dU}{d\sigma_v} = \frac{\partial U}{\partial \sigma_v} + \frac{\partial U}{\partial \log \tilde{v}} \frac{d \log \tilde{v}}{d\sigma_v} = \frac{\partial U}{\partial \sigma_v} - \frac{M_1(\tilde{v}, \sigma_v)}{\tilde{v}} \theta \frac{d \log \mu}{d\sigma_v}$$

Therefore, $\frac{dU}{d\sigma_v} \leq 0$ if and only if

$$\frac{d \log \mu}{d\sigma_v} \geq \frac{1}{\theta} \frac{\partial U / \partial \sigma_v}{M_1(\tilde{v}, \sigma_v)/\tilde{v}}$$

By part (1), $\frac{\partial U}{\partial \sigma_v} \leq 0$. Thus the right-hand side is weakly negative. Therefore a sufficient condition for utilization to decrease under re-optimized prices is $\frac{d \log \mu}{d \sigma_v} \geq 0$. That is, if markups are weakly increasing in volatility, then re-optimized prices reinforce the direct decline in utilization. Conversely, utilization can increase only if markups fall sufficiently fast with volatility. This proves part (2).

Proof of Proposition 2 — Generic Distribution Using the definitions of marginal cost and capacity, we have $MC_j q_j = whb_j$. Therefore, the first order conditions with respect to training, capital, and labor can be rewritten as:

$$
\begin{aligned}
(\tau): \qquad & (1-\alpha) w h_j \frac{\phi'(\tau_j)}{\phi(\tau_j)} \Big[E_v(u_j) + (\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}_j)) \Big] = \chi_j'(\tau_j) b_j \\
(k): \qquad & \frac{\alpha}{k_j} w h_j b_j \Big[E_v(u_j) + (\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}_j)) \Big] = r \\
(b): \qquad & w h_j \Big[-\alpha E_v(u_j) + (1-\alpha)(\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}_j)) \Big] = \chi_j(\tau_j)
\end{aligned}
$$

Training: Since $\phi(\cdot)$ is increasing and concave, $\phi'(\tau_j)/\phi(\tau_j)$ is decreasing in τ_j. Since training costs are convex, $\chi_j'(\tau_j)$ is increasing in τ_j. Hence a decline in

$$
E_v(u_j) + (\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}_j))
$$

shifts the marginal-benefit of training downward and weakly reduces the optimal choice of training. Therefore,

$$
\frac{d \tau_j}{d \sigma_v} \leq 0
$$

whenever

$$
\frac{d}{d \sigma_v} \Big[E_v(u_j) + (\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}_j)) \Big] \leq 0
$$

Capital: Using $MC_j q_j = whb_j$, the capital FOC becomes

$$
k_j = \frac{\alpha w h_j}{r} b_j \Big[E_v(u_j) + (\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}_j)) \Big].
$$

Since α, w, h_j, and r are fixed and positive, capital is weakly decreasing in σ_v whenever

$$
\frac{d}{d \sigma_v} \Big\{ b_j \Big[E_v(u_j) + (\mu(\tilde{v}_j) - 1)(1 - F(\tilde{v}_j)) \Big] \Big\} \leq 0
$$

Number of workers: The term in brackets in the labor FOC is the expected net marginal benefit of hiring an additional worker. Employment is weakly decreasing in σ_v when the expected net marginal benefit of labor falls relative to the per-worker training cost. Since wh_j is fixed and

positive, this condition is

$$\frac{d}{d\sigma_v}\Big[-\alpha E_v(u_j) + (1-\alpha)(\mu(\tilde{v}_j)-1)(1-F(\tilde{v}_j))\Big] \leq \frac{1}{wh_j}\frac{d}{d\sigma_v}\chi_j(\tau_j)$$

as stated in the proposition.

Proof of Proposition 2 (Training) — Log-normal distribution, fixed prices Let $m = E[v]$ and assume $\log v \sim N\left(\log(m - \frac{\sigma_v^2}{2}), \sigma_v^2\right)$, so that $E[v] = m$ is held fixed as σ_v changes. Holding prices fixed, the volatility-sensitive component of the training FOC is:

$$E_v(u) + \mu(1 - F(\tilde{v}, \sigma_v))$$

where μ is fixed. The derivative of this term with respect to σ_v if:

$$\frac{d}{d\sigma_v}E_v(u) + \mu\frac{d}{d\sigma_v}(1 - F(\tilde{v}, \sigma_v))$$

By Proposition 1, $dE_v(u)/d\sigma_v \leq 0$ when prices are fixed. Let $a = \log\left(\frac{\tilde{v}}{m}\right)$. For the fixed-mean lognormal,

$$1 - F(\tilde{v}, \sigma_v) = 1 - \Phi\left(\frac{a}{\sigma_v} + \frac{\sigma_v}{2}\right)$$

Differentiating,

$$\frac{\partial(1 - F(\tilde{v}, \sigma_v))}{\partial\sigma_v} = \phi\left(\frac{a}{\sigma_v} + \frac{\sigma_v}{2}\right)\left(\frac{a}{\sigma_v^2} - \frac{1}{2}\right)$$

so

$$\frac{\partial(1 - F(\tilde{v}, \sigma_v))}{\partial\sigma_v} \leq 0 \quad \Longleftrightarrow \quad a \leq \frac{\sigma_v^2}{2} \quad \Longleftrightarrow \quad \tilde{v} \leq me^{\sigma_v^2/2}.$$

Therefore, if $\tilde{v} \leq me^{\sigma_v^2/2}$, then both $\frac{d}{d\sigma_v}E_v(u) \leq 0$ and $\frac{d}{d\sigma_v}(1 - F(\tilde{v}, \sigma_v)) \leq 0$. Since $\mu > 0$, it follows that training falls when $\tilde{v}$ falls in this region.

References

[1] ACEMOGLU.D AND J.S.PISCHKE (1998) "Why Do Firms Train? Theory and Evidence," *Quarterly Journal of Economics* 113: 79-119.

[2] ACEMOGLU.D AND J.S.PISCHKE (1999) "The Structure of Wages and Investment in General Training," *Journal of Political Economy* 107: 539-72.

[3] 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: 2369-414.

[4] ALFONSI.L, V.BASSI, I.RASUL AND E.SPADINI (2026) Firm Responses to Aggregate Shocks and Implications for Workers: Lessons from the Pandemic, mimeo UCL.

[5] ALI.N, G.DE GIORGI, A.RAHMAN AND E.VERHOOGEN (2025) What Do Market-Access Subsidies Do? Experimental Evidence from Tunisia, NBER WP33985.

[6] AMBERG.N, AND R.FRIBERG AND C.SYVERSON (2026) Have We Got News For You: Theory and Evidence on Firms' Optimal Expected Capacity Utilization, NBER WP33400.

[7] ATKIN.D, A.KHANDELWAL AND A.OSMAN (2017) "Exporting and Firm Performance: Evidence from a Randomized Experiment," *Quarterly Journal of Economics* 132: 551-615.

[8] AUTOR.D.H (2001) "Why Do Temporary Help Firms Provide Free General Skills Training," *Quarterly Journal of Economics* 116: 1409-48.

[9] BANDIERA.O, BASSI.V, BURGESS.R, RASUL.I, SULAIMAN.M AND A.VITALI (2025) "The Search for Good Jobs: Evidence from a Six-year Field Experiment in Uganda," *Journal of Labor Economics* 43: 885-935.

[10] BASSI.V, R.MUOIO, T.PORZIO, R.SEN AND E.TUGUME (2022) "Achieving Scale Collectively," *Econometrica* 90: 2937-78.

[11] BASSI.V, J.LEE, A.PETER, T.PORZIO, R.SEN AND E.TUGUME (2025) Self-Employment within the Firm, NBER WP31740.

[12] BASU.S, J.G.FERNALD AND M.S.KIMBALL (2006) "Are Technology Improvements Contractionary?," *American Economic Review* 96: 1418-48.

[13] BECKER.G.S AND K.M.MURPHY (1992) "The Division of Labor, Coordination Costs, and Knowledge," *Quarterly Journal of Economics* 107: 1137-60.

[14] BLATTMAN.C, N.FIALA AND S.MARTINEZ (2014) "Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda," *Quarterly Journal of Economics* 129: 697-752.

[15] BLATTMAN.C, N.FIALA AND S.MARTINEZ (2020) "The Long-Term Impacts of Grants on Poverty: Nine-Year Evidence from Uganda's Youth Opportunities Program," *AER: Insights* 2: 287-304.

[16] BLOOM.N, (2009) "The Impact of Uncertainty Shocks," *Econometrica* 77: 623-685.

[17] BLOOM.N, B.EIFERT, A.MAHAJAN, D.MCKENZIE AND J.ROBERTS (2013) "Does Management Matter? Evidence from India," *Quarterly Journal of Economics* 128: 1-51.

[18] BLOOM.N, M.FLOETOTTO, N.JAIMOVICH, I.SAPORTA-EKSTEN AND S.J.TERRY (2018) "Really Uncertain Business Cycles," *Econometrica* 86: 1031-1065.

[19] BLOOM.N, R.LEMOS, R.SADUN, D.SCUR AND J.VAN REENEN (2014) "The New Empirical Economics of Management," *Journal of the European Economic Association* 12: 835-76.

[20] BLOOM.N AND J.VAN REENEN (2007) "Measuring and Explaining Management Practices Across Firms and Countries," *Quarterly Journal of Economics* 122: 1351-1408.

[21] BOEHM.C.E AND N.PANDALAI-NAYAR (2022) "Convex Supply Curves," *American Economic Review* 112: 3941-69.

[22] BROOKS.W, K.DONOVAN AND T.R.JOHNSON (2018) "Mentors or Teachers? Microenterprise Training in Kenya," *American Economic Journal: Applied Economics* 10: 196-221.

[23] BROWN.G, HARDY.M, MBITI.I, MCCASLAND.J, AND I.SALCHER (2024) "Can Financial Incentives to Firms Improve Apprenticeship Training? Experimental Evidence from Ghana," *American Economic Review: Insights* 6: 120-36.

[24] BURNSIDE.C, M.EICHENBAUM AND S.REBELO (1993) "Labor Hoarding and the Business Cycle," *Journal of Political Economy* 101: 245-73.

[25] BURNSIDE.C AND M.EICHENBAUM (1996) "Factor-Hoarding and the Propagation of Business-Cycle Shocks," *American Economic Review* 86: 1154-74.

[26] BUTTERS.R.A (2020) "Demand Volatility, Adjustment Costs, and Productivity: An Examination of Capacity Utilization in Hotels and Airlines," *AEJ: Microeconomics* 12: 1-44.

[27] CAI.J AND AND A.SZEIDL (2018) "Interfirm Relationships and Business performance," *Quarterly Journal of Economics* 133: 1229-82.

[28] CAICEDO.S, M.ESPINOSA AND A.SEIBOLD (2022) "Unwilling to Train? Firm Responses to the Colombian Apprenticeship Regulation," *Econometrica* 90: 507-50.

[29] CEFALA.L, P.NASO, M.NDAYIKEZA AND N.SWANSON (2025a) Under-training by Employers in Informal Labor Markets: Evidence from Burundi, mimeo, UC Berkeley.

[30] CEFALA.L, F.IRAKOZE, P.NASO AND N.SWANSON (2025b) The Economic Consequences of Knowledge Hoarding, mimeo, UC Berkeley.

[31] COOLEY.T.F, G.D.HANSEN AND E.C.PRESCOTT (1995) "Equilibrium Business Cycles with Idle Resources and Variable Capacity Utilization," *Economic Theory* 6: 35-49.

[32] COOPER.R, HALTIWANGER.J AND J.L.WILLIS (2007) "Search Frictions: Matching Aggregate and Establishment Observations," *Journal of Monetary Economics* 54: 56-78.

[33] DE MEL.S, D.MCKENZIE AND C.WOODRUFF (2008) "Returns to Capital in Microenterprises: Evidence from a Field Experiment," *Quarterly Journal of Economics* 123: 1329-72.

[34] DE MEL.S, D.MCKENZIE AND C.WOODRUFF (2019) "Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises," *AEJ: Applied Economics* 11: 202-35.

[35] FAGNART.J-F, O.LICANDRO AND F.PORTIER (1999) "Firm Heterogeneity, Capacity Utilization and the Business Cycle," *Review of Economic Dynamics* 2: 433-55.

[36] FILMER.D, AND L.FOX (2014) "Youth Employment in Sub-Saharan Africa," *World Bank and AFD*, Washington, DC.

[37] GOLDBERG.P.K AND T. REED (2023) "Demand-Side Constraints in Development: The Role of Market Size, Trade, and (In)Equality," *Econometrica* 91: 1915-50.

[38] GREENWOOD.J, Z.HERCOWITZ AND G.W.HUFFMAN (1988) "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review* 78: 402-17.

[39] HARDY.M AND J.MCCASLAND (2023) "Are Small Firms Labor Constrained? Experimental Evidence from Ghana," *AEJ: Applied Economics* 15: 253-84.

[40] HAUSMAN.J (1996) "Valuation of New Goods Under Perfect and Imperfect Competition," in *The Economics of New Goods*, University of Chicago Press.

[41] HAUSMAN.J, G.LEONARD AND J.D.ZONA (1994) "Competitive Analysis with Differenciated Products," *Annales d'Economie et de Statistique* 34: 159-80.

[42] HSIEH.C-T AND P.J. KLENOW (2009) "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics* 124: 1403-48.

[43] ILO (2023) "Apprentices in Countries with Large Informal Economies," *Statistical Brief, International Labour Organization*, Geneva.

[44] KABIR.P, E.TAN AND I.VARDISHVILI (2025) "Quantifying the Allocative Efficiency of Capital: The Role of Capital Utilization," *Journal of Monetary Economics*, forthcoming.

[45] LAGAKOS.D, B.MOLL, T.PORZIO, N.QIAN AND T.SCHOELLMAN (2018) "Life Cycle Wage Growth Across Countries," *Journal of Political Economy* 126: 797-849.

[46] MACCHI.E AND J.STALDER (2024) Informal Redistribution Through Work, mimeo, Brown University.

[47] MCKENZIE.D AND A.L.PAFFHAUSEN (2019) "Small Firm Death in Developing Countries," *Review of Economics and Statistics* 101: 645-57.

[48] MCKENZIE.D. AND C.WOODRUFF (2014) "What Are We Learning from Business Training and Entrepreneurship Evaluations around the Developing World?" *World Bank Research Observer* 29: 48-82.

[49] MCKENZIE.D AND C.WOODRUFF (2017) "Business Practices in Small Firms in Developing Countries," *Management Science* 63: 2967-81.

[50] SWANSON.N (2025) All in the Family: Kinship Pressure and Firm-Worker Matching, mimeo, Stanford University.

[51] UMA (2023) An Assessment of Capacity Utilization in Uganda Manufacturing Sector, Policy Report, Kampala, Uganda.

[52] VITALI.A (2025) Consumer Search and Firm Location: Theory and Evidence from the Garment Sector in Uganda, mimeo, NYU.

[53] WALKER.W.M, N.SHAH, E.MIGUEL, D.EGGER, F.S.SOLIMAN AND T.GRAFF (2025) Slack and Economic Development, NBER WP33055.

Tables and Figures

Table 1: Firms at Baseline

Sample: All firms at baseline
Means, standard deviations of continuous variables in parentheses

	Mean/Percentage	Standard Deviation
A. Firm characteristics		
Age of firm	6.42	(4.54)
Monthly sales	499	(667)
Monthly profits [USD]	218	(329)
Manufacturing	34.3	
	Tailoring (14.9%), Welding (14.8%), Construction (4.6%)	
Services	65.7	
	Hairdressing (39.7%), Motor mechanics (9.5%), Catering (7.3%), Electrical (6.0%), Plumbing (3.2%)	
B. Firm owner characteristics		
Age of owner	34.2	(7.30)
Gender of owner (female)	53.8	
Hours spent by firm owner at premises (monthly)	267	(73.0)
C. Capital and labor		
Value of capital stock [USD]	1256	(2361)
Capital per worker [USD]	481	(994)
Number of employees	2.96	(2.31)
Share with stock value < monthly revenues	68.2	
Share with daily liquidity < monthly revenues	95.7	
D. Networks		
Any network firm	.788	
Number of network firms	1.10	(.754)
Share of firms that are professional ties | any network	.368	
Within the network:		
share workers	8.78	
share machines	23.0	
share credit	20.3	
share demand information	9.85	

Notes. All the variables are at the firms level and collected at baseline. The table reports mean outcomes and standard deviations for continuous outcomes in parentheses. In Panel B *Hours spent by firm owner at premises (monthly)* are computed multipling the average number of hours by six, the average number of days workers report being at the firm in a typical day, times four weeks. The data for this statistic is from Control firms at first follow-up. In Panel C, network firms are identified as other businesses in the same area with whom the firm owner shares something (such as inputs, machines, workers, information, credit, advice, etc.) or meets to talk. Professional ties are businesses with whom the owner reports having a business relation (vs. family members, friends, and neighbors). All monetary variables are winsorized at the top 1%and deflated at August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD.

Table 2: Worker Descriptives

Sample: Worker roster from control firms at first follow-up
Means, standard deviations of continuous variables in parentheses

	Mean or Percentage (1)	Standard Deviation (2)
A. Compensation		
Full time	95.3	
Paid piece rate (along with salary or in-kind)	66.8	
Paid piece rate only	49.1	
B. Tenure, hours and earnings		
Tenure at the firm (months)	19.2	(23.3)
Hours worked per day	11.0	(2.15)
Hours workers spent at the premises last month	280	(70.1)
Hours workers spent at the premises last month (trainee)	247	(59.4)
C. Sources of training		
Received on-the-job training (at the firm)	49.8	
Completed vocational training	26.8	
On-the-job training considered useful in other firms	63.0	
D. Time costs of training		
Length of training (months)	9.30	(5.21)
Mainly conducted by owner	45.6	
Conducted by owner and co-workers	45.1	
Conducted by co-workers only	9.07	
Hours owner spent training in the last month | Training	113	(101)
Hours owner spent training in the last month	50.3	(87.8)
E. Labor costs of training		
Trainee was unpaid	51.5	
Trainee was paying the owner	28.8	
Trainee was paid	19.7	
F. Tools/material costs of training		
Any such cost	82.7	
Owner paid most costs | any such costs	55.9	
Employee paid most costs | any such costs	36.2	
Owner-employee share costs | any such costs	7.91	

Notes: All data is from the employee roster for control firms at 1st follow-up. The table reports mean outcomes and standard deviations for continuous outcomes in parentheses. In Panel A the *'paid piece rate'* dummy equals one if the employee is paid a piece-rate (along with any other form of payment – including a base salary or in-kind payment). In Panel B, the *total hours worked per month* is calculated by multiplying hours worked per day by days worked per week, and time four weeks. *Total earnings in the last month* includes both money and in-kind payments made to the employee. In Panel C *On-the-job training considered useful in other firms* is a dummy equal to one if the owner reported that the skills learnt by their employees through on-the-job training received at the firm are "Useful both in this firm as well as in other firms". All monetary variables are winsorized at the top 1% (in-kind and money earnings are winsorized separately before they are summed to calculate total earnings) and deflated at August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD.

Table 3: Firm Provided Training

Sample: Worker roster from control firms at first follow-up (Panels A, B), and second, third, and fourth follow-up (Panel C)

Means, standard deviations of continuous variables in parentheses

	On-the-job trained	Not on-the-job trained	[p-value]
	(1)	(2)	(1) = (2)
Share of type	49.8	50.2	
A. Tenure, hours and earnings			
Full time	95.8	94.2	[.336]
Tenure at the firm (months)	23.7	14.7	[.000]
	(28.3)	(18.5)	
Total earnings last month [USD]	66.8	46.5	[.000]
	(50.3)	(46.7)	
Hours workers spent at the premises last month	281	277	[.486]
	(79.6)	(66.9)	
B. Productivity			
Owner assessed productivity at time of survey (0-10)	8.93	7.89	[.000]
	(1.88)	(2.89)	
Able to perform typical sector specific task at time of survey	0.91	0.81	[.001]
Owner assessed productivity when recruited (0-10)	2.81	4.77	[.000]
	(3.34)	(4.24)	
Able to perform typical task when recruited	0.20	0.50	[.000]
C. Products			
Hours spent working on the typical product/service	5.92	4.63	[.047]
	(11.8)	(7.00)	
Number of typical items produced last month	25.4	43.3	[.009]
	(44.3)	(99.7)	
Price charged for typical product/service [USD]	68.2	53.8	[.061]
	(119)	(105)	
Value of materials used for typical product/service [USD]	37.6	29.2	[.085]
	(78.2)	(63.7)	

Notes: All data is from the employee roster for control firms. The table reports mean outcomes and standard deviations for continuous outcomes in parentheses. Panels A and B use data from the first follow-up. Panel C uses data from the second, third, and fourth follow-up, as questions about the typical product/service done by employees were only asked then. Workers are split into on-the-job trained workers (Column 1) and not on-the-job trained workers (Column 2), based on whether the owner reports the worker ever received training at the firm. Individuals who receievd vocational training, but not on-the-job training are excluded from the sample. Column 3 reports the p-value of a t-test of equality of means between FT and non-FT workers. In Panel A, the *total hours worked per month* is calculated by multiplying hours worked per day by days worked per week, and multiplying by four. *Total earnings in the last month* includes both money and in-kind payments made to the employee. Panel B reports the owner's 0-10 assessment of how productive each employee is relative to all workers the owner knows in the same sector, and an indicator for whether the worker is able to independently perform a sector-specific task. In Panel D, for each worker, the owner reports the typical product the employee works on and provides information on the production and pricing of that product. The number of typical items produced last month is calculated at the worker level as total earnings last month / piece rate for the typical product. The number of typical items produced last month, and the price received by the employee for the typical product (the piece-rate), are therefore only available for worker paid piece-rate. All monetary variables are winsorized at the top 1% (in-kind and money earnings and winsorized separately before they are summed for total earnings) and deflated at August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD.

Table 4: Demand Uncertainty

Sample: Control firms at first follow-up
Means, standard deviations in parentheses

	Min	Max	Expected Value	Standard Deviation	Coefficient of Variation
	(1)	(2)	(3)	(4)	(5)
A. Expected Demand					
Expected monthly revenues, next six months	239	808	467	125	.239
	(288)	(1093)	(582)	(207)	(.104)
Expected customers, next seven days	9.28	28.4	20.2	4.36	.219
	(10.9)	(29.1)	(19.8)	(4.81)	(.097)
B. Realized Demand					
Monthly revenues, last three months	341	602	473	138	.305
	(415)	(728)	(582)	(198)	(.187)

Note: All data is from Control firms at first follow-up. Each firm owner was asked the following questions: (1) Imagine that over the next 6 months the revenues from this business are very high. What is the maximum amount of revenues the ownercould make in the best month over the next 6 months? (2) Imagine that over the next 6 months the revenues from this business are very low. What is the minimum amount of revenues the owner could make in the worst month over the next 6 months? (3) In how many months over the next 6 months, does the owner think that the monthly revenues from this business' activities will be higher than (max-min)/2? We use these three questions to recover the owner's distribution of expectations, assuming a triangular distribution (Attanasio and Kaufmann, 2014). We follow the same approach for epectations regarding the number of customers, which are measured over a seven days period. In Panel B, the firm-level distribution of monthly revenues is calculated from data on the firm's monthly revenues in the three months prior to the survey. All variables are winsorized at the top 1%. Monteray amounts are deflated at August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD.

Table 5: Input Responsiveness to Demand Shocks

Sample: All firms at all follow-ups, monthly and follow-up panels
IV regression coefficients, robust standard errors in parentheses

	Log (Total Value of Assets)	Log (Number of Employees)	Log (Total Value of Assets)	Log (Number of Employees)	Log (Labor Utilization)	Log (Hours in production)	Log (Hours at the premises)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log (Revenues)	.008	-.003	.148**	-.018	.829***	.734**	-0.032
	(.009)	(.016)	(.066)	(.090)	(.368)	(.327)	(0.073)
First stage coefficients:							
Mean log(revenues), other firms	.508***	.504***	.284***	.277***	.277***	.277***	.277***
	(.020)	(.019)	(.069)	(.068)	(.068)	(.068)	(.068)
First stage F-stat	109	109	16.9	16.8	16.8	16.8	16.8
Mean outcome for control firms	6.56	1.05	6.77	1.30	-1.77	3.87	5.57
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of observations	7,743	7,847	1,275	1,296	1,296	1,296	1,296

Notes: The data is from a monthly panel of firms in Columns 1 and 2 and a panel of firms at follow-up restricted to the sample for which labor utilization information is available in Columns 3 to 7. For each firm, monthly revenues are instrumented by the average value of monthly revenues among firms in the same sector, branch, and month, excluding the firm itself. All regressions include firm fixed effects. *Labor utilization* is measured as the number of monthly hours actually worked over the number of monthly hours the worker spent at the firm premises. *Mothly hours at firm premises* are computed multipling the average number of hours by the number of days the worker report being at the firm in a typical day. *Monthly hours actually worked* is calculated as follows. For each workers paid piece-rate, we ask the following three questions regarding a specific product/service: (1) How much time (minutes) would <name> normally take right now to complete his/her work on this product/service alone to a standard of quality considered satisfactory by the average customer of this business? (2) How much of the price charged to customers for this product/service does <name> normally receive? In addition, we have data on the total earnings of the worker in the month prior to the survey. We calculate the number of products/services the worker has worked on in the last month by dividing total earnings by the piece-rate. We then calculate monthly hours actually worked as the number of products/services the workers has worked on in the last month times the average time taken to complete the product/services. Monthly hours worked can only be calculatd available from second follow-up and for workers paid per-piece (67% of all workers). ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level.

Table 6: Demand Uncertainty and Labor Utilization

Sample: Workers panel at first follow-up all firms (Columns 1, 3, and 5), all follow-ups (all other Columns)
OLS regression coefficients, robust standard errors in parentheses

	Labor Utilization		Hours spent in Production (workers)		Hours at premises (workers)	
	OLS	IV, same trade and branch	OLS	IV, same trade and branch	OLS	IV, same trade and branch
	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient of Variation, *Expected* Monthly Revenues	-.262***		-88.2***		-26.1	
	(.086)		(25.7)		(30.8)	
Coefficient of Variation, *Current* Monthly Revenues		-.567***		-157.6***		-3.87
		(.103)		(28.4)		(18.7)
First Stage Coefficients:						
Mean CV of monthly revenues, other firms		.524***		.524***		.524***
		(.083)		(.083)		(.083)
First stage F-stat		40.0		40.0		40.0
Mean Outcome for Controls	.120	.281	35.2	74.0	297	274
Survey wave fixed effects	-	No	-	No	Yes	No
Sector and location fixed effects	Yes	No	Yes	No	Yes	No
Firm fixed effects	No	Yes	No	Yes	No	Yes
Number of observations	527	2,985	527	2,985	527	2,985

Notes: The data is from a monthly panel of worker. The table shows the results from a regression of labor utilization. the number of hours the worker spent in production, and the number of hours the worker spent at the firm premises in the last month on the coefficient of variation of expected and current monthly revenues.The coefficient of variation of expected monthly revenues is only measured at first follow-up, and so the sample in Columns 1, 3, and 5 only includes firms at first follow-up. In Columns 2, 4, and 6, for each firm, the coefficient of variation of firms' monthly revenues is instrumented by the average coefficient of variation of monthly revenues among firms in the same sector, branch, and survey round excluding the firm itself. *Labor utilization* is measured as the number of monthly hours actually worked over the number of monthly hours the worker spent at the firm premises. *Mothly hours at firm premises* are computed multipling the average number of hours by the number of days the worker report being at the firm in a typical day. *Monthly hours actually worked* is calculated as follows. For each workers paid piece-rate, we ask the following three questions regarding a specific product/service: (1) How much time (minutes) would <name> normally take right now to complete his/her work on this product/service alone to a standard of quality considered satisfactory by the average customer of this business? (2) How much of the price charged to customers for this product/service does <name> normally receive? In addition, we have data on the total earnings of the worker in the month prior to the survey. We calculate the number of products/services the worker has worked on in the last month by dividing total earnings by the piece-rate. We then calculate monthly hours actually worked as the number of products/services the workers has worked on in the last month times the average time taken to complete the product/services. Monthly hours actually worked can only be calculatd available for workers paid per-piece (67% of all workers). The *Coefficient of variation of the firm's expected revenues* is calculated as follows. Each firm owner was asked the following questions: (1) Imagine that over the next 6 months the revenues from this business are very high. What is the maximum amount of revenues the ownercould make in the best month over the next 6 months? (2) Imagine that over the next 6 months the revenues from this business are very low. What is the minimum amount of revenues the

Table 7: Demand Uncertainty and Training

Sample: Firm panel at first follow-up (Columns 1, 3), all follow-ups (all other Columns)
OLS regression coefficients, robust standard errors in parentheses

	Hours spent training (Owner)		Hours at premises (Owner)	
	OLS	IV, same trade and branch	OLS	IV, same trade and branch
	(1)	(2)	(3)	(4)
PANEL A: Demand Uncertainty				
Coefficient of Variation, *Expected* Monthly Revenues	-174***		-48.2	
	(38.6)		(30.0)	
Coefficient of Variation, *Current* Monthly Revenues		-61.0*		1.66
		(35.9)		(36.5)
First Stage Coefficients:				
Mean CV of monthly revenues, other firms		.451***		.433***
		(.048)		(.053)
First stage F-stat		88.6		65.7
Mean Outcome for Controls	58.8	51.5	343	327
Survey wave fixed effects	-	No	-	No
Sector and location fixed effects	Yes	No	Yes	No
Firm fixed effects	No	Yes	No	Yes
Number of observations	575	2,331	510	1,926

Notes: The data is from a monthly panel of firms. The table shows the results from a regression of the number of hours the firm owner spent training and the number of hours they spent at the firm premises in the last month on the coefficient of variation of expected and current monthly revenues. The *coefficient of variation of expected monthly revenues* is only measured at first follow-up, and so the sample in Columns 1 and 3 only includes firms at first follow-up. In Columns 2 and 4, for each firm, the coefficient of variation of firms' monthly revenues is instrumented by the average coefficient of variation of monthly revenues among firms in the same sector, branch, and survey round excluding the firm itself. *Hours at the premises in the last month (owner)* are computed multipling the average number of hours by six, the average number of days workers report being at the firm in a typical day, times four weeks. The *coefficient of variation of the firm's expected revenues* is calculated as follows. Each firm owner was asked the following questions: (1) Imagine that over the next 6 months the revenues from this business are very high. What is the maximum amount of revenues the ownercould make in the best month over the next 6 months? (2) Imagine that over the next 6 months the revenues from this business are very low. What is the minimum amount of revenues the owner could make in the worst month over the next 6 months? (3) In how many months over the next 6 months, does the

Table 8: Parameter Estimates and Model Fit

Parameter	Value	Targeted Moment	Data	Model
Estimated				
Cross-section:				
α	0.359	Labor share	0.329	0.329
θ	7.747	Average markup	2.700	3.044
χ	1.571	Share of trained workers	0.497	0.497
σ_v	3.463	Labor utilization	0.569	0.569
γ_1	1.39	Time supplied by workers	0.442	0.441
w	2.213	Number of workers	2.242	2.242
A	12.06	Number of machines	2.306	2.306
Experiment:				
ϕ	0.889	LATE on firm revenues	0.319	0.319
γ_2	1.91	Time supplied | Wage subsidy	0.397	0.416
Calibrated				
r	0.514	Hourly rent (Bassi *et al.* 2022)	0.514	

Table: The table report the estimated model parameter values, together with the moments targeted by the model estimation. For each parameter, the table indicates the moment most closely associated with its identification. The last two columns report the value of each moment observed in the data and its corresponding model-implied estimate.

Table 9: Treatment Effects, Data vs. Model

Sample: All firms at first follow-up
Changes related to controls

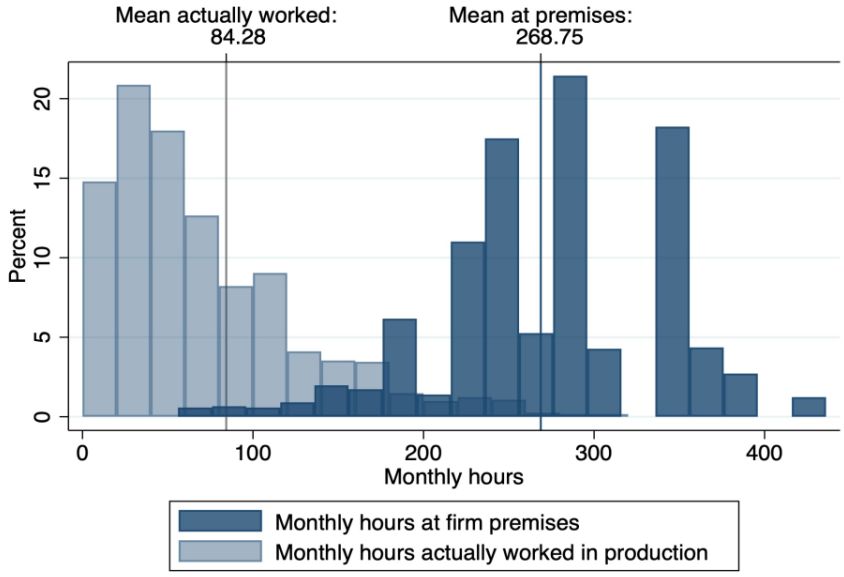
	Number of Workers	Number of Skilled Workers	Log Total Value of Assets	Utilization	Time at Firm Premises	Log Monthly Sales	Log Monthly Profits
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
PANEL A: Data							
Local Average Treatment Effect (LATE), *Levels*	1.21***	1.06***	.037	-.009	-32.4***	.319**	.100
	(.262)	(.270)	(.106)	(.029)	(8.94)	(.155)	(.133)
Local Average Treatment Effect (LATE), *% changes*	.539***	.477***	.037	-.088	-.115***	.319**	.100
	(.117)	(.122)	(.106)	(.268)	(.032)	(.155)	(.133)
PANEL B: Model, *% changes*							
Simulated Treatment Effect (k Fixed)	.446	.491	.000	-.043	-.047	.194	.133
Simulated Treatment Effect (k Variable)	.446	.482	.154	-.051	-.056	.319	.240

Notes: The data is from the first follow-up. The table shows treatment effects for the wage subsidy intervention in the data (Panel A) and the model (Panel B). In Panel A, Local Average Treatment Effects are estimated instrumenting the take-up of the wage subsidy intervention with treatment assigment. For percentage changes, estimates for the following variables are from a regression that uses the logarithm of the variable as an output: total value of assets, monthly sales, and monthly profits. Panel B shows results from two separate simulations of the wage subsidy intervention in the model: one where we hold capital to the same level as control group (k fixed), and one in which we allow firms to reoptimize capital (k variable). ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level.

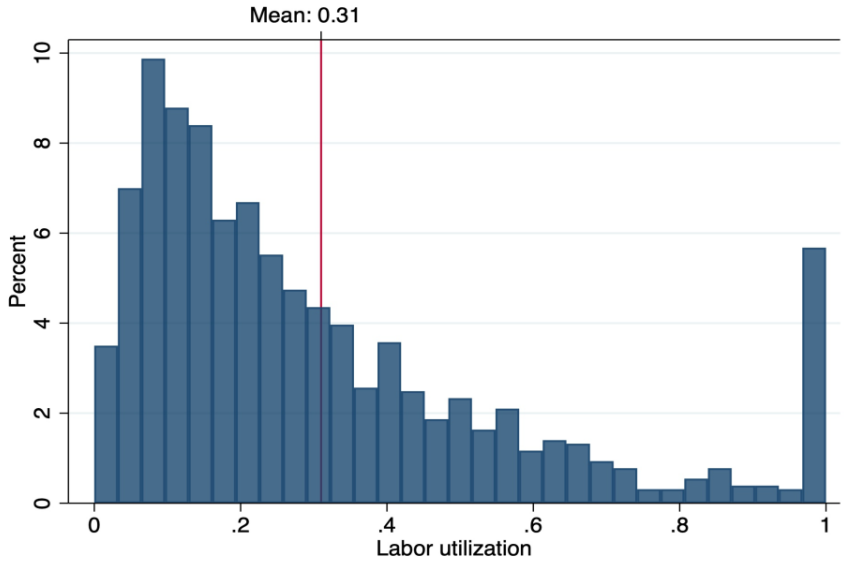
Figure 1: Labor Utilization

Sample: Worker panel from control firms, follow-ups 2-3-4

A: Hours Worked and Hours at Business Premises



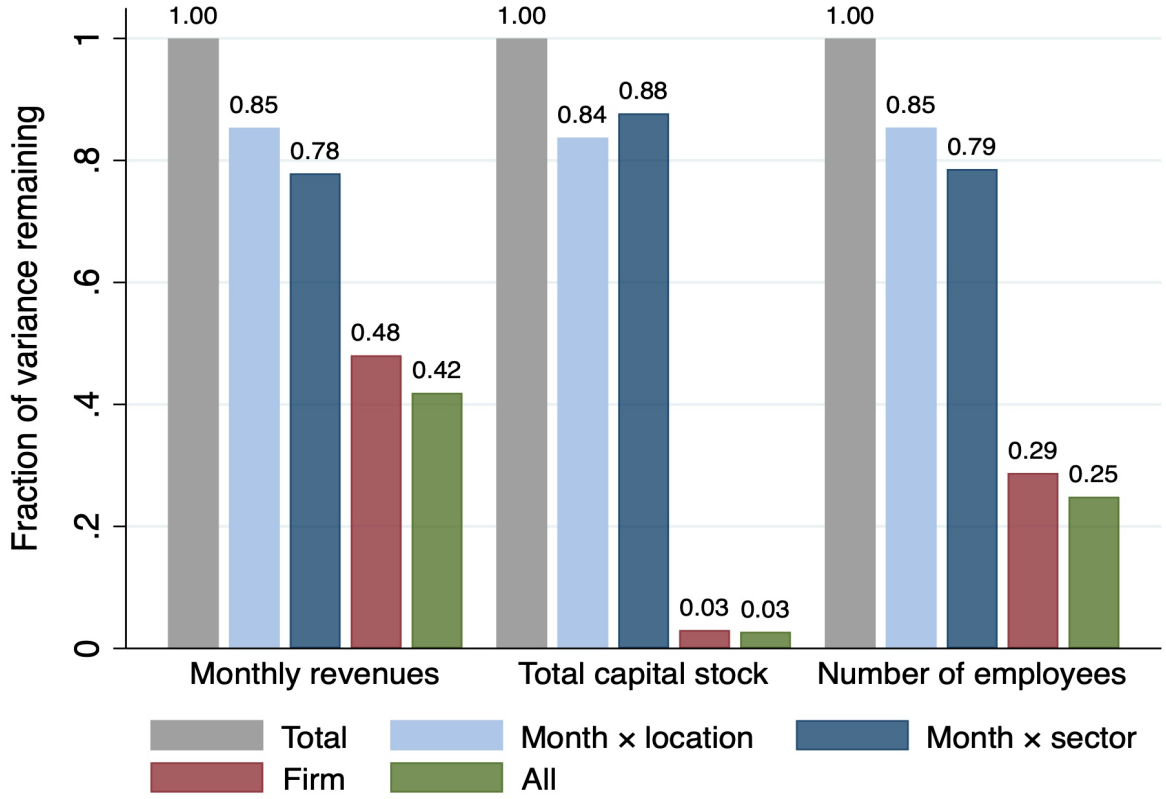
B: Labor Utilization Rate



Note: The figures use worker-level data from the Control group at second, third and fourth follow-ups. Panel A shows the distribution of the monthly number of hours workers typically spend at the firm premises (in dark blue) and the monthly number of hours actually worked in production (in light blue). *Hours actually worked* are calculated as follows. For each workers paid piece-rate, we ask the following three questions regarding a specific product/service: (1) How much time (minutes) would <name> normally take right now to complete his/her work on this product/service alone to a standard of quality considered satisfactory by the average customer of this business? (2) How much of the price charged to customers for this product/service does <name> normally receive? In addition, we have data on the total earnings of the worker in the month prior to the survey. We calculate the number of products/services the worker has worked on in the last month by dividing total earnings by the piece-rate. We then calculate monthly hours actually worked as the number of products/services the workers has worked on in the last month times the average time taken to complete the product/services, which therefore only reflects the hours spent on production. *Hours at the premises* is calculated using information on: (i) the number of hours the worker usually spends on the firm premises per day, (ii) the number of days the worker spends on the premises per week. The total number of hours spent on the premises in a given month is then calculated as the product of these two values, multiplied by four. Panel B shows labor utilization, computed as share of hours at the firm premises spent in productive activities.

Figure 2: Residual Demand Uncertainty

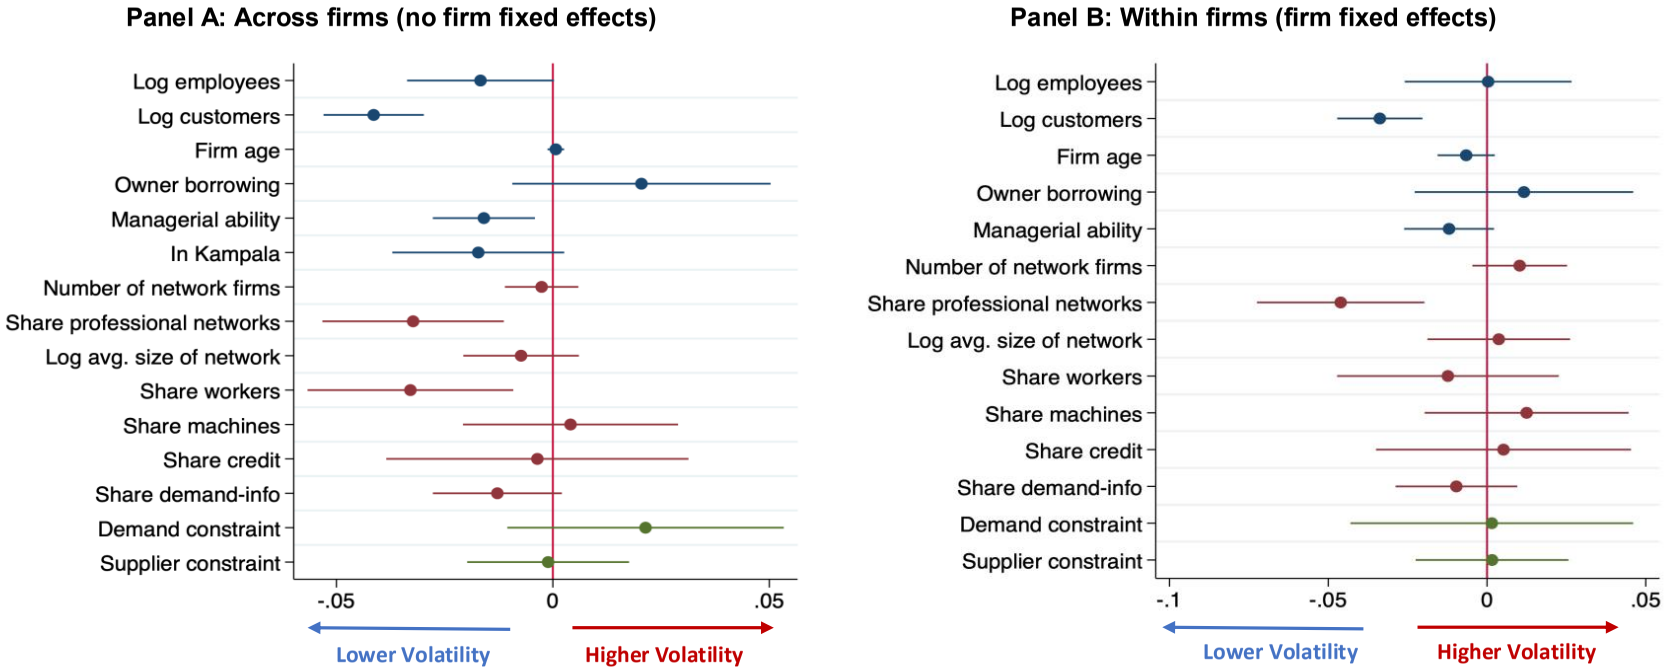
Sample: Panel of Control firms, all follow-ups



Notes: The figure uses monthly data from the Control group at all follow-ups to show the share of variation in revenues, number of workers and value of capital that is predicted by month by branch, month by sector, and firm fixed effects. For each outcome, the initial variance is normalized to one. Each bar shows the share of the initial variance remaining after controlling for each fixed effect. In the second, third, and fourth bar, regressions only include the fixed effect indicated in the legend. In the last bar (All), the remaining variance is obtained from a regression where we include all fixed effects.

Figure 3: Correlates of Demand Volatility

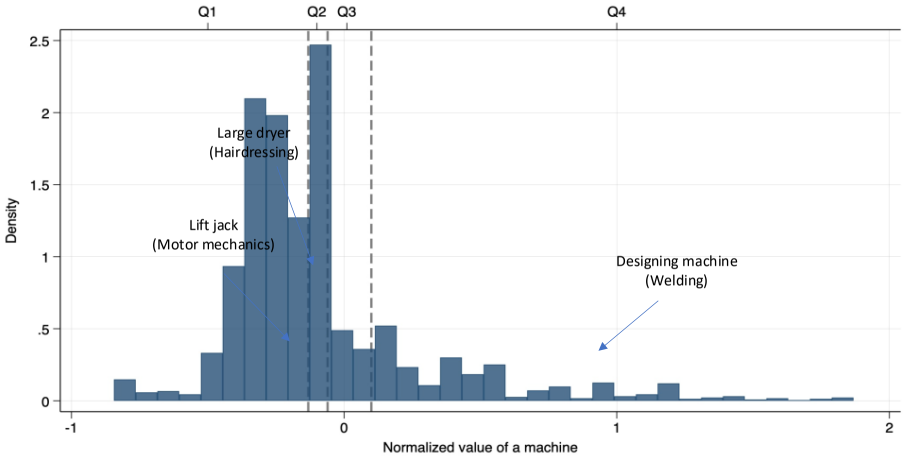
Sample: All firms, all follow-ups



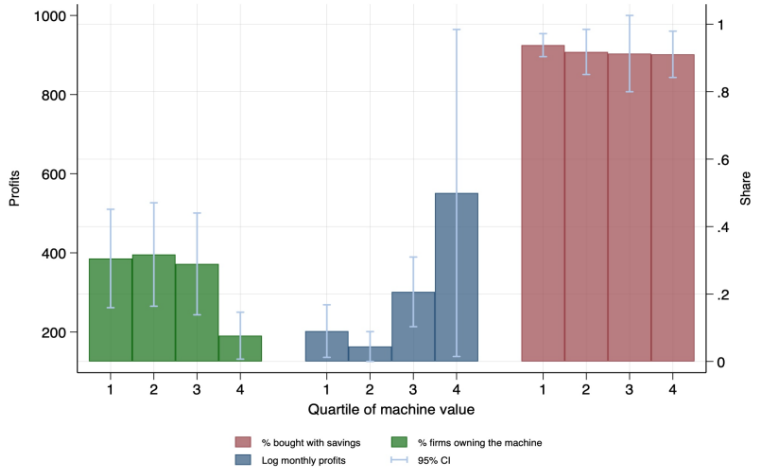
Notes: The figures show coefficients and 95% confidence intervals from regressions of the coefficient of variation of firm monthly sales on the covariate of interst, one at the time. These covariates span three categories: (i) firm characteristics, (ii) firm networks, (iii) self-reported constraints. Panel A shows coefficients from a regression that includes sector, branch, survey round, and treatment fixed effects. Panel B shows coefficients from a regression that only includes firm fixed effects. *Log customers* is the logarithm of the number of customers who made purchases/orders from the business in the 7 working days prior to the survey. *Managerial ability* is measured using an standardized index of managerial practices constructed from a battery of questions on record keeping, marketing, stock control, business performance review, and financial planning. Network firms are identified as other businesses in the same area with whom the firm owner shares something (such as inputs, machines, workers, information, credit, advice, etc.) or meets to talk. Professional ties are businesses with whom the owner reports having a business relation (vs. family members, friends, and neighbors). *Supplier constraint* is a dummy equal to one if the firm reported it would hire more workers or make new investments if either it was easier to find suppliers that have available all the materials needed by this business or it was easier to find suppliers that are reliable and always deliver materials in the quantity, quality, timing and price agreed upon. *Demand constraint* is a dummy equal to one if the firm reported it would hire more workers or make new investments if people would buy more in the place where this business sells.

Figure 4: Machines and Perceived Constraints to Adjusting Inputs

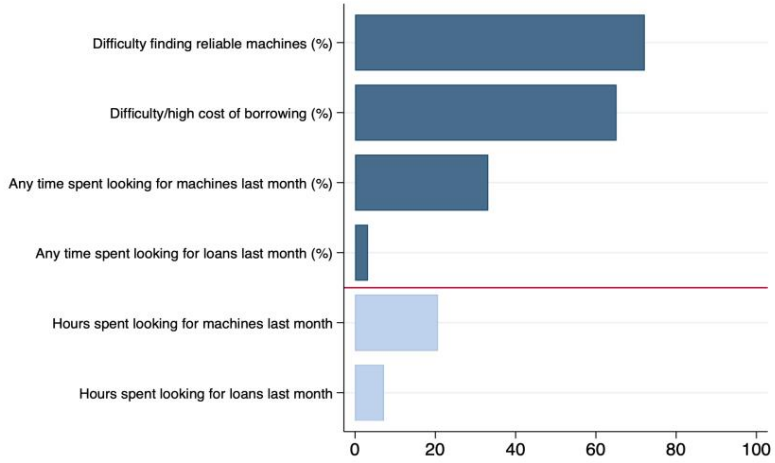
A: Machine Costs



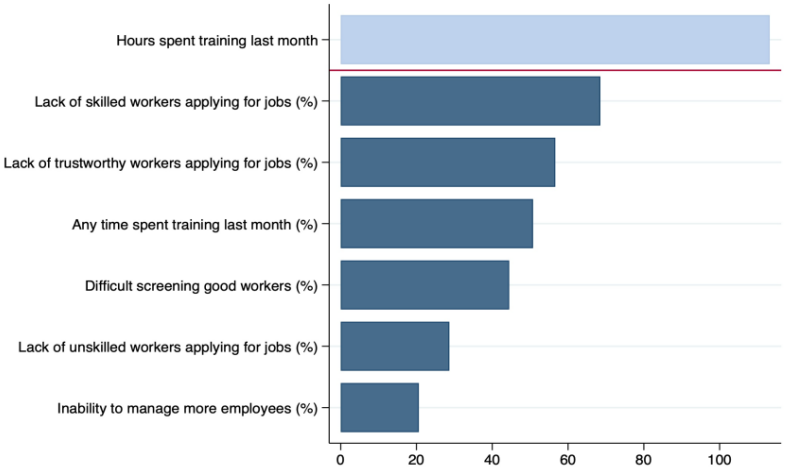
B: Profits, Ownership and Financing of Machines, by Quartile of Machine Cost



C. Constraints to Adjusting Capital

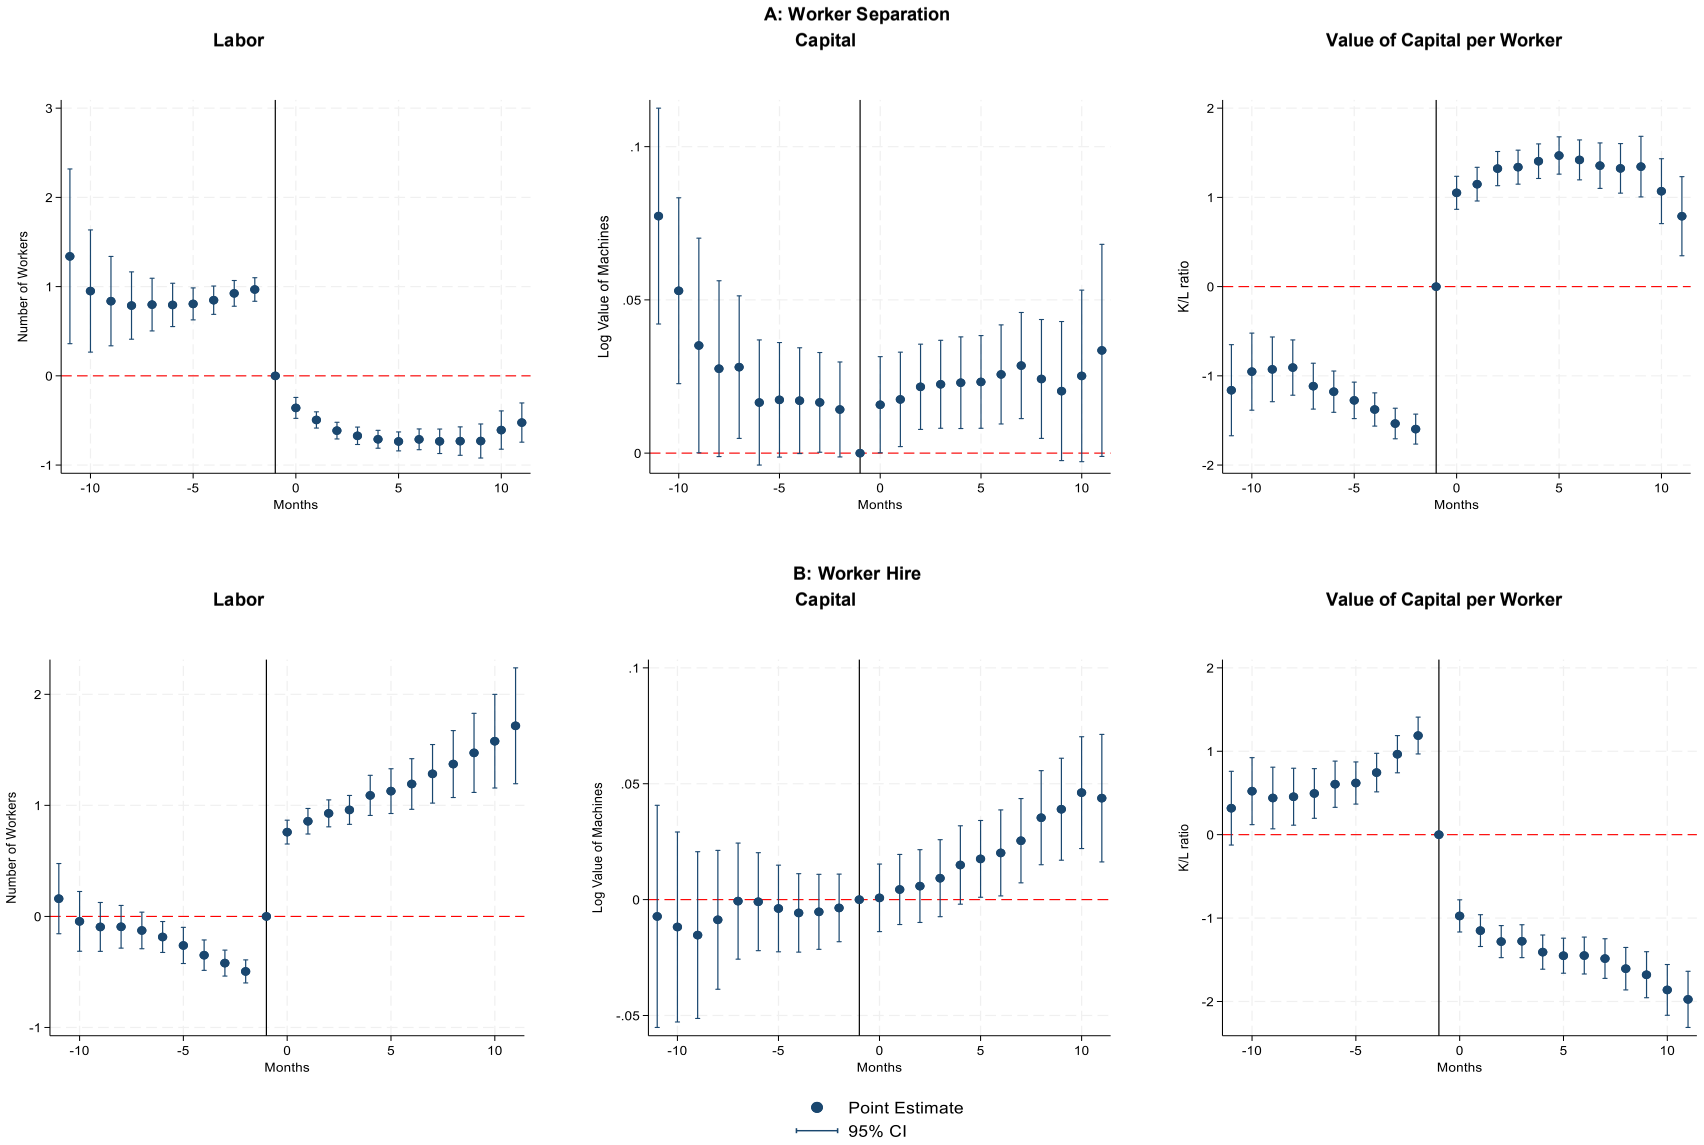


D. Constraints to Adjusting Labor



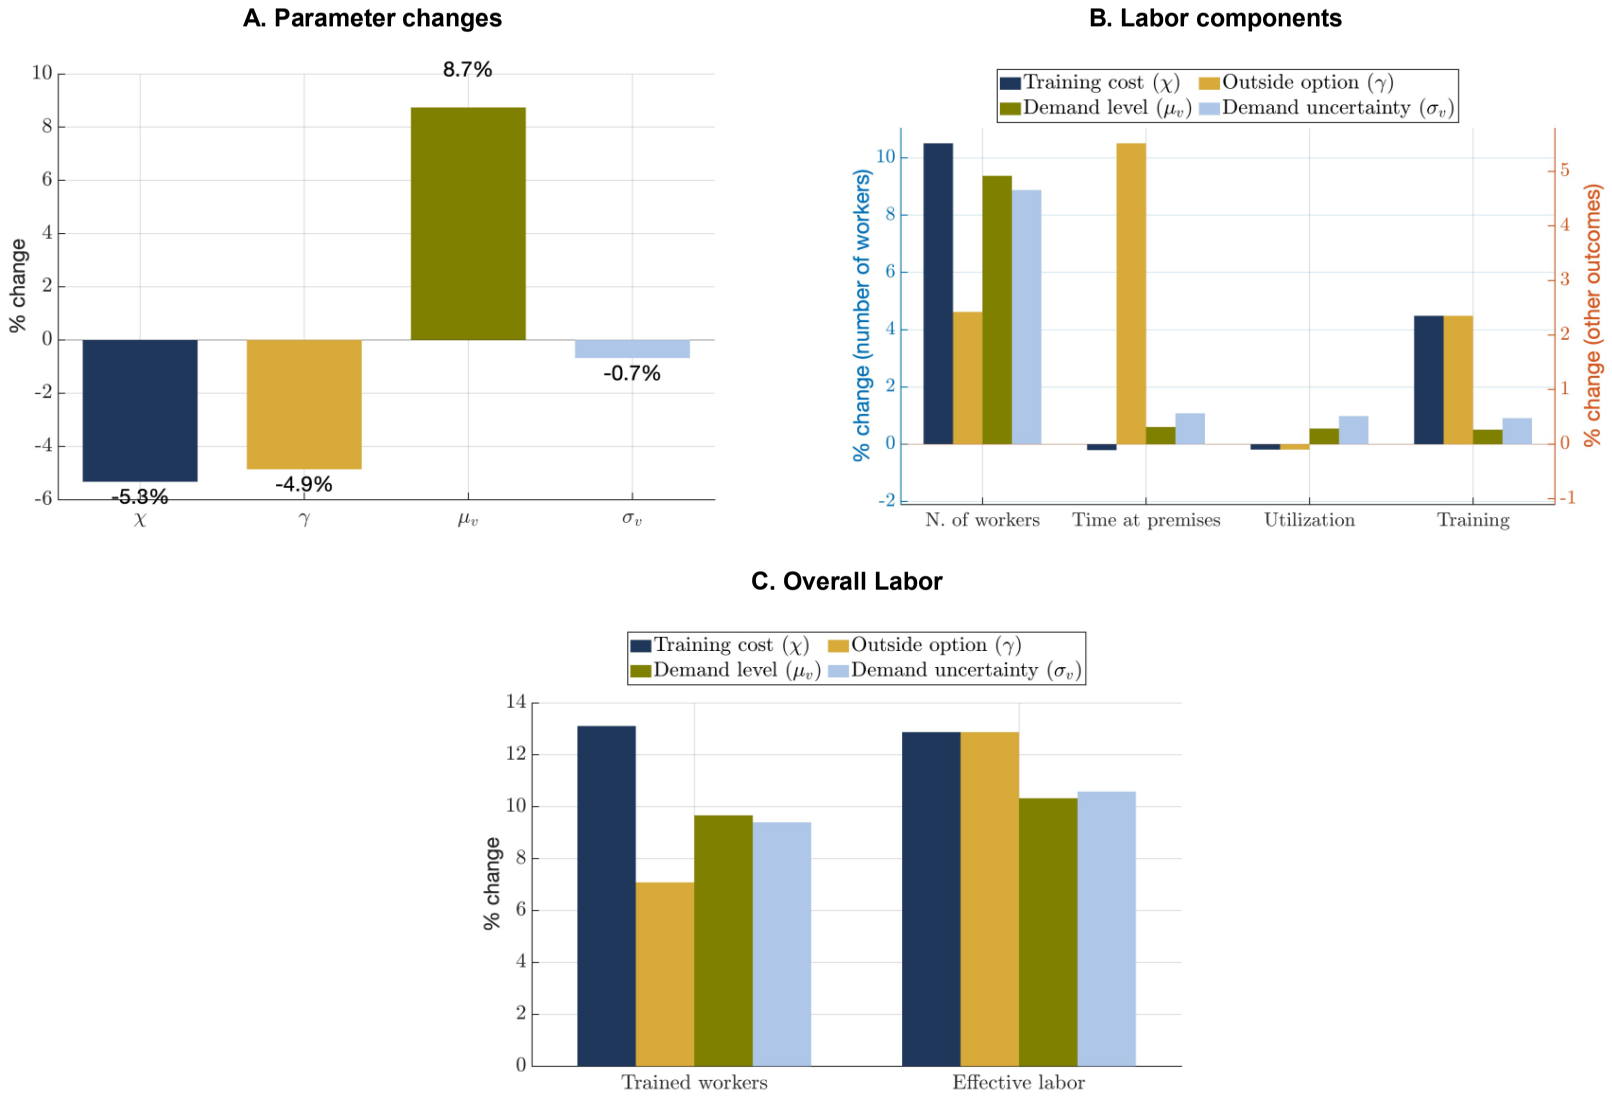
Notes: Panel A plots the distribution of machine values at baseline. The unit of observation is a specific machine owned by a given firm. Machine values are normalized by the mean and standard deviation of machine values in the sector in which the firm operates. In Panel B, we first calculate the average (normalized) value of a given machine type (e.g. lift jacks) across all firms who have access to that machine. We then rank the machines according to their average value and assign them to quartiles. Finally, we calculate average profits, share of ownership and share of machies boughts with savings/reinvested profits among firms who own a machine in a given quartile of the machine value distribution. The dotted lines in Panel A indicate the thresholds corresponding to the quartiles of machine values used in Panel B. Panel C and D use data from all firms at first follow-up (time variables) and at baseline (self-reported constraints related to machines and workers). All variables refer to the month prior the survey. The number of hours spent training, looking for machines and looking for loans in the last month are conditional on the owner having spent at least one day training, looking for machines or looking for loans in the last month.

Figure 5: Worker Separation and Hire Event Studies



Notes: The figures show event studies for the hiring of a new worker (Panel A) and the firing or loss of a worker (Panel B). The event is the hiring of a new worker (Panel A), and the firing or loss of a worker (Panel B) in the 12 month before a given follow-up survey. The graphs show the results of the analysis pooled across all follow-ups. The outcomes are the number of workers employed at the firm, including the firm owner (Labor), the logarithm of the firm's total value of assets (Capital) and the logarithm of the firm's total value of assets over the number of workers employed at the firm, including the firm owner (K/L ratio).

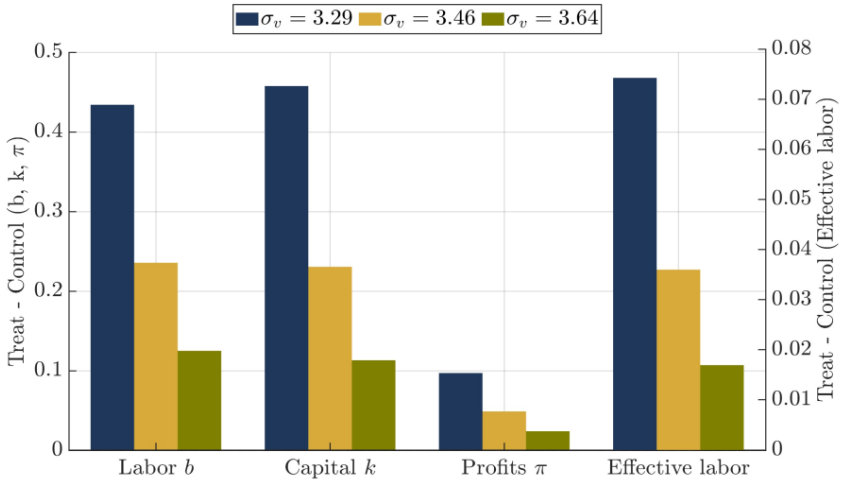
Figure 6: Counterfactual Effect of Reducing Training Frictions



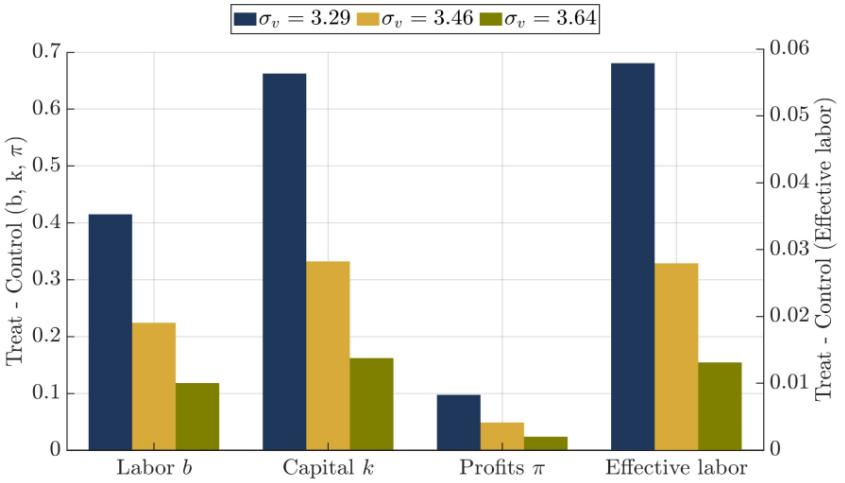
Notes: The figure shows changes in worker and firm outcomes relative to the baseline estimation in counterfactual simulations where we change each parameter (chi, sigma, gamma, and mu) to induce a 10% increase in firm sales. Panel A shows the magnitude of the parameter changes needed to induce such increase in sales. Panel C shows the changes in the number of trained workers (b X tau) and effective labor (b X tau X h X l) induced by changes in each of the parameters.

Figure 7: Counterfactual Effects of Policies under Uncertainty

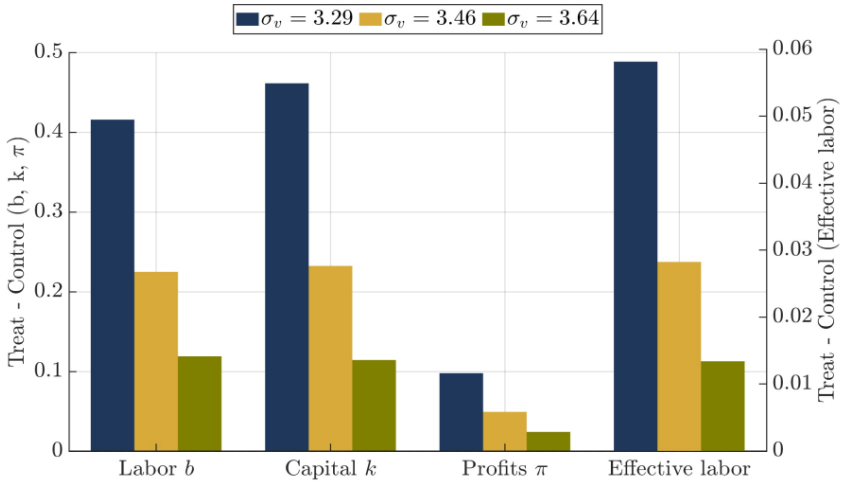
Panel A: Training Cost Reduction



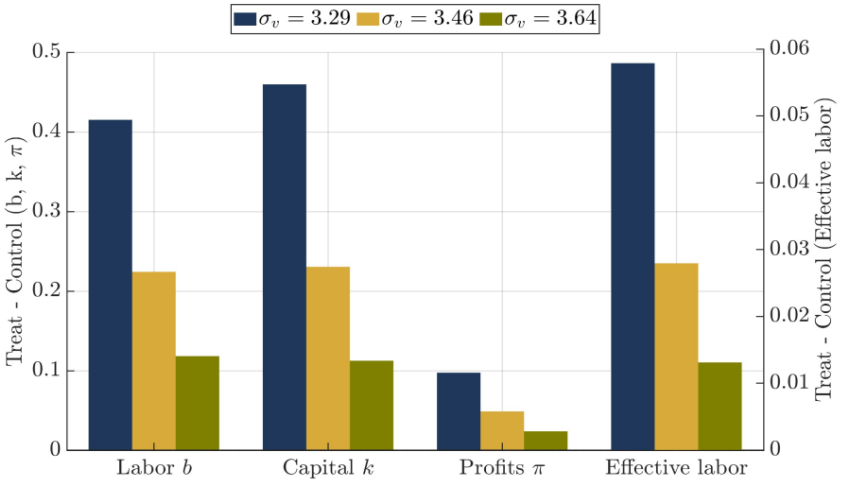
Panel B: Rental Cost Reduction



Panel C: Improved Market Access

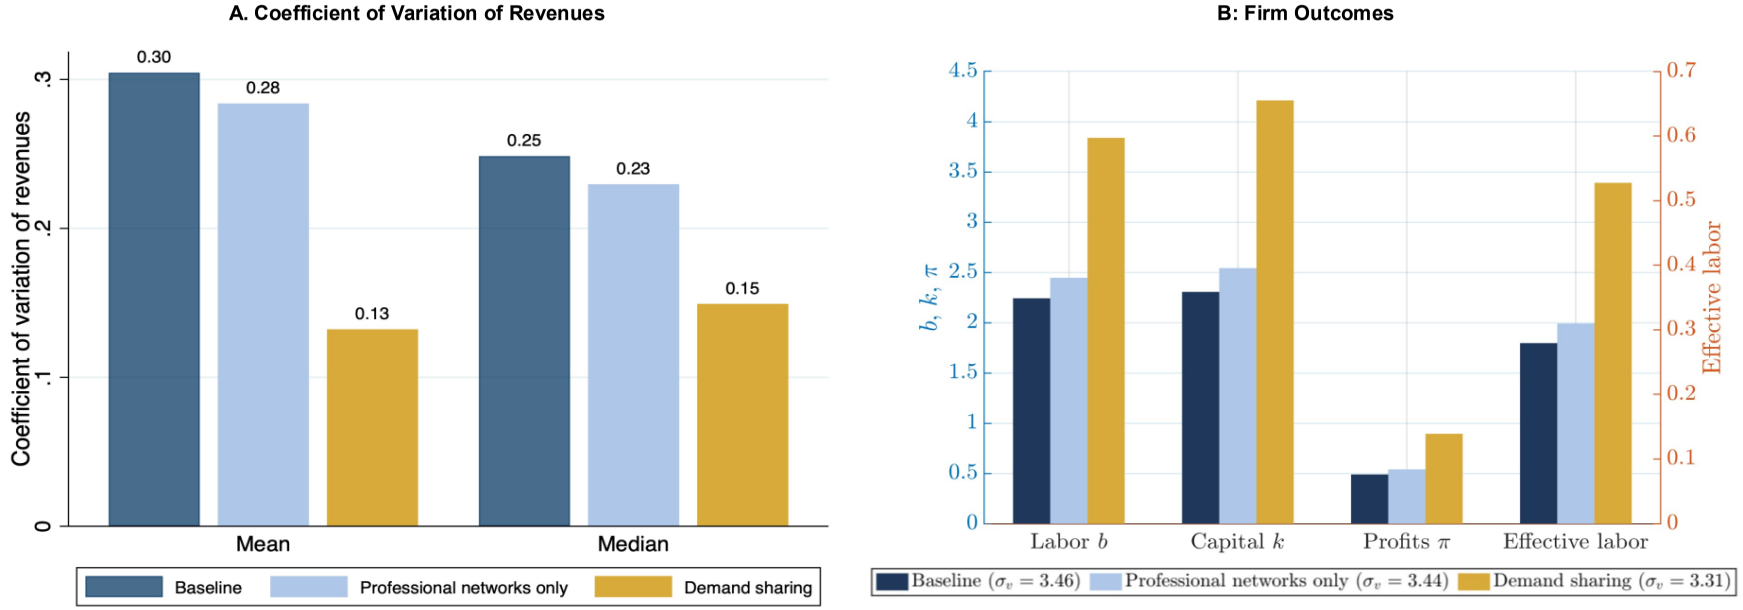


Panel D: Management Training



Notes: The figure shows results from simulatations of different types of policy interventions across firms facing heterogeneous levels of demand uncertainty (sigma_v). In Panel A, the figure shows changes in outcomes induced by the a reduction in training costs. In Panel B, it shows outcomes for firms facing lower rental prices (r). In Panel C, it shows outcomes for firms facving a higher mean demand level (mu_v). In Panel D, it shows outcomes for firms with a higher TFP (A). Policy impacts given the baseline level of volatility (3.46) are shown in yellow. Blue and green bars show policy impacts in a settimg with 5% lower and 5% higher sigma_v respectively.

Figure 8: Counterfactual Effects of Sharing Workers and Demand



Notes: The figure simulates the impact of reducing demand volatility (i) by assuming all existing networks between firms are professional networks, and (ii) by aggregating the sales of firms operating within the same sector and branch. Panel A shows the mean and median coefficient of variation of sales of firms in the baseline scenario vs. the counterfactual scenarios of professional networks only and demand sharing. Panel B shows average inputs (labor, capital, effective labor) and profits at the current level of sales volatility vs. the level of volatility firms would face if they only had professional networks or shared demand.

Table A1: Attrition of Firms

OLS regression coefficients, robust standard errors in parentheses, and F-tests in brackets

Dependent Variable: Firm attrited or closed in each follow-up

Attrited by:	1st follow-up	2nd follow-up	3rd follow-up	4th follow-up
	(1)	(2)	(3)	(4)
Monthly profits (log)	-.042**	-.013	-.006	-.014
	(.017)	(.016)	(.017)	(.021)
Number of employees	-.013**	-.014**	-.012**	-.016**
	(.006)	(.006)	(.006)	(.008)
Number of unskilled employees	.018*	.014	.008	.001
	(.010)	(.009)	(.009)	(.011)
Capital stock (log)	.010	-.007	.001	.003
	(.011)	(.011)	(.012)	(.014)
Firm's age	-.004	-.005*	-.009***	-.014***
	(.003)	(.003)	(.003)	(.003)
Attrition Rate	.160	.181	.186	.322
Owner and firm characteristics	Yes	Yes	Yes	Yes
Sector and branch fixed effects	Yes	Yes	Yes	Yes
F-test on firm characteristics	[.005]	[.041]	[.007]	[.000]
F-test on owner characteristics	[.678]	[.414]	[.200]	[.623]
Number of observations	1,025	1,025	1,025	1,025

Notes: The table shows results from a regression of attrition on firm characteristics. Firms that were found shut-down were not interviewed and are here considered attriting. The two stratification variables used for the randomisation of firms at baseline are: trade and BRAC branch. All regressions include trade and branch fixed effects. All covariates are measured at baseline. The top 1% of profit and capital values are trimmed. All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Missing values are replaced with the sample average and a dummy for missing value for each variable is included in all regressions. ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level.

Table A2: Owner's Time Use

Sample: All firm owners at baseline (Panel A), Control firm owners at first follow-up (Panel B)

	% Owners	Hours/day
A. Days spent at the premises & hours worked		
Every day	85.0	10.8
		(2.56)
Several times a week	10.8	8.91
		(2.80)
Several times a month	2.39	8.25
		(3.18)
Once a month or less	1.79	5.59
		(4.39)
Average		10.4
		(2.84)

	% Owners	Hours/30 days	Hours/30 days | Spending any time on it
B. Time use at work			
Training	44.5	50.3	113
		(87.8)	(101)
Assisting workers	60.5	63.7	105
		(95.2)	(103)
Looking for workers	5.74	1.32	22.9
		(17.6)	(71.3)
Looking for machines	33.2	6.84	20.6
		(22.8)	(36.0)
Looking for loans	3.21	.230	7.14
		(1.82)	(7.59)
Interacting with government	20.2	1.56	7.73
		(5.72)	(10.7)

Notes: The data is from all firm owners at baseline in Panel A, and firm owners in control firms at first follow-up on Panel B. In Panel A, the variable hours/day is winsored at the top1% (among non-0 values). In Panel B, the hours/30 days variable is calculated as days spent doing a given activity in the last 30 days*hours/day spent doing that activity on a normal day. The list of activities were pre-specified.

Table A3: External Validity

Mean, standard deviation in parentheses

P-value on t-test of equality of means between baseline and GEM samples in brackets

	Number of firms	Owner gender [Female=1]	Owner age	Owner education	Owner started the business [Yes=1]	Number of owners	Firm Age [Years]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
2012 Baseline Sample	2306	.532	34.5	2.48	.949	1.20	6.58
		(.499)	(7.49)	(.765)	(.220)	(2.43)	(4.76)
2010 GEM Uganda Sample	762	.545	34.5	2.47	.934	1.43	6.36
		(.498)	(11.0)	(.801)	(.248)	(2.17)	(7.69)
p-value		[.534]	[.943]	[.773]	[.123]	[.027]	[.359]

Notes: The table reports summary statistics on basic owner and firm characteristics from baseline data and a comparison between firms at baseline and firms in the 2010 Adult Population Survey for Uganda conducted by the Global Entrepreneurship Monitor (GEM) project. In the baseline sample, the Owner education variable measures years of education. Education levels are instead coded in the GEM data using 1997 UN ISCED codes. These are: 2=Up to lower secondary education, 3=Upper secondary education, 4=Post-secondary non-tertiary education, 5=First stage of tertiary education, 6=Second stage of tertiary education. For the comparison with the GEM sample, the Owner education variable is recoded in the sample of firms from baseline to match 1997 UN ISCED education codes. Education levels in baseline data are translated into 1997 UN ISCED groups using the conversion table provided here: http://www.uis.unesco.org/Education/ISCEDMappings/Pages/default.aspx. The first two rows in columns (2) to (7) report means and standard deviations of each of the variables in the column headings for the baseline sample. Due to missing values, the actual sample sizes for the descriptives in the first and second row are the following (numbers in parentheses refers to column number): (2) 1,794, (3) 1,738, (4) 1,772, (5) 1,773, (6) 1,789, (7) 1,753. The fourth and fifth rows in columns (2) to (7) report means and standard deviations for the corresponding variables (when available) among the entire sample of firms included in the GEM sample (the sample here inlcudes all sectors, apart from agriculture/animal rearing). The sample is restricted to firms employing between 0-25 employees. Due to missing values, the actual sample sizes for the descriptives in the fourth and fifth row are the following (number in parentheses refers to column number): (2) 762, (3) 758, (4) 762, (5) 761, (6) 755, (7) 736. The p-values in brackets are from the t-test of equality of means between firms in GEM sample and firms in the baseline sample. ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level.

Table A4: Measuring Labor Utilization

Sample Column 1: control firms at follow-ups 2-3-4

Sample Columns 2 and 3: workers at control firms at follow-ups 2-3-4

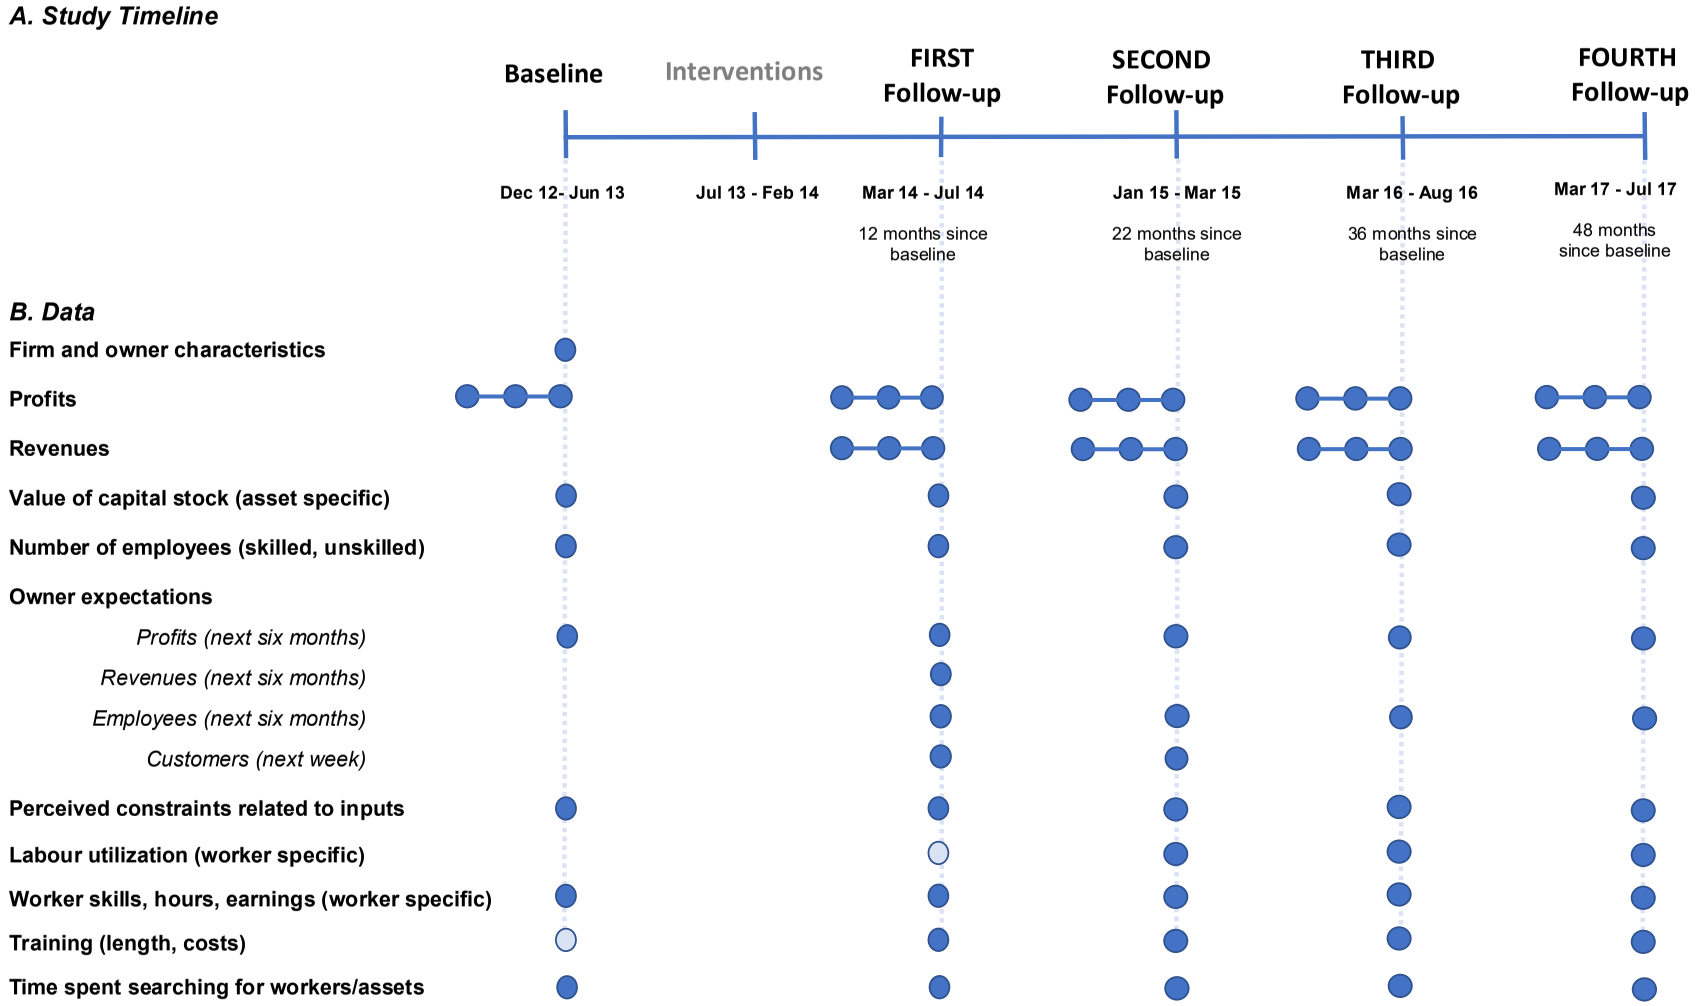
Sample Columns 4 and 5: workers with labor utilization measured at control firms at follow-ups 2-3-4

OLS regression coefficients, robust standard errors in parentheses

	Labor Utilization Measure Can be Constructed			Labor Utilization	
	(1) Firms	(2) Workers	(3) Workers and Firm Fixed Effects	(4) Workers	(5) Workers
Monthly profits (log)	.016				
	(.021)				
Number of employees	.016**				
	(.008)				
Number of unskilled employees	-.086***				
	(.015)				
Capital stock (log)	.010				
	(.015)				
Firm's age	.004				
	(.004)				
Age		.001	-.002	.002	.005
		(.002)	(.003)	(.003)	(.004)
Sex (female)		-.034	-.011	.018	.032
		(.035)	(.046)	(.044)	(.077)
Tenure (months)		.001	.001**	-.001	-.002**
		(.000)	(.001)	(.001)	(.001)
Primary school		.102*	093	030	044
		(052)	(065)	(048)	(075)
Junior Secondary school		072	069	033	032
		(051)	(065)	(046)	(076)
High Secondary school		056	050	010	010
		(061)	(074)	(059)	(084)
Certificate or higher		- 006	021	041	096
		(059)	(079)	(054)	(098)
Skilled		.327***	.336***	.035	.021
		(.033)	(.037)	(.037)	(.051)
Owner's family member		.020	.035	-.003	-.014
		(.030)	(.035)	(.031)	(.037)
Productivity (1-10)		.031***	.021***	-.006	-.002
		(.006)	(.006)	(.006)	(.008)
Received training		-.007	-.022	.013	.039
		(.027)	(.033)	(.025)	(.036)
Mean outcome	59 9	475	475	310	310
Survey wave fixed effects	Yes	Yes	Yes	Yes	Yes
Location fixed effects	Yes	Yes	No	Yes	No
Sector fixed effects	Yes	Yes	No	Yes	No
Other Firm Characteristics	Yes	-	-	-	-
Firm Fixed Effects	-	No	Yes	No	Yes
Number of firms	368	361	361	256	256
Number of observations	867	1788	1788	849	849

Notes: The unit of observation is the control firms (Column 1) and workers in control firms (cols 2-3-4-5) at follow-ups 2-3-4. All regressions include survey wave FE. Columns 1, 2, 4 include location (branch) and sector FEs. Columns 3 and 5 inlcude firms FE. In column 1 missing values for age, capital, firm's licence, profits, and num. of network firms are replaced with 0 and a dummy for missing values is included for each of the variables. The dependent variable in columns 1-2-3 is a dummy that equals 1 if the firm/worker have nonmissing value for labor utilization. The dependent variable in columns 4-5 is labor utilization. Labor utilization is computed as the share of hours workers typically spend at the firm during which they actually worked. Hours actually worked is calculated as follows. For each workers paid piece-rate, we ask the following three questions regarding a specific product/service: (1) How much time (minutes) would <name> normally take right now to complete his/her work on this product/service alone to a standard of quality considered satisfactory by the average customer of this business? (2) How much of the price charged to customers for this product/service does <name> normally receive? In addition, we have data on the total earnings of the worker in the month prior to the survey. We calculate the number of products/services the worker has worked on in the last month by dividing total earnings by the piece-rate. We then calculate monthly hours actually worked as the number of products/services the workers has worked on in the last month times the average time taken to complete the product/services, which therefore only reflects the hours spent on production. Manufacturing is a dummy that equal 1 if the firm's sector is tailoring, construction, or welding. ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level.

Figure A1: Study Timeline and Key Variables



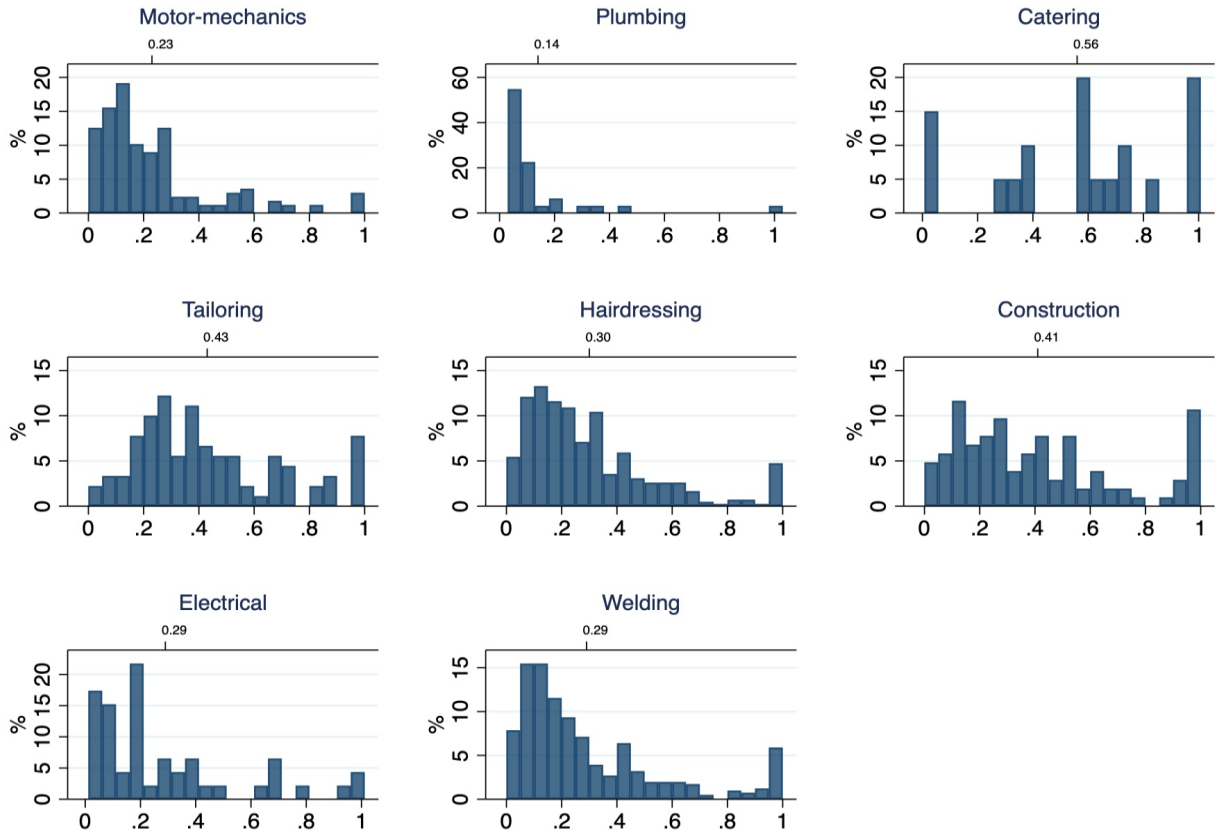
Notes: The figure shows the study timeline. Panel A indicates the timing of each survey round, as well as the time when the interventions were implemented. Firm level interventions include the wage subsidy, plus additional interventions discussed in Alfonsi et al. (2020) and Bandiera et al. (2025). Panel B shows which outcomes were measured at each survey round. An empty dot means a specific outcome was not measured at a given follow-up.

Figure A2: Sector-Specific Assets Module

Trade specific equipment: Ask only about the equipment corresponding to the enterprise type				
Item	5.16 Does this business have access to **<item>**? ***See Response Options Below*** ***If 5.16=4 → next line***	5.17 How many **<items>** does the business have access to?	5.18 How much would it cost (UGsh) to purchase all **<items>** the business has access to in a similar condition today?	5.19 Is this item also used by the owner's household for non-business uses? 1=Yes 2=No
Hairdressing				
Hand hairdryer				
Scissors				
Tongs				
Hot comb				
(large) Dryer				
Sink				
Catering				
Pot				
Pan				
Stove				
Chafing dish				
Electrical Wiring				
Grinder				
Drilling machine				
Multi meter				
Tailoring				
Sewing machine				
Scissors				
Knitting machine				
Tape measure				
Responses for 5.16 1=Yes, own item 2=Yes, pay for the service but do not own 3=Yes, do not own item, but business has access without paying 4=No, do not own and do not have access				

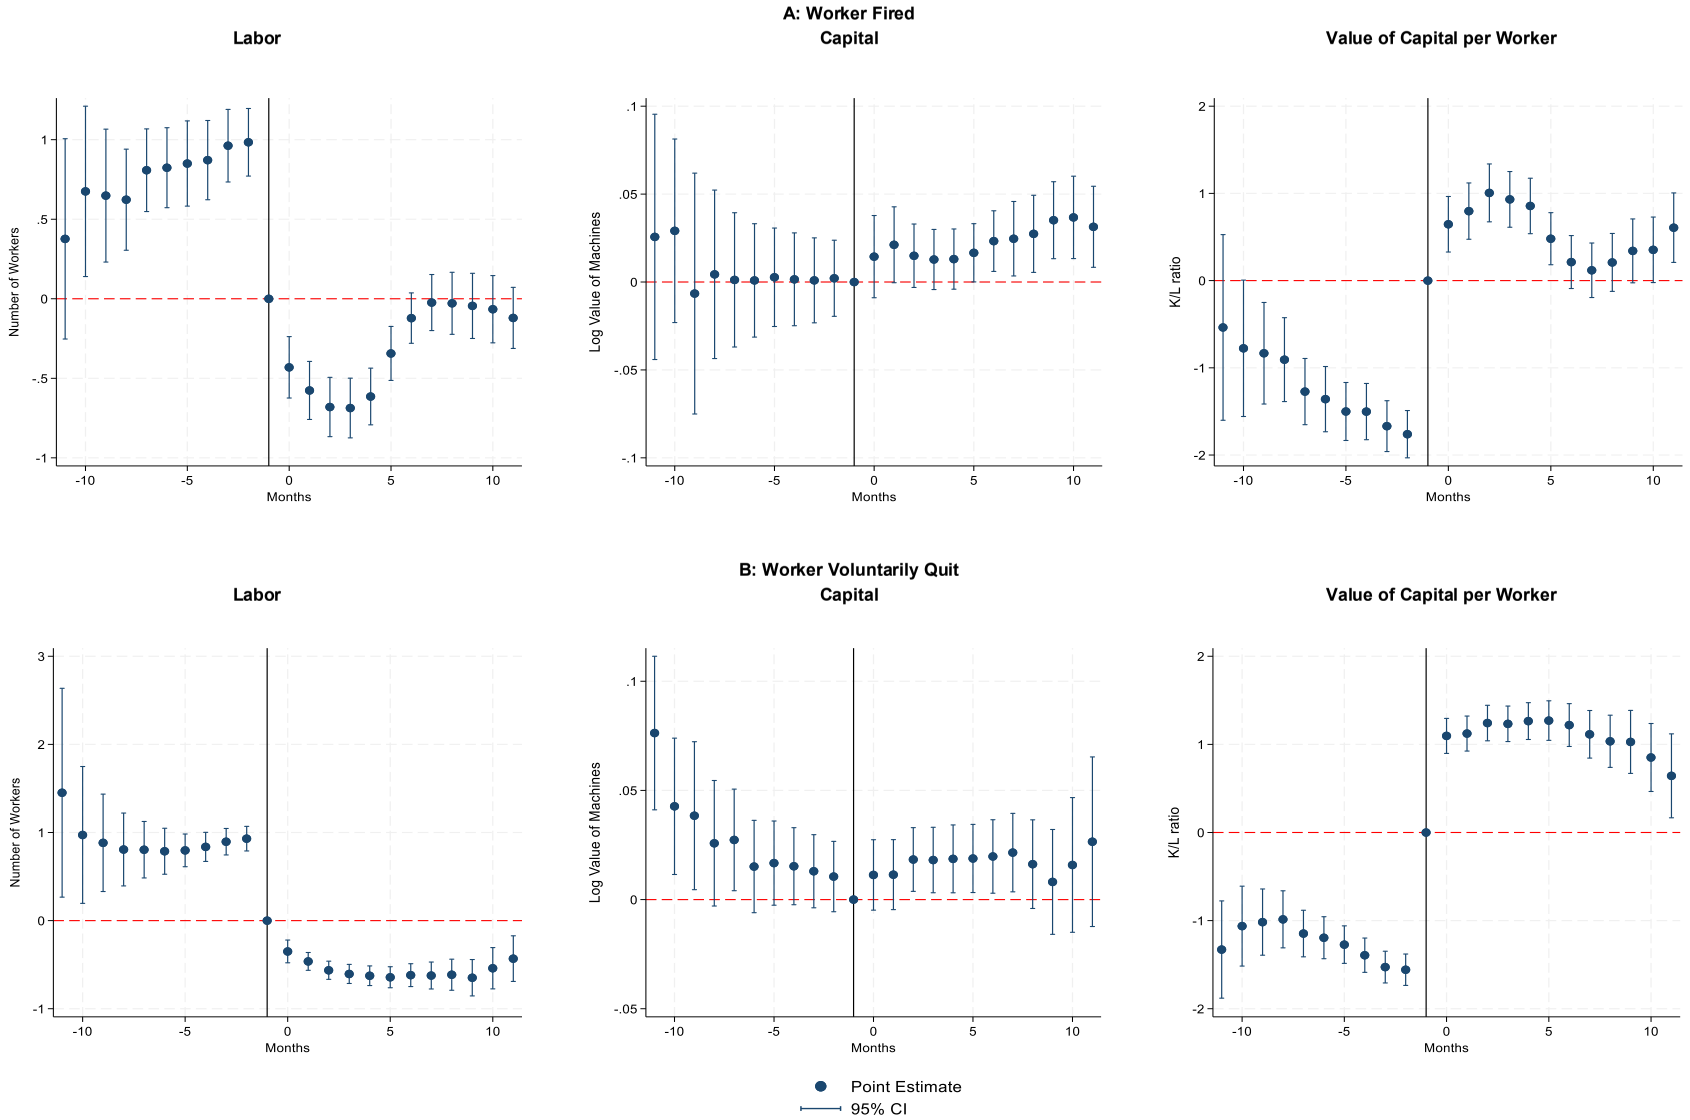
Figure A3: Labor Utilization, by Sector

Sample: Worker panel from control firms, follow-ups 2-3-4



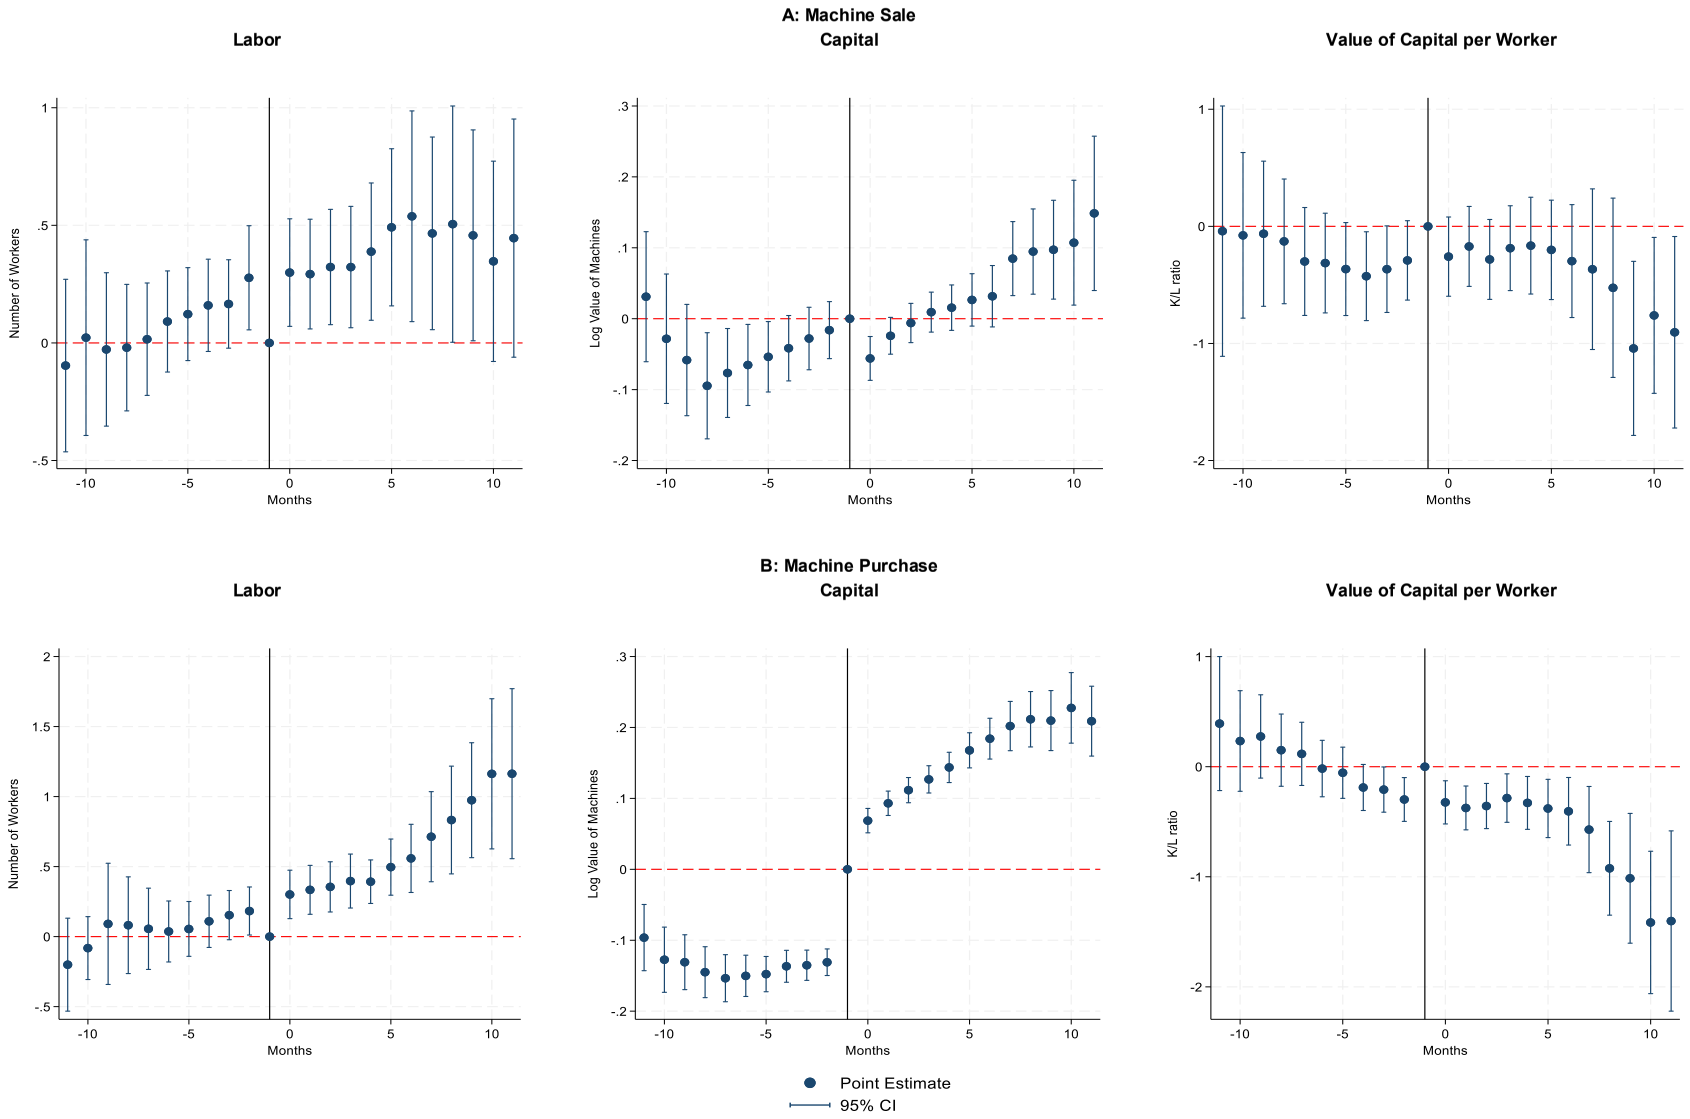
Note: The figure uses worker-level data from the Control group at second, third and fourth follow-ups to show the distribution of labor utilization. Labor utilization is computed as the share of hours workers typically spend at the firm during which they actually worked. Hours actually worked is calculated as follows. For each workers paid piece-rate, we ask the following three questions regarding a specific product/service: (1) How much time (minutes) would <name> normally take right now to complete his/her work on this product/service alone to a standard of quality considered satisfactory by the average customer of this business? (2) How much of the price charged to customers for this product/service does <name> normally receive? In addition, we have data on the total earnings of the worker in the month prior to the survey. We calculate the number of products/services the worker has worked on in the last month by dividing total earnings by the piece-rate. We then calculate monthly hours actually worked as the number of products/services the workers has worked on in the last month times the average time taken to complete the product/services, which therefore only reflects the hours spent on production.

Figure A4: Reasons for Worker Separation Event Studies



Notes: These figures show event studies for the firing (Panel A) or the loss of a worker due to voluntary workers' exits (Panel B). The event is the firing and the the loss of a worker due to voluntary workers' exits in the 12 month before a given follow-up survey. The graphs show the results of the analysis pooled across all follow-ups. The outcomes are the number of workers employed at the firm, including the firm owner (Labor), the logarithm of the firm's total value of assets (Capital) and the logarithm of the firm's total value of assets over the number of workers employed at the firm, including the firm owner (K/L ratio).

Figure A6: Machine Sale and Purchase Event Studies

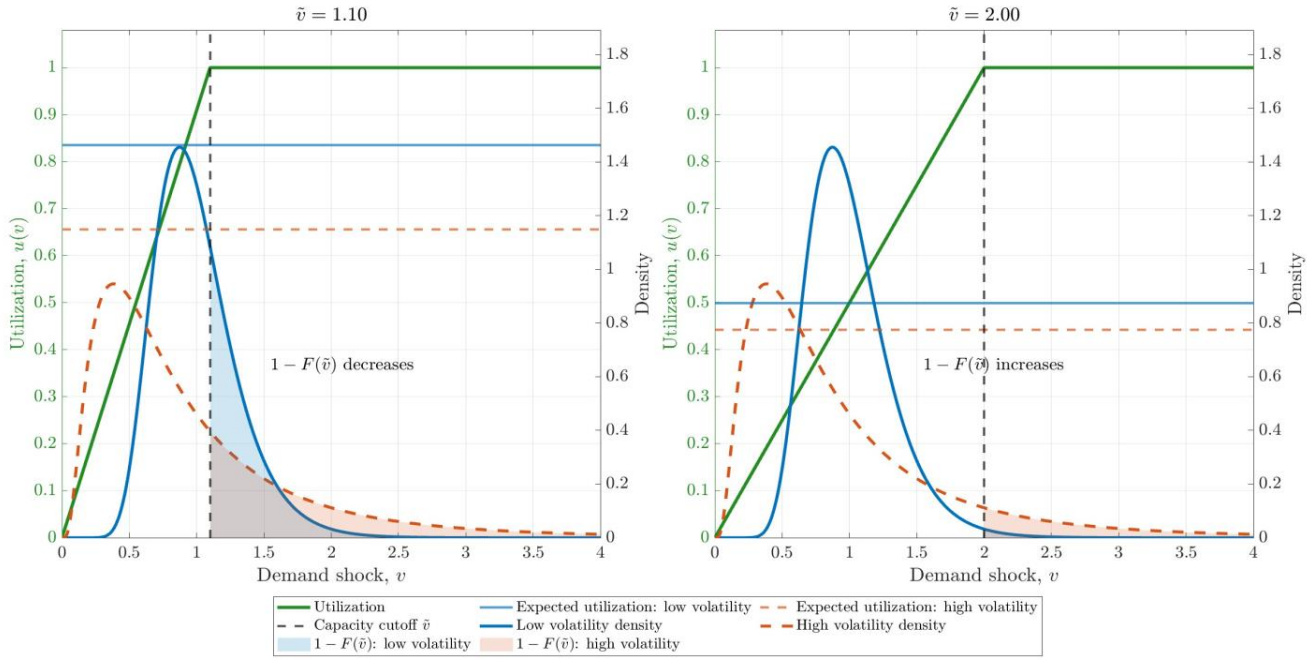


Notes: These figures show event studies for the sale (Panel A) and the purchase of a new machine (Panel B). The event is the sale or the purchase of a new machine in the 12 month before a given follow-up survey. The graphs show the results of the analysis pooled across all follow-ups. The outcomes are the logarithm of the firm's total value of assets over the number of workers employed at the firm, including the firm owner (K/L ratio), the number of workers employed at the firm, including the firm owner (Labor) and the logarithm of the firm's total value of assets (Capital).

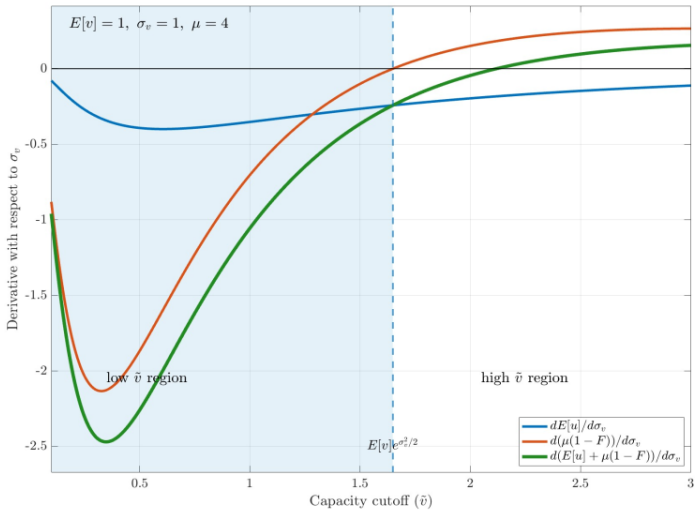
Figure A6: Illustration of Propositions 1 and 2

A. Slack states, Capacity states and Volatility

A1. Low Cutoff **A2. High Cutoff**



B. Returns from Training



Notes: Panel A shows expected utilization E(u) and tail probability 1-F as a function of demand uncertainty, holding the expected value of v constant. Panel A1 shows these values for a low level of the demand cutoff v_tilde, while Panel A2 shows the corresponding values for a high demand cutoff. Panel B shows the derivative of (i) expected utilization E(u), tail probability 1-F, multiplied by markups, and the sum of the two as a function of the demand cutoff sigma_v.