Self-employment within the firm *

Vittorio Bassi[†] Jung Hyuk Lee[‡] Alessandra Peter[§]
Tommaso Porzio[¶] Ritwika Sen[‖] Esau Tugume[**]

May 16, 2025

Abstract

We study the internal organization of manufacturing firms in Uganda. We measure what people do within firms and find limited specialization, far below what is feasible given the prevailing production process and average firm size of 5.7 workers. We build and estimate an occupational choice model in which firm size, productivity, and specialization arise endogenously. The model shows that firms in this setting are largely "self-employment in disguise" and generate just a 20% productivity gain over literal self-employment. In a counterfactual economy with full specialization, the same aggregate output can be produced with an average firm size of only 1.6.

JEL Codes: O11, O17, L23, L25

* We would like to thank Johannes Boehm, Martina Kirchberger, David Lagakos, and Claudia Macaluso for insightful discussions of the paper. We also thank Oriana Bandiera, Lauren Bergquist, Nick Bloom, Emily Breza, Thomas Chaney, Kevin Donovan, Pascaline Dupas, Dennis Egger, Doug Gollin, Supreet Kaur, Jonas Hjort, Joe Kaboski, Ted Miguel, Michael Peters, Simon Quinn, Esteban Rossi-Hansberg, Todd Schoellman, Nick Tsivanidis, John Van Reenen, Chris Woodruff and participants at many seminars and conferences from which we greatly benefited. Elena Spadini and Sai Zhang provided outstanding research assistance. We gratefully acknowledge financial support from CEGA, IGC, PEDL, UKRI, USC and World Bank. All errors are our own.

† University of Southern California, BREAD, CEPR, NBER, J-PAL and IGC; vbassi@usc.edu
‡ Ministry of Economy and Finance, Republic of Korea; jay.junghyuk.lee@gmail.com
§ New York University, CEPR and NBER; alessandra.peter@nyu.edu
¶ Columbia University, BREAD, CEPR and NBER; tommaso.porzio@columbia.edu
‖ Ludwig Maximilian University of Munich; ritwika.sen@econ.lmu.de
** BRAC NGO; tugumeesau@gmail.com

1 Introduction

A defining feature of the development process is the transformation in how production is organized: as economies grow, self-employment gives way to employment in firms (Gollin, 2002; Buera et al., 2011; Bandiera et al., 2022). This shift is underway even in the world's poorest countries, where workers are beginning to come together in firms—primarily *small firms* with fewer than ten employees (Hsieh and Olken, 2014; Quinn and Woodruff, 2019). While the emergence of small firms is encouraging in principle, it raises the question of whether this represents a meaningful reorganization of production that significantly boosts productivity.

Firms can facilitate productivity gains through labor specialization and economies of scale. However, these gains are not guaranteed—they depend on whether working together in a firm truly differs from working alone. In the extreme, a group of individuals performing identical tasks side by side could be labeled "a firm," with one person designated the entrepreneur, yet the arrangement may yield no meaningful reorganization of production and little improvement in productivity.[1] This raises a central question for development: do the small firms emerging in low-income countries capture the productivity gains from organizing labor into firms, or are they mostly self-employment in disguise? Until now, the absence of direct evidence on what workers do inside these firms has made this question difficult to answer.

We address this gap by studying the *internal organization* of manufacturing firms in Uganda. We collect time-use data from a representative sample of firms in three sectors and show that although the average firm has about six workers, labor specialization is limited: horizontal specialization across employees is virtually absent, and we find only limited vertical specialization between employees and entrepreneurs.

To assess the implications of limited specialization for aggregate output, we develop an occupational choice model in which entrepreneurs choose the assignment of workers to tasks and the size of their firm. In the model, firms serve two purposes: they enable labor specialization and facilitate economies of scale through the sharing of fixed costs. The framework nests two polar cases: (i) "self-employment within the firm", where each worker's output depends solely on their own ability and firms exist only to share fixed costs; and (ii) fully specialized, "scalable firms" where entrepreneurs leverage their talent by employing

[1] A large literature documents that much of the world's self-employed operate at low levels of productivity, leading to underutilized resources, low returns to time, and persistent poverty (Banerjee et al., 2024; Breza et al., 2021; Agness et al., 2025; Walker et al., 2024).

others.

Estimating the model with our data, we find that firms in Uganda are much closer to the benchmark of "self-employment within the firm" and that fixed costs are small. As a result, aggregate output in the current equilibrium—where the average firm has 5.7 workers and less than 18% of individuals are entrepreneurs—is identical to that of an economy composed of "scalable firms" with an average size of 1.6 workers and over 62% self-employment. These findings suggest that in low-income countries, the transition from self-employment to firm-based employment may not necessarily entail substantial changes in the *actual* organization of labor. As such, relying solely on the firm size distribution—without accounting for how labor is organized within firms—could risk overstating the productivity gains from the move away from self-employment.

We survey 1,115 firms in three sectors: furniture, metal products, and grain milling. We chose these sectors because they are large (accounting for 30% of Ugandan manufacturing employment), and because they include both small and relatively large firms with about ten employees. We show that the firm size distribution in our sample covers most of the support of Ugandan manufacturing firms.

The key innovation of our survey is to measure time use within the firm, tracking how entrepreneurs and their employees allocate each hour of their workday to 17 pre-specified tasks, including both "production" tasks (e.g., specific steps of the production process) and "non-production" tasks (e.g., interacting with customers, supervision, input procurement).

We use this data to document new facts on labor specialization within the firm and how it varies with firm size. Initially, we pool the data for furniture and metal products as these two sectors have similar production and output market characteristics and make up 85% of firms in our sample. We later contrast the results with grain milling.

We start by showing that the set of tasks performed by firms, as well as the average share of time spent on each task, are remarkably similar across the size distribution. All firms spend about 20% of time on non-production tasks and the allocation of time across various production steps does not vary across the firm size distribution. That is, larger firms simply do more of the same tasks.

We then argue that there is little evidence of horizontal, or "Smithian", specialization: on average, 85% of employees work on each production step for the products we focus on in our survey, and this percentage varies little with firm size. We build an empirical benchmark of full horizontal specialization and find that each step would need to be performed by only 25%–30% of the employees. Low specialization is thus not a mechanical byproduct of firms being small or using simple production processes.[2]

Next, we turn to vertical specialization, measured as the extent to which entrepreneurs spend more time on non-production tasks than their employees.[3] Vertical specialization is more prevalent than horizontal specialization, especially in large firms, where entrepreneurs spend twice as much time as employees on non-production tasks. However, it is far from complete: even in firms with more than five employees, entrepreneurs spend only 50% of their time on non-production tasks, even though there would be enough of these tasks at the firm level to fill the entrepreneur's day.

Having established that labor specialization is limited, we investigate several potential explanations suggested by the literature. Our data primarily support the hypothesis that a lack of product standardization in furniture and metal products limits specialization. In grain milling, which produces a standardized good (flour), we find much higher specialization.[4] We further validate this hypothesis by showing that customization entails significant communication and coordination costs within the firm and makes it costly to "unbundle" the production process into separate tasks that different individuals can perform.

The empirical evidence reveals minimal horizontal and limited vertical specialization. However, the data alone cannot quantify how "limited" specialization is, that is, to what extent firms in our setting help realize productivity gains over self-employment. To answer this, we develop a model that formalizes the link between specialization, productivity, and firm size. The model also allows us to disentangle barriers to specialization from other external factors that may constrain firm growth.

We consider an economy populated by individuals with heterogeneous abilities. The model features three nested components: (i) an assignment problem of workers to tasks within each firm; (ii) a firm size decision by each entrepreneur, subject to a reduced-form *hiring cost* that captures any external constraints; and (iii) a standard occupational choice between entrepreneurship and wage employment in the tradition of Lucas (1978).

The core of the model lies in the assignment problem. Production involves completing tasks of varying complexity. When working together in a firm, individuals can *unbundle* the production process and assign the most complex tasks to the most skilled individuals—

[2]Using data at the product level, we also show that even in large firms, each finished product unit is only worked on by 2-2.5 employees, and the average employee works on 75-85% of all the product types made by the firm.

[3]To validate this measure, we show that entrepreneurs are more skilled, and non-production tasks are more skill-intensive.

[4]Differences between the sectors in capital intensity or skill requirements do not explain differences in specialization rates.

typically the entrepreneur. Such specialization raises productivity by increasing the extent to which worker output depends on entrepreneurial ability, which we refer to as the "pass-through". However, labor specialization entails an *unbundling cost*. This cost is a reduced-form way of capturing internal barriers originating from either the demand or the supply side—for example, communication costs linked to product customization.[5]

When the unbundling cost is low, entrepreneurs can easily specialize in complex tasks and the pass-through is high. Firms serve as vehicles for scaling talent, with high-ability entrepreneurs optimally running larger firms. In contrast, when the unbundling cost is high, each worker is effectively *self-employed within the firm*, and firm productivity equals the average ability of all its members. In this case, firms exist primarily to share fixed costs, and their optimal size is smaller.

The internal organization of firms has equilibrium effects that ripple through the economy. When the pass-through of entrepreneurial ability is high, talent can be easily leveraged and is highly valued. As a result, only large, productive firms operate, while marginal entrepreneurs become workers, attracted by the higher equilibrium wage.

We estimate the model using data from the furniture and metal products sectors. To avoid overstating the role of labor specialization, we extend the model to allow entrepreneurial ability to affect worker output independently of time allocation, modulated by a parameter λ.[6] Hence, in the estimated model, the pass-through of entrepreneurial ability to worker productivity depends both on the endogenous level of specialization and the exogenous parameter λ. The estimation targets a rich set of moments capturing the allocation of labor to tasks, and heterogeneity in firm size, revenues, and worker earnings within and across firms. All parameters are jointly estimated and we provide a heuristic identification argument validated through model simulations.

Two features of our setting are particularly important for identification. First, having data on time allocation within firms, in addition to standard measures of firm size and revenues, allows us to separately identify internal (unbundling cost) and external (hiring cost) barriers. Second, in both the model and the data, workers' wages are primarily piece-rate and hence a measure of individual productivity. Using wages and profits, we can therefore estimate the overall pass-through of entrepreneurial ability to worker output and, given the

[5]The model in the main text includes the "unbundling cost" as a technological parameter in the production function. In Appendix B.3, we show how this formulation can be alternatively micro-founded starting from the demand side.

[6]This channel captures, for example, that the entrepreneur may come up with a particular product design or enjoy a good reputation.

unbundling cost, recover λ.

We then use the estimated model to conduct a series of quantitative exercises. First, we ask how much more or less output the economy would produce if average firm size remained constant but internal organization changed. This allows us to cleanly assess how much firms in our context contribute to aggregate productivity over literal self-employment. In the estimated model, entrepreneurs pass through 25% of their productivity to workers. We find that reducing the unbundling cost so as to achieve full specialization and hence a pass-through of 100% would triple aggregate output. On the other hand, reducing the pass-through to zero would lower output by only 20%. In this sense, firms in our setting are much closer to the benchmark of self-employment within the firm than to fully specialized firms in which all complex tasks are performed by the entrepreneur.[7]

Second, to understand how large Ugandan manufacturing firms truly are, we calculate the average firm size needed to generate the same output in an economy with fully specialized firms. We find that the "true" firm size is much smaller than the observed average of 5.7 workers: under full specialization, the same aggregate output could be achieved with firms averaging only 1.6 workers. This result is also driven by the fact that fixed operating costs are small in our setting and hence there are only minor technological gains from working together in a firm.[8]

Finally, we turn to policy-relevant counterfactuals. We evaluate the aggregate gains of reducing unbundling and hiring costs within an empirically relevant range and compare the impact of internal barriers (unbundling costs) to external ones (hiring costs). In our setting, internal barriers have substantially larger effects: lowering unbundling costs to match the specialization observed in grain milling increases output by 10%, compared to only 3% when reducing hiring costs to generate the same increase in firm size. Because selection into entrepreneurship is held constant by construction, the larger gains from reducing unbundling costs reflect increased pass-through of entrepreneurial ability to workers and a reallocation toward more productive entrepreneurs. We also find that internal and external frictions are complementary: the returns to reducing hiring costs are much larger when unbundling costs are low. These results highlight the importance of understanding *why* firms are small when evaluating whether interventions to expand firm size can generate large aggregate gains.

[7] A corollary of this result is that internal barriers, such as unbundling costs, substitute for external ones: once we account for low pass-through, the hiring costs needed to rationalize observed firm sizes are one-fifth as large.

[8] The presence of an active rental market as documented by our earlier work (Bassi et al., 2022b) can explain why our estimated fixed costs are small.

Related Literature and Contribution. This paper is most closely related to the extensive literature on the emergence of firms in developing countries and the external frictions hindering their growth (Gollin, 2002; Buera et al., 2011; Moll, 2014; Quinn and Woodruff, 2019; Bandiera et al., 2022). We contribute to this literature by emphasizing the importance of studying the internal organization of firms to better understand the firm size distribution, the allocation of talent, and aggregate productivity.[9]

Specifically, we make two main contributions to this literature. First, our results highlight that the simple rise of multi-person firms may not translate into significant productivity gains if labor specialization within these firms remains limited. Second, we show that accounting for internal frictions reduces the importance of external frictions in explaining small firm sizes and low productivity. For instance, if internal frictions prevent talented entrepreneurs from effectively scaling their ability, our model can reproduce the low correlation between managerial ability and firm size observed in developing countries without relying solely on large external frictions (Bloom et al., 2022).

Many before us have empirically studied the role of labor specialization for firm productivity and growth in higher income countries (Caliendo et al., 2015; Freund, 2022; Kohlhepp, 2023).[10] We contribute by showing how measuring and modeling labor specialization is quantitatively important to understand the nature of small firms in developing countries and their potential role as engines of development.[11]

Our findings reinforce the view that demand-side constraints could play a key role for productivity and growth (Lagakos, 2016; Jensen and Miller, 2018; Hjort et al., 2020; Startz, 2019; Bold et al., 2022; Goldberg and Reed, 2022; Vitali, 2022). We argue that one specific feature of demand—the prevalence of customization—curtails specialization by increasing coordination costs and making it difficult to unbundle production tasks.[12]

[9]Our paper complements a small but growing literature studying the organization of large firms and factories in low-income settings and across countries at different stages of development (Bloom et al., 2012; Hjort, 2014; Atkin et al., 2017a; Porzio, 2017; Bandiera et al., 2022; Ghosh, 2022; Boehm and Oberfield, 2023; Adhvaryu et al., 2023, 2024; Minni, 2023; Hjort et al., 2025).

[10]More broadly, we contribute to a classic literature in organizational economics, mostly theoretical, which has long emphasized the role of firms in improving coordination and specialization (Coase, 1937; Williamson, 1971, 1973, 1979; Alchian and Demsetz, 1972; Chandler, 1990; Yang and Borland, 1991; Becker and Murphy, 1992; Bolton and Dewatripont, 1994; Garicano and Rossi-Hansberg, 2006).

[11]In addition, we show the importance of collecting time-use data: in a developing country setting, relying on coarse occupational data, as the literature typically does, would not have allowed us to identify the key patterns of limited specialization. In doing so, we build on the seminal work of Bandiera et al. (2020).

[12]Jensen and Miller (2018) is particularly related to ours in that they show that small firms specialize labor as they grow larger. While we also find that small firm size reduces specialization, our key focus is to show and quantify the reverse relationship: barriers to specialization keep firms small in the first place.

Finally, this research is part of a broader agenda examining the boundary of the firm in low-income countries.[13] Our own previous work (Bassi et al., 2022b) shows that firms can achieve capital-side scale economies through rental markets, making them "larger than they appear". This paper reveals that, from the perspective of labor organization, firms are instead "smaller than they appear". Together, these two studies underscore that a firm is a nuanced concept in developing economies.

Structure of the Paper. In Section 2, we describe the survey and sample. Section 3 shows evidence on labor specialization. Section 4 develops the model, Section 5 describes the estimation, and Section 6 reports our quantitative results and counterfactuals. Section 7 concludes. Additional results are in the Online Appendix.[14]

2 Survey and Setting

We describe the survey and present descriptive statistics of our sample of firms to motivate the analysis in the rest of the paper.

2.1 Sampling

Our sample consists of manufacturing firms in the following three ISIC codes: (i) 3100 "manufacture of furniture", (ii) 2511 "manufacture of structural metal products", and (iii) 1061 "manufacture of grain mill products". For brevity, we refer to the three sectors as furniture, metal products, and grain milling. We chose these sectors because they are large (employing about 30% of workers in manufacturing) and include smaller and, by Ugandan standards, larger firms, which allows us to study labor specialization across the size distribution.[15]

We selected a representative sample of 52 sub-counties, stratifying by population and by whether the sub-county is in Kampala, the capital city.[16] We first conducted a complete listing within each sub-county and found close to 3,000 firms overall. We then randomly sampled about 1,000 firms from the listing.[17] We interviewed the firm owner and all em-

[13]Two recent papers in this broader agenda are Anderson and McKenzie (2022) and McCasland et al. (2024). Anderson and McKenzie (2022) find that training Nigerian entrepreneurs in new skills is less effective than helping them hire new specialized workers or connecting them with external market providers. McCasland et al. (2024) show that employees in Ghana contribute to the firm activities with their own personal capital equipment.

[14]Additional results not intended for publication can be found in a Supplemental Appendix posted on the authors' website and available at bit.ly/SEWIF_supplemental.

[15]The latest Census of Business Establishments shows that these three sectors comprise 32% of total manufacturing employment and 27% in firms with five or more employees (UBOS, 2011).

[16]The average sub-county consists of 5,285 households and spans 4.4 square miles.

[17]We over-sampled firms with more than five workers to ensure enough observations among relatively

ployees in the firm working on pre-specified "core" products that are common in each sector: doors in furniture, windows in metal products, and maize flour in grain milling. Our final sample includes 1,115 entrepreneurs and 2,883 employees.[18]

The main survey wave was collected in person by our enumerators in 2018–2019. We then followed up with a brief, focused phone survey in 2022.

2.2 Survey Design

Measuring Labor Specialization. Our key innovation was to collect detailed measures of labor specialization within the firm. To do so, we designed two novel modules in the main survey to measure labor specialization, each directed at both the entrepreneur and the employees.

The first is a time-use module. To start, respondents were asked to report all hours worked for the firm on the last day worked. For each hour, they were then asked which specific tasks they performed, from a pre-specified list with three categories: "production", "non-production", and "idle" time. Among production tasks, we differentiated between working on the core product or another product, and in the case of the core product, we also asked about the specific production steps performed. The list of non-production tasks encompasses all managerial/organizational activities typically needed to run a business, such as customer interactions, supervision and training, sourcing of inputs, book-keeping and financial management, maintenance of machines, and management of stock. Finally, for idle time, we recorded the time spent eating/resting, or away from the firm for non-business reasons.

The second module complements this information by asking which production steps the respondent *usually* performs on the core product (not limited to the last day worked), as well as the hours they spend on each step.[19]

In Appendix Table A.1, we list all tasks and production steps, together with the share of time the average firm spends on each.

Additional Information on Production, Demand, and Customization. We also collected detailed information on the production process of firms and the economic environment

large firms. All results are appropriately weighted to reflect our sampling strategy. We define a "firm" as one entrepreneur and all employees working under their supervision in the same premises. This is the same definition as in typical surveys of informal firms, such as the World Bank Informal Sector Enterprise Surveys.

[18]We use either the terms "entrepreneur" and "firm owner" interchangeably since in most cases they are the same person.

[19]As not all production steps for one product may be completed on the same day, this information allows us to study which steps each individual typically works on.

in which they operate. Specifically, we asked firm owners about (i) production steps and machines used to produce the core product; (ii) characteristics of employees and entrepreneurs, including an index of managerial practices (as in McKenzie and Woodruff (2017)); and (iii) features of the output market, including prices and customers.[20] We complement this data with qualitative information which we collect in a follow-up survey designed with the purpose of gathering details on product characteristics and interactions with customers. We use this additional information to shed light on the prevalence of product customization and how this may limit labor specialization.[21]

2.3 Basic Descriptives

In line with Hsieh and Olken (2014), most firms in our sample employ less than 10 workers. However, they are *not* micro-enterprises: the median firm has six workers, and hence there is scope for labor specialization.[22]

Table 1 reports summary statistics by sector for firms below and above median size. The table is divided into two parts. Panel A shows that firms in our sample are well established. They have been in business for about 10 years and make monthly profits of \$130 to \$400 (for context, GDP per capita in Uganda was around \$60 a month at the time of the study). Firms also offer relatively stable jobs (average employee tenure is around 3.5 years).

In line with other studies of small manufacturing firms in similar contexts, we find that around 90% of workers are paid piece-rate.[23] Despite the prevalence of piece-rate arrangements, there are clear employment relationships between each entrepreneur and their workers. In the follow-up survey, we find that, while employees do occasionally work for other firms, they earn on average 79% of their weekly income and spend 87% of their time working for their main employer. Further, when asked what they would do if they came to work one

[20] These additional survey modules feature in our previous work and are described in detail in Bassi et al. (2022b).

[21] Compliance with the initial survey was over 90%. The follow-up survey was a phone survey, and the attrition rate is about 32% for entrepreneurs and 41% for employees. This survey is used to provide additional qualitative evidence. As described in Appendix C, none of the moments used for estimation come from this survey; we rely on it only for one calibrated parameter. See the Supplemental Appendix (available at bit.ly/SEWIF_supplemental) for details on attrition and a summary of which specific tables and figures from the main text use data from the follow-up survey.

[22] Throughout the paper, we include the entrepreneur in the definition of firm size and worker count. The size distribution in the three sectors is reported in Appendix A.2 (see Figure A.1).

[23] See Hardy et al. (2024) and Macchi and Stalder (2023). In Table 1 we report the share of employees paid any piece-rate component. Moreover, in Appendix Table A.2 we show that few workers are paid a combination of piece-rates and salary (or in-kind payments), so that the share of workers paid pure piece-rates is only slightly lower.

day and there were no orders to work on, only 22% of employees answer that they would look for work elsewhere. The majority stated that they would remain at the premises waiting for orders to arrive and/or try to source more orders for the firm (32%), or do other productive work such as repairing or increasing stock (46%).

Panel B of Table 1 shows that, in line with the literature, almost no firms export, and the majority of sales are to customers within the district.[24] Most firms sell on order, but the underlying reasons differ across sectors: in furniture and metal products, customers buy custom-made products, while in grain milling, customers bring their own maize to be processed into flour. In furniture and metal products, we also find much larger price dispersion for the same product within the firm, which firms report to be primarily a result of customization.

Panel B thus highlights a key difference in the output market between the sectors. In grain milling, goods are relatively standardized: firms mostly turn maize into flour. In furniture and metal products, on the other hand, both smaller and larger firms produce customized goods tailored to the needs of individual consumers.

2.4 Representativeness

Before turning to the main empirical analysis, we study the representativeness of our sample by comparing it to the latest available Census of Business Establishments for Uganda (UBOS, 2011). Figure 1 reports the size distribution: (i) for our sample (pooling across the three sectors); (ii) for the same three sectors in the census; (iii) for all firms in the census.

The figure shows two main results. First, within our study sectors, our sample includes relatively large firms. This is likely due to our focus on urban and semi-urban areas, where firms tend to be larger.[25] Second, the sectors we targeted have larger firms on average.

Overall, the figure shows that our sample has good coverage of firm sizes between 1-10 workers, which represent the bulk of economic activity in Uganda.[26] Our focus on sectors and geographical areas with relatively large firms is an explicit choice, as this allows us to shed light on labor specialization across the size distribution.

In Appendix A, we compare our sample to the typical firm targeted in development inter-

[24] See, for instance, Startz (2019), Bassi et al. (2022b), Bassi et al. (2022a) and Vitali (2022).

[25] The fact that firm size is smaller in the census could also be because the Census of Manufacturing is from 2010, which is eight years before our survey.

[26] In all three samples considered in Figure 1, less than 1% of firms have more than 20 employees. While it would be interesting to study labor specialization in the few very large firms that operate in Uganda, doing so would shed light only on how production is organized in a minimal segment of the size distribution. Our paper instead aims to shed light on how firms representing the bulk of economic activity are organized.

Table 1: Firm Characteristics by Firm Size

	Furniture		Metal Products		Grain Milling	
Number of workers	<=6	>6	<=6	>6	<=6	>6
Sample size (firms)	333	189	254	179	86	74
Panel A. Well-established firms						
Firm age (yrs.)	10.1	11.2	8.2	10.7	13.5	10.0
Monthly revenues (USD)	1,120	1,516	1,174	2,386	1,046	3,181
Monthly profits (USD)	206	260	209	367	131	403
Employee tenure (yrs.)	3.3	4.0	3.2	3.7	4.0	3.6
Monthly employee earnings (USD)	72	79	68	80	42	66
Employees paid piece-rate (%)	94	91	94	91	78	84
Panel B. Local sales and customization						
Sells outside Uganda (%)	0.0	2.3	1.4	0.1	0.0	0.3
Most sales outside district (%)	9.4	9.3	8.4	14.9	3.1	7.3
Sales made to order (%)	73	80	88	90	73	63
Buy on order to customize (0/1)	0.63	0.71	0.66	0.62	0.08	0.35
Buy on order to bring own inputs (0/1)	0.06	0.00	0.05	0.04	0.72	0.41
Price dispersion w/i firm (same prod.)	1.40	1.51	1.29	1.37	1.08	1.16
Why diff. prices? Customization (0/1)	0.45	0.44	0.54	0.54	0.14	0.19
Why diff. prices? Qty discounts (0/1)	0.23	0.14	0.23	0.35	0.52	0.53

Notes: Means are reported by median firm size. Monthly profits (revenues): average reported profits (revenues) in the three months preceding the survey (trimmed at top 1%). 1 USD = 3,800 UGX for monetary amounts. All variables in Panels A and C are derived from baseline survey data. Panel B: Rows 1-3 and row 6 report information based on sales in the last three months prior to the baseline survey. Rows 4–5 report dummies for the main reason why customers buy on order in the follow up survey (we label as "Customization" the two answer options "Customers want to choose the materials/inputs" and "Each customer wants a different product"). Rows 7–8 report dummies for reasons listed among the top three for charging different prices for the same product at baseline. Panel C, row 1: to compute the firm-level share of time, we sum across the entrepreneur and all surveyed employees.

ventions. In 14 of the 16 studies we identified, firms have 10 or fewer employees, confirming that our sample represents the "typical" manufacturing firms that are studied in low-income countries (see Appendix Figure A.2 for details).

3 Evidence on Labor Specialization

In this section, we analyze the organization of labor inside the firm and how it varies across the size distribution. As a preliminary step, we describe the set of tasks firms do and then turn to our main empirical results on who does what within the firm. Given the similarity between furniture and metal products in terms of the characteristics of the output market and customization (see Table 1), we first pool the data for these sectors. We later contrast

Figure 1: Representativeness of the Study Sample

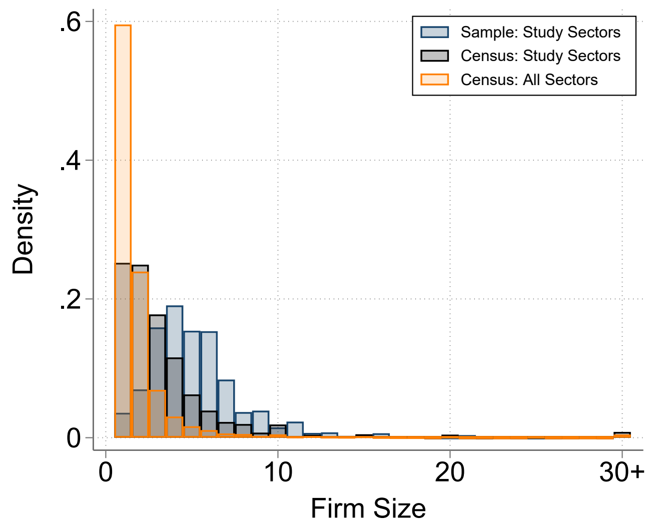

Notes: Figure 1 shows the firm size distribution in our data, using all firms included in the listing and pooling across sectors, and in the 2010 Census of Business Establishments, both restricting to our three sectors and overall. We censor firm size at 30 workers.

the results with grain milling to explore potential mechanisms behind the degree of labor specialization.

3.1 Task Composition: What Do Firms Do?

We document which tasks firms do and show that these do not vary by firm size.

Which Production Steps Do Firms Do? Since we collected data on production steps for the core product only, we limit the sample to the 80% of firms that make that product. For each individual production step, we compute the share of firms that perform that step. We then average across steps to create the share of firms performing the representative step.[27] Panel (a) of Figure 2 shows that: (i) each step is done by most firms, and (ii) this does not vary across the size distribution.[28]

[27]The core product has 10 production steps in furniture and 7 steps in metal products. See Appendix A.1 (Table A.1) for details. We average across steps, weighting by the average share of time each production step accounts for in the data, so that steps that represent a larger fraction of total production time get a higher weight. We then average across the two sectors.

[28]We censor firm size at 10 workers as very few firms are larger than that (see Appendix Figure A.1).

How Do Firms Allocate Time Across Tasks? Panels (b)-(e) of Figure 2 plot, for each firm size, the share of time spent on different tasks.[29] All firms, irrespective of their size, spend about 60% of their time in production activities, 20% in non-production, or "managerial" tasks, and the remaining 20% idle (Panel (b)). Even within managerial activities (Panel (c)) or within production across steps (Panels (d) and (e)), there is very little variation in task composition by firm size.[30]

No Specialization Across Firms. These facts have two broad implications. First, firms do not specialize in different tasks. For instance, we do not find evidence that some firms specialize in production and sell to other firms, which then specialize in customer sales. Second, there is no evidence of scale economies driven by changes in task composition, such as an overhead cost in terms of managerial time: larger firms simply do more of the same tasks.[31]

3.2 Task Allocation: Who Does What Within the Firm?

Next, we study the division of labor inside the firm. We focus on two margins of specialization: (i) within production across steps, and (ii) between production and non-production tasks. (i) is motivated by the classic "Smithian", or *horizontal*, specialization: as in the pin factory described by Adam Smith, individuals can increase their productivity by specializing in a narrow production task. (ii) is motivated by the fact that non-production tasks are more skill-intensive and entrepreneurs are more skilled than employees, as we verify in Appendix A.5 (see Table A.5). This second margin corresponds to *vertical* specialization based on skill, as in the literature on the organization of knowledge into hierarchies (Garicano (2000); Garicano and Rossi-Hansberg (2006)).

3.2.1 Labor Specialization Within Production Tasks

In Figure 3, we plot the share of employees performing the different production steps for the core product by firm size, separately in furniture and metal products. In both sectors, the

[29] Panels (d) and (e) use information from the survey module asking which production steps the respondent usually performs. This survey module was presented only to firms with at least one employee, thus explaining why the x-axis starts at a firm size equal to 2.

[30] The one exception is that one-person enterprises, reassuringly, spend little to no time on supervision or training (see Panel (c)).

[31] In the Supplemental Appendix, we show that the results in Figure 2 hold when we disaggregate the production steps and the time shares completely to reflect all individual production steps, non-production categories, and idle-time categories. In Appendix A.3 we also report versions of Panels (a) and (b) of Figure 2 with confidence intervals for the shares by firm size; these confirm that we cannot reject that task composition is constant across the firm size distribution (see Figure A.4).

Figure 2: Task Composition across the Size Distribution

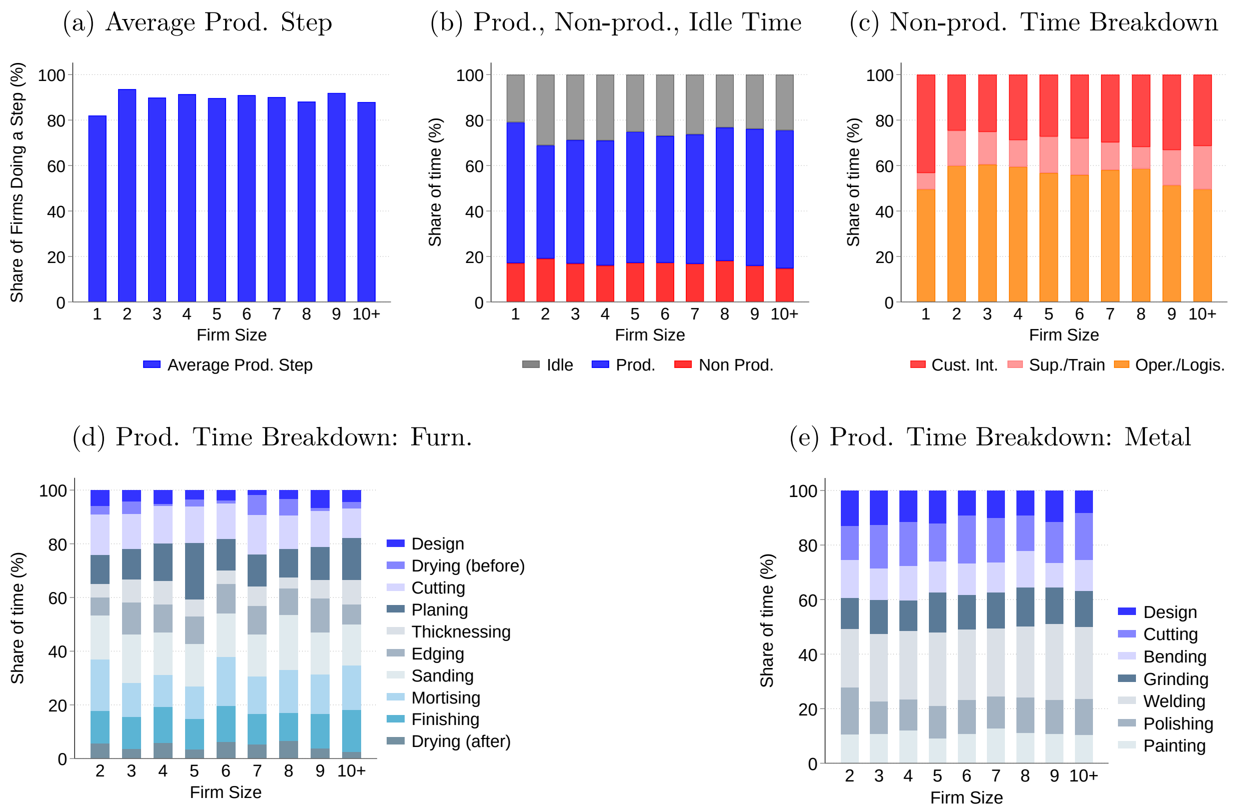

Notes: Sample: furniture and metal fabrication. Panel (a): share of firms doing the representative step, computed as described in the text. Panel (b): share of firm-level time in Production, Non-Production, and Idle tasks. Panel (c): breakdown of the non-production time into customer interaction, supervision, and operations/logistics. The category operations/logistics includes all tasks listed between bookkeeping and Other non-production tasks from Table A.1. Panels (d) and (e): breakdown of the production time of the core product into the different production steps in furniture and metal fabrication, respectively. Panels (a), (d), and (e): sample is restricted to firms making the core product.

share of employees working on each step is high and barely decreases with firm size: about 85% of employees work on the representative step in firms of size 6, and the share remains close to 80% even in firms of size 8–10. Further, there is little heterogeneity across steps, especially for the important ones.[32]

To interpret the magnitudes, we build an empirical benchmark corresponding to the share of employees working on a production step if there was full specialization. To do so, we reassign employees across steps to minimize the overlap between employees while keeping the firm-level time on each task constant. Comparing the actual allocation with the full

[32] In Appendix A.4, we present a version of Figure 3 with confidence intervals for the representative production step (see Figure A.5). This demonstrates that there are no statistically significant differences in specialization between mid-sized and larger firms in furniture and metal products.

specialization benchmark highlights that horizontal specialization is limited relative to what would be potentially attainable given the firm size distribution and the complexity of the production process. In Appendix A.4 we show that specialization across production steps is also limited for entrepreneurs throughout the size distribution (Figure A.6).

Figure 3: Task Allocation within Production Across the Size Distribution

(a) Furniture | (b) Metal Products

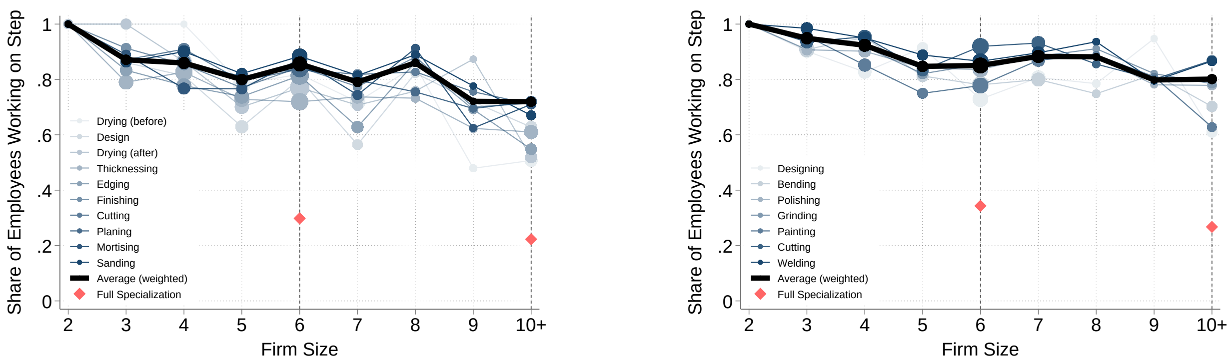

Notes: Sample: furniture (Panel (a)) and metal products (Panel (b)). The figures report the share of employees working on each production step, where darker blue colors indicate steps that correspond to a larger share of production time. We also include – in black – the share of employees performing the representative step, which is computed following the same procedure as for Figure 2, Panel (a). The red diamond markers represent the full specialization benchmarks computed for firms of size 6 and 10 (see main text for definition). To build the figures, we use information on which production steps individuals *usually* perform, rather than information from the time-use diary for the last day worked. We do so because not all production steps for one product may be completed on the same day.

Product-level Analysis. In Appendix A.4, we use a new survey of firms in the same sector and locations conducted in 2022-23 to show that even in firms with 8-10 workers: (i) only 2-2.5 different employees work on each product unit; and (ii) the average employee still works on 75-85% of all product types made by the firm (see Figure A.7). These additional results address the potential concern that focusing on the allocation of employees' time across production tasks for a core product, as we do in this section, might under-estimate specialization if workers rotate across tasks over a day (so that each product unit might be worked on by many employees) or if workers specialize across products: whichever way we try to measure it, horizontal labor specialization is limited and only weakly increasing with firm size.

3.2.2 Labor Specialization Between Production and Non-production Tasks

In Figure 4, we compare the time the entrepreneur and the average employee spend on each task. The y-axis shows the different tasks: blue ones are related to production, red ones to non-production, and grey ones to idle time. Each bar reports the share (normalized to 100%)

of that task done by the entrepreneur (the dark portion of the bar) and the average employee (the light portion). If the entrepreneur and the average employee were to spend the same time on a given task, the dark and light bars would each amount to 50%. Figure 4 offers two takeaways. First, entrepreneurs specialize in non-production tasks, and employees in production ones. This shows that there is some (vertical) specialization along this margin, and justifies our partitioning of tasks into production and non-production. Second, even though there is some specialization, there is also substantial overlap between entrepreneurs and the average employee in terms of time allocation.[33]

Figure 4: Time Allocation Between Production and Non-production Tasks

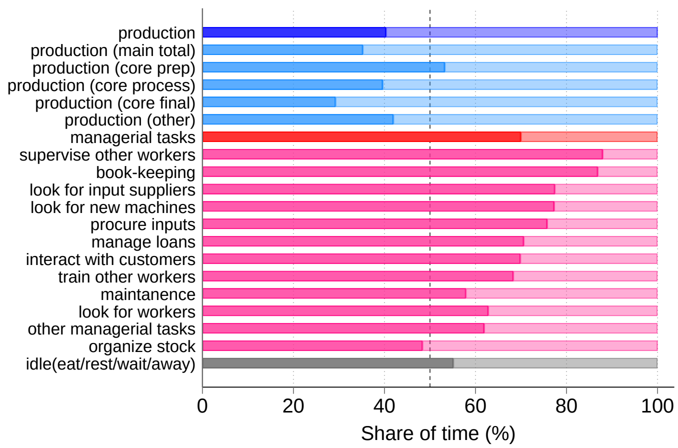

Notes: The figure compares the time spent on each task by the entrepreneur (dark bars) and the average employee (light bars). Blue bars: production. Red bars: Non-production. Grey bars: Idle time. "Production (core prep)", "Production (core process)" and "Production (core final)" include the following production steps: "Preparation (core prep)": (i) Furniture: Design & Drying (before production), (ii) Metal products: Design; "Production (core process)": (i) Furniture: Cutting-Mortising, (ii) Metal Products: Cutting-Welding; "Production (core final)": (i) Furniture: Finishing & Drying (after painting), (ii) Metal Products: Polishing & Painting. Sample: furniture and metal products.

More Vertical Specialization of Entrepreneurs in Larger Firms. In Figure 5, we study how vertical specialization varies across the size distribution. To do so, we aggregate all non-production tasks, and plot the average individual's share of time spent in these tasks as a function of firm size, for both employees and entrepreneurs.[34] The figure shows that the share of employees' time in non-production tasks is constant at around 20%. Instead, entrepreneurs spend about twice as much time in non-production tasks, and the gap relative

[33] In Appendix Figure A.9, we compare the time allocations of different employees in the same firm. We find that the high-skilled spend a bit more time on non-production tasks, but the overlap is substantial: overall, differences across employees are less pronounced than differences between employees and entrepreneurs.

[34] In Figure 5 we only consider production and non-production time and instead drop idle time.

to employees increases in firm size. Hence, larger firms are more vertically specialized.[35]

Even Large Firms Are Far from Full Vertical Specialization. Panel (b) of Figure 5 shows that vertical specialization increases weakly with firm size: going from a firm of size one to a firm with five workers, the share of time in non-production activities only increases from about 34% to 45%. Even in large firms, the entrepreneur spends only about *half* of her time on non-production activities.[36]

One possibility is simply that there are not enough non-production tasks to keep entrepreneurs busy. To show that this is not the case, we compute, for each firm, the (hypothetical) share of time that the entrepreneur would spend in non-production tasks if she had fully specialized. To do so, we reassign the total firm-level time on non-production tasks to the entrepreneur.[37] The observed relationship between specialization and firm size is far from this empirical full-specialization benchmark (in pink); in fact, it is closer to a flat line. This highlights that limited vertical specialization is not a mechanical artifact of firms being small.

One direct implication of the limited specialization of entrepreneurs is that most non-production activities in larger firms are done by employees, not the entrepreneur. For example, in firms of size six, close to 70% of non-production activities are done by employees, despite the entrepreneur having enough time to do all of them.

Robustness: Focusing on the Most Complex Non-Production Tasks. In Figure 5 we are aggregating all non-production tasks. However, as shown in Appendix A.5, there is variation in the complexity of non-production tasks, with customer interaction, supervision, and input procurement being the most complex ones. Motivated by this, in Appendix A.5 we show that the results in Figure 5 are robust to focusing only on these most complex non-production tasks (see Table A.4).[38]

In sum, we find that labor specialization is limited. Still, vertical specialization of entrepreneurs in complex non-production tasks seems relatively more important than horizontal specialization of labor across production tasks. As a result, we focus on vertical specialization

[35] In Appendix A.5 we show that the results in Panel (a) of Figure 5 are not driven by some workers specializing in production while others specialize in non-production (Figures A.9 and A.10).

[36] In Appendix A.5, we also show that entrepreneurs specialize in the more difficult steps within production, although again only to a limited extent (Figure A.11).

[37] The counterfactual share of time in non-production tasks stays at 100% in firms with more than six workers.

[38] In the Supplemental Appendix, we also show that our measurement of non-production tasks is consistent across the size distribution, and report several pieces of evidence suggesting that more specialized firms are more productive (even conditional on a rich set of controls).

Figure 5: Task Allocation Between Production and Non-production by Firm Size

(a) Employees

(b) Entrepreneur

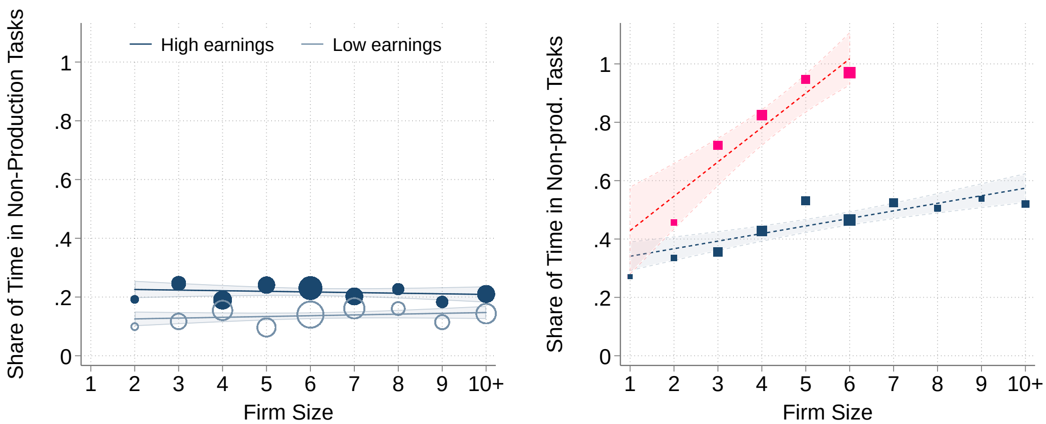

Notes: Sample: furniture and metal products. Shaded areas: 95% confidence intervals. The sizes of dots and squares represent the number of firms in each size group. Time use reported by interviewed entrepreneurs and employees. Idle time is excluded. Panel (a): Employee share of time in non-production tasks. Employees are classified as high and low earners within each firm (above or below the median). Panel (b): Entrepreneur share of time in non-production tasks. The pink squares represent the benchmark of full specialization, as described in the text.

in the model in the next section.

3.3 Correlates of Specialization

Our results so far raise the question of why there is low specialization in this context. Figure 5 rules out that firms are simply too small: we observe low specialization even in firms that are large enough to specialize. In this section, we use our detailed survey data to explore other potential drivers of low specialization.

3.3.1 Correlates of Specialization within Sector

We begin by exploring firm-level correlates of specialization within sector. To do so, we estimate the following regression for firm i in sector s and region r, pooling the data for furniture and metal products:

$$specialization_{isr} = \alpha + \beta X_i + \gamma size_i + \delta_s + \eta_r + \epsilon_{isr}, \tag{3.1}$$

where $specialization_{isr}$ measures either horizontal or vertical specialization. Following Figures 3 and 5, horizontal specialization is defined as the share of employees performing the representative step, and vertical specialization as the gap between the share of the entrepreneur's and average employee's time in non-production tasks. We regress these on various firm-level characteristics X_i as well as region (η_r) and sector (δ_s) fixed effects. Since we explored the role of firm size in the previous section, we also control for firm size, allowing us to focus on other potential factors explaining variation in specialization.[39] Nevertheless, in Appendix A.6 we show that our conclusions are robust to running these regressions without controlling for firm size (see Figure A.12).

The results are in Figure 6. Each row shows a separate regression for a different firm characteristic X_i. To interpret the results more easily, we split the sample based on each characteristic and report the mean predicted specialization in the two groups of firms. The bars correspond to the 90% confidence intervals of the difference in means.

We find suggestive evidence in support of multiple theories for low specialization proposed by the literature. Firms with high-ability managers (defined as those scoring above the median on our index of managerial practices) are more specialized, which is in line with the literature on managerial practices and firm productivity (see, e.g., Bloom et al. (2013)). Firms in which employee absenteeism is more prevalent are less specialized (along the horizontal dimension), consistent with absenteeism increasing coordination costs (Atencio-De-Leon et al., 2023).[40] Firms where all employees are hired through family and friends (labeled "family firms") are less specialized. In principle, we might have expected more or less specialization in these firms: for instance, the literature finds that firms where family members of the owner run the firm tend to be less well managed (Bennedsen et al., 2007), but may also have stronger trust between owners, managers and employees (Akcigit et al., 2021; Bloom et al., 2013), which could facilitate specialization.[41]

While the majority of firms pay everyone piece rates, those that pay at least some of their employees a fixed salary (without a piece-rate component) are more specialized. This is consistent with moral hazard leading employers to prefer piece-rate contracts, but this type of compensation hindering specialization (Holmstrom and Milgrom, 1991, 1994).[42] Finally, firms in Kampala are more vertically specialized, which is consistent with access to larger markets facilitating standardization and hence specialization (Piore and Sabel, 1984; Holmes and Stevens, 2014).

We find no significant difference in specialization by employee tenure. This suggests

[39] Consistent with the analysis in the rest of the paper, we censor firm size at 10 workers.

[40] The average employee is absent 1.8 days per month.

[41] 28% of firms are "family firms" according to our definition. To define family firms, we exploit a survey question where for each employee we know if they were hired through family or friends. We compare firms where everyone was hired through family/friends with firms where no employee was hired in this way.

[42] 9% of firms pay at least some employees a fixed salary. For this analysis, we exclude unpaid workers and workers paying the owner of the firm for training (less than 2% of workers in total).

that lack of specialization is not driven by apprenticeship or training motives, whereby entrepreneurs spend time in production to train employees (Hardy and McCasland, 2023) or to help them with more complex production tasks.[43] We also find no evidence that more mechanized firms are more specialized. An important literature in development finds that larger firms employ more capital-intensive technologies as they face different factor prices (see, e.g., Söderbom and Teal (2004)), something that we verified also in our previous work (Bassi et al., 2022b). Our results complement this literature by showing that mechanization does not appear to affect labor specialization in this context.[44]

Although not causal, these results do uncover systematic heterogeneity in specialization. At the same time, the magnitude of differences within sector are small, implying that limited specialization is a pervasive feature of how firms operate in this setting. Even in the best-managed firms in furniture and metal products, workers and entrepreneurs are not highly specialized.

3.3.2 Heterogeneity in Specialization across Sectors

We contrast the within-sector heterogeneity in specialization with differences across sectors. To do so, in the last row of Figure 6 we run a version of equation (3.1) with the full sample and compare grain milling vs. furniture/metal products (solid red bars). The results are striking: grain millers are much more specialized both horizontally and vertically, and the magnitude of these cross-sectoral differences swamps all within-sector heterogeneity. For example, while high ability managers in furniture and metal products are 5pp more vertically specialized than low ability managers, grain milling firms are 27pp more specialized than furniture/metal products firms (controlling for firm size).

Importantly, the last row of the figure also shows that the larger specialization in grain milling is robust to controlling for *all* the firm-level covariates X_i at the same time (dashed red bars). This confirms that the different rates of specialization between grain milling and the other two sectors cannot be explained, for instance, by grain milling firms being more capital intensive (as we control for mechanization rates).[45]

[43] See the Supplemental Appendix for additional analysis ruling out that apprenticeship motives can explain the documented limited specialization.

[44] We define the mechanization rate as the share of machines used among the pre-specified list of machines asked in the survey. The literature has shown the importance of mechanization for labor specialization in the industrial revolution in the U.S., thus highlighting that mechanization may play an important role in a different context (Atack et al., 2019, 2023).

[45] The results of these specifications controlling for all covariates X_i simultaneously are reported in detail in the Supplemental Appendix.

Figure 6: Correlates of Specialization

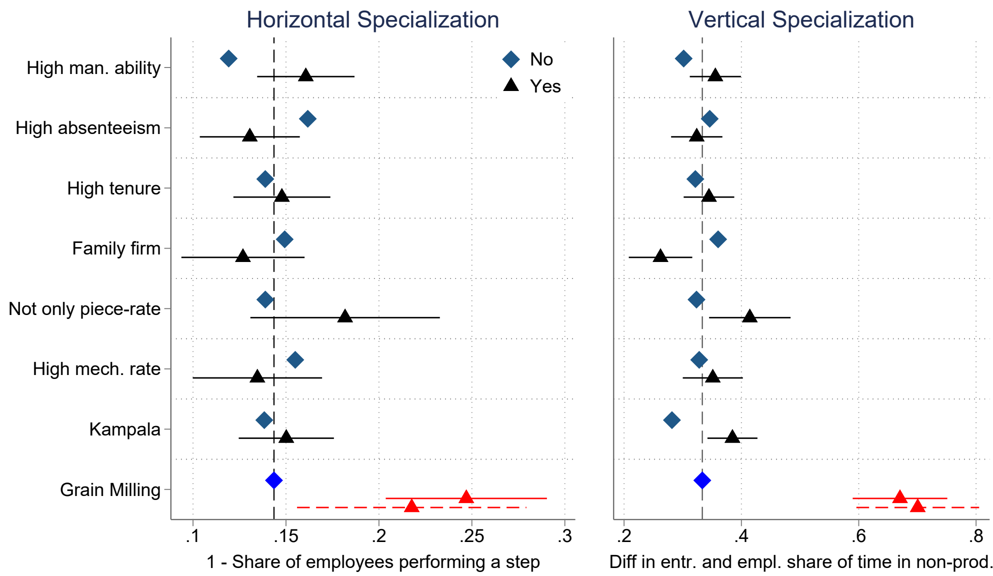

Notes: Sample: rows 1-7: furniture and metal products; rows 8: all sectors. Vertical dotted line: mean horizontal (left panel) and vertical (right panel) specialization in furniture and metal products. Results of OLS estimation of equation 3.1. Definition of horizontal and vertical specialization: see main text. The y-axis lists the independent variables of interest in each regression (see main text for definitions). For managerial ability, absenteeism, tenure and mechanization rate we split firms by below/above median. Diamonds: predicted mean in the comparison group (e.g., below median managerial ability firms). To predict this mean, we subtract from the average specialization in furniture and metal products the coefficient on the relevant characteristic of interest (e.g. dummy for above median managerial ability), weighted by the share of observations with that characteristic. Triangles: predicted mean in the comparison group plus estimated coefficient on the characteristic of interest. Bars: 90% confidence intervals.

In Figure 7 we show how labor specialization varies with firm size in the different sectors. The results are again striking: while in small firms specialization is similar across sectors, larger grain millers are substantially more specialized.[46] In Appendix A.6 we also show that there is less idle time in grain milling and that it decreases faster with firm size, consistent with higher labor specialization resulting in better time coordination (Figure A.13).[47]

[46]Panel (a) of Figure 7 confirms that in grain milling the empirical benchmark of full horizontal specialization is similar, implying that differences in the shares of employees performing the representative step can be interpreted as differences in horizontal specialization across sectors. In Appendix A.4 we also show that in grain milling each finished product unit is worked on by more workers than in the other two sectors, and this is especially the case in larger firms (Figure A.7). This is again consistent with more horizontal specialization in grain milling.

[47]In the Supplemental Appendix we report the average levels of horizontal and vertical specialization for entrepreneurs and employees in the three sectors, and their correlation with firm size. These confirm the results of Figure 7. There we also show that employees in grain milling spend a smaller share of time in non-production tasks than in furniture/metal products and that this share does not vary with firm size. Therefore, the results in Panel (b) of Figure 7 underestimate the actual differences in vertical specialization across sectors.

Figure 7: Heterogeneity in Specialization Across Sectors

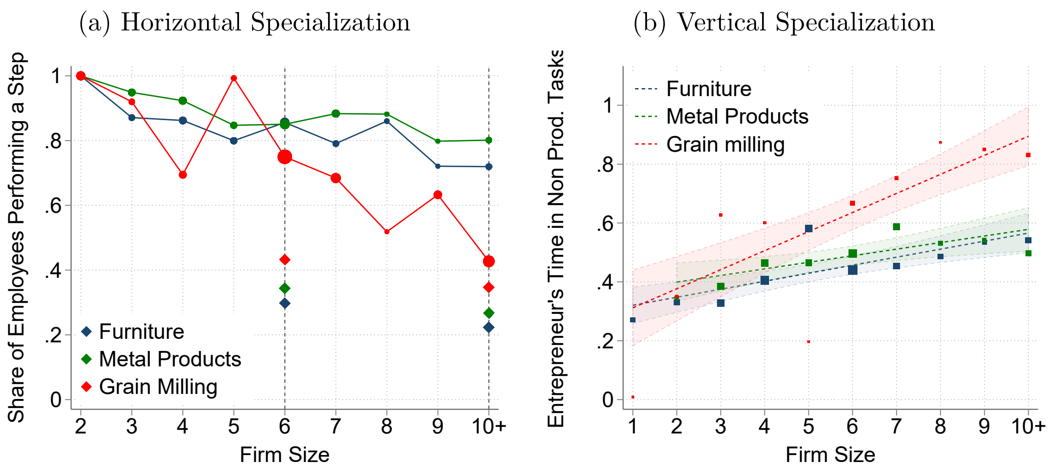

Notes: Sample: all sectors. Panel (a): replication of Figure 3 by sector, focusing on the representative step only. See Figure 3 for details. Panel (b): replication of Figure 5 Panel (b) by sector. See Figure 5 for details.

3.3.3 Customization and Labor Specialization

As highlighted in Section 2, customization is significantly more common in furniture and metal products than in grain milling. Combined with the sectoral differences in labor specialization documented above, this suggests that limited product standardization may explain why firms in furniture and metal products operate with a production technology that makes specialization difficult. We probe this further in Appendix A.6, where we show that within furniture and metal products, firms that are more engaged in customized production (as measured by the proxies for customization used in Panel B of Table 1) are less specialized on the vertical dimension (see Figure A.14).

Why does product customization limit labor specialization? In practice, customization may entail communication and coordination costs that make it costly to specialize.[48] In Table 2, Panel A, we use our survey to corroborate this hypothesis. First, we show evidence consistent with the idea that customization increases communication and coordination costs: furniture makers and metal fabricators are twice as likely to report that customers buy on order because they want to discuss the details with the person producing. Direct phone communication with the person producing is also twice as common as in grain milling. In addition, "independent orders" are more prevalent in furniture and metal products, an arrangement in which a single employee manages the entire production of the order as well

[48]This would be consistent with a large literature studying the link between standardization, specialization, and scale of operation (Piore and Sabel, 1984; Holmes and Stevens, 2014; Vickery et al., 1999; Dessein and Santos, 2006).

as the relationship with the customer, thereby eliminating any communication costs within the firm.

Second, as each door or window can have different features, it may be difficult to set up a production line with fully specialized workers. To gauge this, the last row of Panel A shows that, according to our data, there are 13 different varieties of doors and 7 varieties of windows.[49] By contrast, in grain milling there are only 4 varieties of flour, so that setting up standardized processes is likely simpler, as firms produce a larger quantity of each product.

Table 2: Implications and Drivers of Customized Production

	Furniture	Metal	Grain Mill.
Panel A. Communication and Coordination Costs			
Customers want to discuss order with person producing	52%	48%	22%
Customers have phone number of person producing	23%	23%	10%
Workers perform independent orders	49%	53%	26%
Potential varieties of core product	13	7	4
Panel B. Product Complexity			
Potential number of machine types for main product	24	20	13
Minimal time needed to produce main product (mins.)	433	351	56
Median days to complete typical order	4.0	4.0	0.6

Notes: Means. Data is from our firm surveys. All statistics are at the firm level. Panel A, row 1: dummy if discussing details with person producing is among top 3 reasons why customers buy on order. Panel A, row 3: dummy if employees perform independent orders. Panel A, row 4: number of different varieties of doors, windows, and flour in the sample. Rows 1–3 of Panel A are conditional on the firm having at least one employee. Panel B, row 1: number of machine types used to produce doors, windows, and flour.

Product Complexity as a Driver of Customization. Why are products more customized in furniture and metal products than in grain milling? Our data suggests that product complexity is a key difference across sectors. As shown in Panel B, many kinds of machines can be used to make products in furniture and metal products. Products also take several hours to make, usually over multiple days. The scope for customization and quality variation is therefore high.[50] This is not the case in grain milling, as flour is a much simpler

[49] Note that these statistics just refer to product varieties (e.g., two-panel vs. four-panel doors), and not to the customization that is conducted on top of this (e.g., the precise size of the two-panel door), and are thus a lower bound on the number of possible product varieties.

[50] Building codes could facilitate standardization in sectors such as furniture and metal products. Anecdotal evidence from our field visits confirms that building codes are present in Uganda but loosely enforced. For example, one larger furniture maker reported that even when they get orders of doors for formal buildings, the size of the door frames usually vary from building to building, and this uncertainty is a key reason why they have not been able to set up a production line. When asked about how the firm is organized, he said:

and more standardized product.[51]

In sum, the sectoral heterogeneity is consistent with the notion that the limited specialization we document is a byproduct of firms producing customized products rather than standardized goods. This heterogeneity also reassures us that the limited specialization in furniture and metal products is not simply due to measurement error, as the measurement of time use is the same across the three sectors.

4 Model

The key takeaways from our empirical findings are that there is virtually no horizontal specialization in the furniture and metals sectors and that vertical specialization is limited. To quantify what "limited" means and assess aggregate implications, we develop a model of vertical specialization within the firm. The model formalizes the relationship between specialization, productivity, and firm size and allows us to address three questions: (i) how far we are from full specialization vs literal self-employment within the firm, (ii) what are the consequences of barriers to specialization for firm size, allocation of talent, and aggregate output, and (iii) how do barriers to specialization interact with other, external, frictions keeping firms small.

The heart of the model is an assignment problem: the entrepreneur chooses how many complex tasks each of her workers performs and how many she takes on. This trade-off is governed by an "unbundling cost", which is a reduced-form way of capturing barriers to specialization, such as a customized production process.[52]

4.1 Environment

Agents. The economy is populated by a measure 1 of agents who differ in their ability $z \in [0, z_{\max}]$, distributed according to $G(z)$. Each agent supplies one unit of labor inelastically and has linear utility over consumption. Individuals can start a firm and become

"I wish we had a production line, but now it is more like a big workshop." When asked if and how he would reorganize production if he could be sure that all doors had the same size, he replied: "I would set up a production line. In fact, I also have a firm that makes pre-packaged snacks, and there we have a production line".

[51]In exploiting differences in complexity across products, we relate to a literature on the role of product complexity for trade frictions (Juhász and Steinwender, 2018) and for building capabilities and specialization through trade (Atkin et al., 2021).

[52]Like the seminal work on the organization of knowledge into hierarchies (Garicano (2000); Garicano and Rossi-Hansberg (2006)), we focus on vertical specialization based on comparative advantage. We depart from these models by allowing people of different abilities to perform equally complex tasks in equilibrium.

entrepreneurs (owners) o, or join the labor market as employees (workers) w. The resulting distributions of ability in the two occupations are F_o and F_w with $G(z) = F_o(z) + F_w(z)$. We refer to the ability of a generic firm owner as $\hat{z}$ and to that of a generic worker as z.

Labor Market. Entrepreneurs choose the mass $n-1$ of workers they would like to hire.[53] Before hiring, they need to pay a sunk cost, given by a function of hired workers—$\chi(n-1)$. This represents all other expenses incurred by the firm, including hiring costs, capital expenditures, credit frictions, and any other auxiliary costs or frictions that scale with firm size.

There is a single labor market that randomly matches workers and entrepreneurs. Hence, the distribution of hired workers is equal to the equilibrium distribution of abilities $F_w(z)$. Upon matching, wages are determined via Nash-Bargaining between each individual worker and the employer.

Production. There is a single good in the economy, chosen as the numéraire. Each agent has access to a technology that transforms one unit of time into output. Production requires two types of tasks, production, or "simple" (e.g., thicknessing wood), and non-production, or "complex" (e.g., negotiating with customers). Since the task composition is independent of firm size in our setting (Figure 2), we assume that complex tasks take up a fixed fraction of time, α. The difference between simple and complex tasks lies in their skill sensitivity: all agents are equally good at simple tasks but have ability z to do the complex ones. If an agent completes all tasks, she produces z units of the good.

When working together in a firm, individuals can trade tasks. Since agents differ only in their ability to perform complex tasks, there are gains from swapping — i.e., there are gains from specialization of labor. In principle, the problem of assigning workers to tasks is infinite-dimensional: for each pair of individuals $\{z, z'\}$, it specifies the fraction of z's complex tasks that are performed by z'.[54] We describe the full problem in Appendix B and focus here on a special case that we prove arises under two empirically relevant assumptions.

Assumption **1.** Each entrepreneur spends some of their time on simple tasks.

Assumption **2.** The workers' bargaining weight is sufficiently small.

Assumptions 1 and 2 are joint assumptions on the parameters of the model.[55] They

[53] We use n to refer to the size of the firm, hence $n-1$ is the mass of hired workers.

[54] We do not explicitly specify the assignment of simple tasks. Since all agents are equally productive at those and labor supply is inelastic, their assignment is irrelevant.

[55] Formally, Assumption 2 requires workers' bargaining weight ω to satisfy $\omega \leq \left(\frac{\partial(\max_{\boldsymbol{\mu}} y(\hat{z}, z, \boldsymbol{\mu}))}{\partial z}\right)^{-1} \forall\{z, \hat{z}\}$. $y(\hat{z}, z, \boldsymbol{\mu})$ is defined in (4.1).

are consistent with the empirical evidence discussed in Section 3 and the fact that there is relatively little variation in worker compensation.

Under Assumption 2, there is sorting by ability in equilibrium: owners are the most skilled individuals in the firm.[56] Assumption 1 guarantees that entrepreneurs have time to take on more complex tasks at the margin. Together, they imply that the owner takes over any complex task delegated within the firm. The assignment thus simplifies from an infinite-dimensional problem to choosing one number for each employee z: the fraction of complex tasks the entrepreneur takes over $\hat{z}$, $\mu(z,\hat{z}) \in [0,1]$. We collect all such assignments in the firm-level vector $\boldsymbol{\mu}$.

Using this result, the output of a firm of size n, owned by an individual with ability $\hat{z}$ who chooses task assignment $\boldsymbol{\mu}$, simplifies to

$$Y(\hat{z},n,\boldsymbol{\mu}) = y(\hat{z},\hat{z},\boldsymbol{\mu}) + (n-1)\int y(z,\hat{z},\boldsymbol{\mu})\frac{dF_w(z)}{F_w(z_{\max})}, \tag{4.1}$$

$$\text{where} \quad y(\hat{z},\hat{z},\boldsymbol{\mu}) = \hat{z}$$

$$y(z,\hat{z},\boldsymbol{\mu}) = \hat{z}^{\mu(z,\hat{z})} z^{1-\mu(z,\hat{z})}(1-\kappa(\mu(z,\hat{z})))$$

Since the owner is the most productive individual in the firm, she performs all her complex tasks and her "production line" yields output $\hat{z}$. The output generated by an employee's production line depends on the average ability with which the complex tasks are performed, net of an *unbundling cost*. The more tasks are delegated, that is, the larger is $\mu(z,\hat{z})$, the larger the weight on owner ability, $\hat{z}$.

Re-assigning tasks is costly. To assign parts of a production line to a different person, tasks must be *unbundled*. For example, if the entrepreneur negotiates all orders, she must then communicate precisely what customers want to the employee producing the order. The cost of unbundling, $\kappa(\mu(z,\hat{z}))$, is increasing in the share of complex tasks delegated by z. In Appendix B, we micro-found this functional form in a simple model in which consumers have preferences over different types of goods and delegating within the firm may lead to quality losses, thus explicitly linking the unbundling cost to customized production.

Pass-Through of Entrepreneurial Talent. The extent to which the firm owner matters for the output of workers can be summarized by the *pass-through* of entrepreneurial ability to worker productivity:

[56] We show this formally in Lemma 1 below.

$$\frac{\partial \log y(z, \hat{z}, \boldsymbol{\mu})}{\partial \log \hat{z}} = \mu(z, \hat{z}). \tag{4.2}$$

In this model, pass-through is equal to labor specialization—the share of complex tasks that the entrepreneur takes over for each worker.[57] The lower the cost of specialization $\kappa(.)$, the higher this share and the more closely worker productivity tracks entrepreneurial talent.

4.2 Choices

We next describe the choices of economic agents: whether to be workers or start a firm, how many workers to hire, and the assignment of individuals to tasks.

Profits. An entrepreneur with ability $\hat{z}$ chooses firm size n and task assignment $\boldsymbol{\mu}$ to maximize profits:

$$\pi(\hat{z}) = \max_{\{\boldsymbol{\mu}, n \geq 1\}} Y(\hat{z}, n, \boldsymbol{\mu}) - (n-1) \int w(z, \hat{z}, \boldsymbol{\mu}) \frac{dF_w(z)}{F_w(z_{\max})} - \chi(n-1) - \chi_f \tag{4.3}$$
$$\text{s.t.}(4.1),$$

where χ_f corresponds to the fixed cost of setting up a firm.

Wages. Workers' outside option is equal to the wage level $\overline{w}$, which is endogenous and adjusts to clear the labor market. The owner's outside option is equal to profits when producing with one fewer worker, similar to Stole and Zwiebel (1996). The surplus of the match is therefore a function of worker and entrepreneurial ability as well as task assignment $\boldsymbol{\mu}$: $y(z, \hat{z}, \boldsymbol{\mu}) - \overline{w}$.[58] The worker has bargaining power ω, his wage is

$$w(z, \hat{z}, \boldsymbol{\mu}) = (1-\omega)\overline{w} + \omega y(z, \hat{z}, \boldsymbol{\mu}) \tag{4.4}$$

Since task assignment maximizes match surplus, workers and owners agree on the choice of task assignment $\boldsymbol{\mu}$.

Occupational Choice. Each agent observes their ability z and chooses whether to be a worker or an entrepreneur. Profits are known, since firm owners hire a representative sample of workers. Wage earnings, on the other hand, depend on who the worker matches with.

[57] When estimating the model, we allow for entrepreneurial ability to affect worker output through a second channel that is independent of time allocation. This channel captures the fact that some entrepreneurial input may be non-rival, such as a design idea or good reputation, and hence not picked up by our time use data. Adding it to the model prevents overstating the role of specialization for pass-through. The full model in Appendix B includes this second channel; we present the model without it in the main text for expositional clarity.

[58] The hiring cost $\chi(n)$ is sunk and therefore not directly included in the surplus.

An individual with ability z starts a firm if and only if profits are higher than the expected wage in the labor market, where expectations are taken over which entrepreneur a worker matches with:

$$\mathbb{I}_o(z) = 1 \quad \iff \quad \pi(z) \geq \int w(z, \hat{z}, \boldsymbol{\mu}) \, \frac{(n(\hat{z}) - 1) \, dF_o(\hat{z})}{\int (n(\hat{z}) - 1) \, dF_o(\hat{z})} \tag{4.5}$$

4.3 Equilibrium

Finally, we define an equilibrium in our setting, which requires that all agents maximize and the wage level clears the labor market.

Definition of Competitive Equilibrium *A competitive equilibrium is a wage level* $\overline{w}$*, size and task assignment for each ability* $\hat{z}$ $\{n(\hat{z}), \boldsymbol{\mu}(\hat{z})\}_{\forall \hat{z}}$*, an occupational choice function* $\mathbb{I}_o(z)$*, and distributions* $F_o(z)$*,* $F_w(z)$ *such that:*

1. *firm owners choose size and task assignment to maximize profits as in* (4.3)*;*
2. *individuals choose their occupation according to* (4.5)*;*
3. *the labor market clears:* $\int (n(z) - 1) dF_o(z) = \int dF_w(z)$*;*
4. $F_o(z)$*,* $F_w(z)$ *are consistent with the occupational choice:* $F_w(z) = \int (1 - \mathbb{I}_o(z)) dG(z)$ *and* $F_o(z) = \int \mathbb{I}_o(z) dG(z)$*.*

4.4 Characterization

In this section, we analyze how the cost of specialization $\kappa(.)$ affects the allocation of talent within firms and consequently how it shapes firm size and aggregate productivity. We start by describing the occupational choice and then turn to the within-firm assignment problem and its implications for firm size and productivity. Finally, we discuss the properties of the economy's equilibrium. All proofs are in Appendix B.

We maintain Assumptions 1 and 2 throughout and additionally assume that the unbundling cost μ takes on the following functional form:

Assumption **3.** The cost of unbundling a fraction μ of complex tasks is given by $\kappa(\mu) = 1 - \exp\{-\hat{\kappa}(\mu)\}$, where $\hat{\kappa}(\mu) = {\kappa_0}^{1/\kappa_1} \frac{\mu^{1+1/\kappa_1}}{(1+1/\kappa_1)}$.

This guarantees closed-form solutions and allows us to parameterize the unbundling cost by a key parameter, κ_0.

Occupational Choice. Assumption 2 guarantees that entrepreneurship is more skill-sensitive than wage work. As a result, individuals with higher ability sort into starting firms.

LEMMA **1** (Occupational Choice)**.** *In equilibrium, there is a cutoff z_0 such that an individual z chooses to become an entrepreneur if and only if $z \geq z_0$.*

Labor Specialization. Each worker delegates complex tasks until the marginal unbundling cost equals the marginal benefit—the difference in abilities between worker and entrepreneur. Using (4.1), the optimal share delegated by employee z thus solves:

$$\log \hat{z} - \log z \ = \kappa_0^{1/\kappa_1} \mu(z, \hat{z})^{1/\kappa_1} \tag{4.6}$$

The level of specialization in the firm is governed by κ_0, while the curvature, κ_1, modulates the extent to which delegation depends on the ability gap between worker and owner. To map directly to the data, we now characterize the resulting time use of entrepreneurs and workers.

LEMMA **2** (Labor Specialization)**.** *Consider a firm of size n. The time spent on complex tasks by a worker z and an entrepreneur $\hat{z}$ is equal to*

$$\theta(z, \hat{z}) = \alpha \left(1 - \frac{1}{\kappa_0} (\log \hat{z} - \log z)^{\kappa_1}\right) \tag{4.7}$$

$$\hat{\theta}(\hat{z}, n) = \alpha \left(1 + \frac{n-1}{\kappa_0} \int (\log \hat{z} - \log z)^{\kappa_1} \frac{dF(z)}{F_w(z_{max})}\right). \tag{4.8}$$

Figure 8 illustrates the relationship between specialization, firm size, and the unbundling cost κ_0. For ease of exposition, we set $\kappa_1 = 0$, implying that the share of complex tasks each worker delegates to the entrepreneur is independent of ability. Both panels of Figure 8 plot the share of time spent on complex tasks as a function of firm size for workers of any z (left), and for entrepreneurs of any $\hat{z}$ (right).

The share of time each worker spends on complex tasks is independent of firm size. Since the entrepreneur has capacity left to take on complex tasks (Assumption 1), optimal delegation only depends on the unbundling cost. Entrepreneurs' time on complex tasks, however, is increasing in firm size. This relationship between size and specialization is mechanical: the bigger the firm, the more "low-hanging" complex tasks for the entrepreneur.

When the unbundling cost is higher ($\kappa_0 \uparrow$), each worker spends more time on his own complex tasks. For entrepreneurs, a higher unbundling cost only affects the slope of $\hat{\theta}(\hat{z}, n)$ with respect to size n. In a firm of size one, the unbundling cost has no impact on time allocation, since the entrepreneur is the sole worker. But each employee delegates fewer

Figure 8: Labor Specialization, Firm Size, and Unbundling Cost κ_0

(a) Workers Time in Complex Tasks (b) Entrepreneurs Time in Complex Tasks

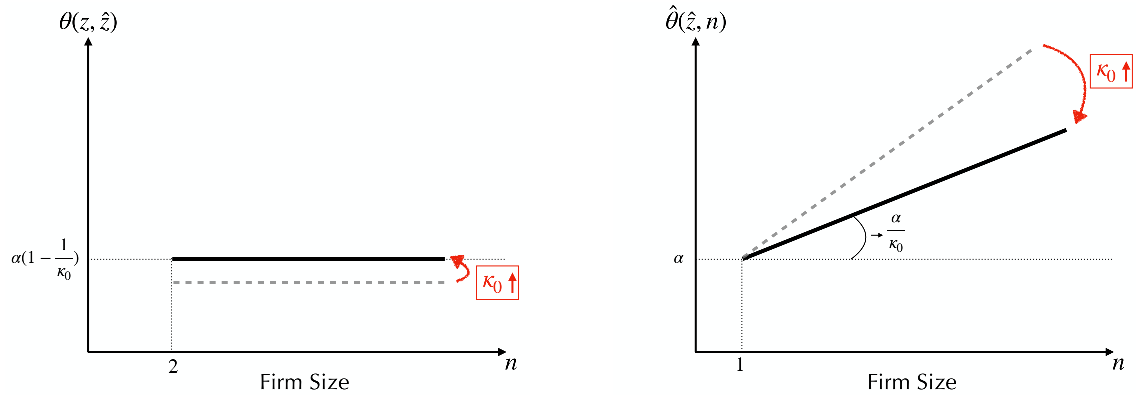

tasks, and hence the entrepreneurs' share of time spent on complex tasks rises more slowly with firm size. Average labor specialization—the difference between the right and the left panels—is therefore decreasing in κ_0, especially for large firms.

Firm Productivity. Conditional on the distribution of worker ability, firm productivity is pinned down by the within-firm assignment of tasks.

LEMMA **3** (Firm Productivity)**.** *The output of a firm of size n, run by an entrepreneur of ability $\hat{z}$, can be written as $Y(\hat{z}, n) = \mathbb{Z}(\hat{z}, n, \boldsymbol{\mu})\, n$ where*

$$\mathbb{Z}(\hat{z}, n, \boldsymbol{\mu}) = \frac{1}{n}\hat{z} + \underbrace{\frac{n-1}{n} \int \tilde{z}(z, \hat{z}, \boldsymbol{\mu}) \frac{dF_w(z)}{F_w(z_{max})}}_{\textit{dilution from firm size}},$$

$$\tilde{z}(z, \hat{z}, \boldsymbol{\mu}) = \hat{z}^{\mu(z,\hat{z})} z^{1-\mu(z,\hat{z})} \left(1 - \kappa(\mu(z, \hat{z}))\right),$$

$\mathbb{Z}(\hat{z}, n, \boldsymbol{\mu})$ is strictly decreasing in n as long as $\kappa_0 > 0$.

Firm productivity is equal to the average ability of all individuals doing the complex tasks. Since the entrepreneur completes all her own tasks, her productivity is $\hat{z}$. The task productivity of each employee, however, is lower than the entrepreneur's—as long as there is less than full specialization ($\kappa_0 > 0$). The entrepreneur's and her workers' task productivity is aggregated to the firm level with weights of $1/n$ and $(n-1)/n$. Increasing the size of the firm would therefore decrease its productivity: more weight is given to the lower productivity of workers. This highlights that low pass-through of entrepreneurial ability (due to high κ_0) leads to stronger decreasing returns to scale.

Optimal Firm Size. Labor specialization and firm size are closely intertwined. Lemma 2 showed one side of this two-way relationship: there is less labor specialization in small firms. Lemma 4 shows that there is also a reverse relationship: barriers to labor specialization reduce the optimal firm size.

LEMMA **4** (Firm Size)**.** *The optimal firm size n of an entrepreneur with ability $\hat{z}$ solves*

$$\left[\mathbb{Z}(\hat{z}, n, \boldsymbol{\mu}) + \underbrace{\frac{\partial \mathbb{Z}(\hat{z}, n, \boldsymbol{\mu})}{\partial n} n}_{\textit{prod. dilution}<0}\right] = \overline{w}(\hat{z}, \boldsymbol{\mu}) + \underbrace{\chi'(n)}_{\textit{hiring cost}} .$$

It is declining in the marginal hiring cost $\chi'(n)$ and in the unbundling cost κ_0.

At the optimal size, the marginal cost of hiring is equal to marginal revenue. The marginal cost equals the average wage plus the additional hiring cost. The first component of the marginal benefit is the standard increase in firm output from hiring an additional worker. The second component is unique to our framework. As shown in Lemma 3, as long as $\kappa_0 > 0$, firm-level productivity is decreasing in size, since each additional worker is less skilled than the entrepreneur and they are performing some of the firm's complex tasks. In choosing firm size, the entrepreneur takes into account the decreasing returns arising from productivity dilution driven by the unbundling cost.

Two frictions keep firms small. The hiring cost $\chi(n)$ directly reduces optimal firm size by making expansion costly. The unbundling cost κ_0 reduces firm size through productivity dilution and because it reduces average productivity for any firm bigger than 1. There is also an apparent complementarity between the two. The benefit from relaxing the external friction $\chi'(n)$ is limited if internal barriers to labor specialization generate strong decreasing returns to scale.

What Is a Firm? Two Polar Cases. To complete the intuition behind our model of the firm, Lemma 5 characterizes two polar cases.

LEMMA **5** (What is a Firm?)**.** *The the size of the unbundling cost (κ_0) spans two polar types of firms:*

1. ***Scalable Entrepreneurial Talent.*** *If $\kappa_0 = 0$, then $\frac{\partial \log y(z,\hat{z},\boldsymbol{\mu})}{\partial \hat{z}} = 1$, $Y(\hat{z}, n, \boldsymbol{\mu}) = \hat{z}\, n$, and optimal firm size is increasing in $\hat{z}$.*

2. ***Self-Employment within the Firm.*** *If $\kappa_0 \to \infty$, then $\frac{\partial \log y(z,\hat{z},\boldsymbol{\mu})}{\partial \hat{z}} = 0$, then $Y(\hat{z}, n, \boldsymbol{\mu}) = \overline{z}(\hat{z})\, n$, with $\overline{z}(\hat{z}) \equiv \frac{1}{n}\left(\hat{z} + \frac{n-1}{n} \int z \frac{dF_w(z)}{F_w(w_{max})}\right)$, and optimal firm size is constant in $\hat{z}$.*

When delegation is costless ($\kappa_0 = 0$), entrepreneurs fully pass through their ability to workers and firm productivity is equal to their ability. This benchmark resembles the typical firm problem dating back to Lucas (1978), in which labor is a commodity and firms are vehicles for leveraging the entrepreneur's talent. In this world, organizing labor into just a few large firms run by talented entrepreneurs yields large aggregate productivity gains.

In the opposite extreme, delegation is prohibitively costly ($\kappa_0 \to \infty$) and entrepreneurs cannot pass through any of their ability to workers. Individuals in the firm work independently, completing all tasks required for their production line and firm productivity is simply the average ability of all their members. In this benchmark, firms only exist to share fixed costs. As a result, the entrepreneur's identity is inconsequential and all firms are identical in size. Regardless of the firm size distribution, aggregate productivity closely tracks the overall distribution of talent in the population rather than the right tail of entrepreneurs.

Equilibrium and Aggregate Implications. So far, we have characterized the solution to the problem of one entrepreneur. Next, we turn to the overall economy. We prove the main proposition for the case of $\kappa_1 = 0$, where no assumptions on the population distribution of talent, $G(z)$, are required. Our estimation finds that κ_1 is small and the insights of Proposition 1 are relevant for thinking about the data.

PROPOSITION **1** (Aggregate Effects of the Unbundling Cost κ_0)**.** *Suppose that $\kappa_1 = 0$. As long as the aggregate labor supply curve is increasing in the wage level and ω is sufficiently small, a decline in κ_0 leads to an increase in:*

1. *the time each entrepreneur spends on complex tasks $\hat{\theta}(\hat{z}, n)$;*
2. *the slope of the relationship between the entrepreneur's time on complex tasks and firm size;*
3. *the average ability of firm owners;*
4. *the average firm size $\overline{n} \equiv \int n(z) \frac{dF_o(z)}{F_o(z_{max})}$, where $n(z)$ is the optimal firm size;*
5. *the average firm productivity $\overline{\mathbb{Z}} \equiv \int \mathbb{Z}(z, n(z), \boldsymbol{\mu}(z)) n(z) \frac{dF_o(z)}{F_o(z_{max})}$;*
6. *the wage $w(z, \hat{z}, \boldsymbol{\mu})$ of all workers z in all firms $\hat{z}$.*

Proposition 1 shows that reducing the unbundling cost transforms the way firms are organized internally with effects that ripple through the economy in equilibrium. Higher labor specialization increases firm productivity and thus the demand for labor. As a result, wages increase, leading some marginal firm owners to become workers. This further increases aggregate productivity through a classic selection effect. Overall, managerial ability is highly priced in the economy, as talent can be leveraged by taking over more and more complex tasks.

Next Steps: Using the Model. Lemma 5 provides two benchmark cases of the nature of the firm. In Section 6, we will assess where the actual economy stands between these two extremes, thus quantifying the productivity gains achieved by firms in our context relative to a world of pure self-employment. Proposition 1 further shows how reducing unbundling costs affects the entire economy, though it remains silent on the magnitude of these effects. The estimated model will allow us to bridge this gap and show that even increasing specialization to the levels observed in grain milling can yield meaningful aggregate gains.

5 Bringing the Model to the Data

We extend and parameterize the model for quantitative analysis. We then discuss model identification and show the estimation results. Appendix C.1 uses heterogeneity across sectors and regions to validate the theoretical predictions from Section 4.

5.1 Extensions and Parameterization

We extend the model along four dimensions. First, we augment the firm production function (4.1) to allow entrepreneurial ability to affect worker output independently of time allocation. In particular, the output produced by a worker z in a firm run by entrepreneur $\hat{z}$ who chooses task allocation $\boldsymbol{\mu}$ is given by

$$y(z, \hat{z}, \boldsymbol{\mu}) = \hat{z}^{\lambda} \quad \left[\hat{z}^{\mu(z,\hat{z})} z^{1-\mu(z,\hat{z})} (1 - \kappa(\mu(z, \hat{z})))\right]^{1-\lambda}. \tag{5.1}$$

Pass-through of entrepreneurial ability to worker productivity is now given by

$$\frac{\partial \log y(z, \hat{z}, \boldsymbol{\mu})}{\partial \log \hat{z}} = \lambda + (1 - \lambda)\mu(z, \hat{z}).$$

As long as $\lambda > 0$, worker productivity depends on the entrepreneur even in the absence of any labor specialization. Since this channel of pass-through does not require any time input from the owner, it is non-rival within the firm. λ captures, for example, the fact that the entrepreneur may have come up with a specific design, or has a great reputation. All workers in the firm benefit from this, even if they perform all complex tasks themselves. Adding this additional feature avoids overstating the role of specialization.

Second, to match the joint distribution of firm sizes and revenues, we allow entrepreneurs to differ not only in their managerial ability z but also in the hiring cost χ_0.[59] We assume that z is drawn from a generalized Pareto distribution with scale and location normalized to one and shape given by σ_z, and that χ_0 follows a normal distribution with mean $\overline{\chi_0}$ and standard deviation σ_χ.

Second, to match the time use within firms, we allow for a firm-level overhead of non-production tasks α_o. This overhead time must be supplied by the entrepreneur and does not affect productivity.

Third, we specify the functional form of the hiring cost. We assume that the entrepreneur has to pay a fixed cost to operate a firm χ_f, and an additional cost for hired labor $(n-1)$. The overall "hiring" cost for a firm of size n is therefore $\chi(n) = \chi_f + \chi_0^{1/\chi_1}(n-1)^{1+1/\chi_1}(1+1/\chi_1)^{-1}$. For simplicity, we still refer to the composite cost $\chi(n)$ as the hiring cost.

Last, we assume that the variable "managerial ability index" discussed in Section 3 is a noisy proxy of true managerial ability, denoted $s(z)$. Specifically, we let the (normalized) managerial index be equal to the (normalized) log of managerial ability plus an additive, normally distributed term.[60]

Table 3 summarizes the economic environment that we take to the data and links each economic block to the main parameters modulating it.

5.2 Targeted Moments and Identification

The model has 12 parameters. We specifically designed our survey to measure firms' start-up and fixed operating costs. We can thus directly calibrate χ_f.[61] The remaining 11 parameters do not have direct empirical counterparts and are jointly estimated.

Targeted moments. We target 150 moments, computed using pooled data for furniture and metal products, as explained in details in Appendix C.2. Table 4 lists 21 summary moments that capture the main relationships we are targeting.[62]

Our choice of moments is guided by two principles. First, the model should be consistent with the key features of the economic environment, as described in Sections 2 and 3.

[59] Matching this joint distribution is important since the dispersion of talent across individuals is a key driver of aggregate losses from barriers to labor specialization. If all potential entrepreneurs were of similar skills, the inability of relatively high-skilled entrepreneurs to leverage their talent would be inconsequential.

[60] This assumption allows us to accommodate enough heterogeneity in managerial ability to match the empirical distribution of log revenues while also matching the observed empirical relationship between firm revenue, workers' earnings, and the managerial index.

[61] See Appendix C.2 for details.

[62] For example, while we target the deciles of the distributions of firm sizes and revenue, we include in Table 4 only their means and standard deviations.

Table 3: Summary of the Economic Environment and Parameters

	Equation	Parameters
Final Output	$Y = \int Y(\hat{z}) dF_o(\hat{z})$	
Firm Output	$Y(\hat{z}) = \mathbb{Z}(\hat{z}, n(\hat{z}), \boldsymbol{\mu})\, n(\hat{z})$	
Firm Productivity	$\mathbb{Z}(\hat{z}, n, \boldsymbol{\mu}) = \hat{z}^{\lambda} \left(\frac{1}{n} \hat{z}^{1-\lambda} + \frac{n-1}{n} \int \tilde{z}(z, \hat{z}, \boldsymbol{\mu})^{1-\lambda} dF_w(z) \right)^{\frac{1}{1-\lambda}}$	λ
Net Task Productivity	$\tilde{z}(z, \hat{z}, \boldsymbol{\mu}) = \hat{z}^{\mu(z,\hat{z})} z^{1-\mu(z,\hat{z})} \left[1 - \kappa\left(\mu(z, \hat{z})\right)\right]$	κ_0, κ_1
Heterogeneity	$\log z \sim N(1, \sigma_z)$, $\chi_0 \sim N(\overline{\chi_0}, \sigma_\chi)$	$\sigma_z, \overline{\chi_0}, \sigma_\chi$
Unbundling Cost	$\kappa(\mu) = 1 - \exp\left\{ -{\kappa_0}^{1/\kappa_1}\, \mu^{1+1/\kappa_1} (1 + 1/\kappa_1)^{-1} \right\}$	κ_0, κ_1
Hiring Cost	$\chi(n) = \chi_f + \chi_0^{1/\chi_1}\, n^{1+1/\chi_1}\, (1 + 1/\chi_1)^{-1}$	$\chi_f, \overline{\chi_0}, \sigma_\chi, \chi_1$
Worker Earnings	$w(z, \hat{z}, \boldsymbol{\mu}) = (1-\omega)\overline{w} + \omega \hat{z}^{\lambda} \tilde{z}(z, \hat{z}, \boldsymbol{\mu})^{1-\lambda}$	ω
Measurement Error	$s(z) = \log z + \epsilon$, $\epsilon \sim N(0, \sigma_\epsilon)$	σ_ϵ
Complex Share (Workers)	$\theta(z, \hat{z}) \equiv \alpha(1 - \mu(z, \hat{z})) = \alpha \left(1 - \frac{1}{\kappa_0} (\log \hat{z} - \log z)^{\kappa_1} \right)$	$\alpha, \kappa_0, \kappa_1$
Complex Share (Entrepr.)	$\hat{\theta}(\hat{z}, n) = \alpha_o + \alpha \left(1 + (n-1) \frac{1}{\kappa_0} \int (\log \hat{z} - \log z)^{\kappa_1} \frac{dF_w(z)}{F_w(z_{max})} \right)$	$\alpha_o, \alpha, \kappa_0, \kappa_1$

Therefore, we target a rich set of moments describing time allocation within the firm as well as heterogeneity across firms in terms of size and productivity.[63] Second, we need to include enough targets to be able to identify the parameters modulating the returns to labor specialization. For this purpose, as we explain below, it is important to include moments on the distribution of workers' earnings within and between firms, and their relationship with our measure of managerial ability.

Estimation Procedure. We estimate the model using indirect inference and simulated method of moments. We minimize the distance between data moments and their exact model counterparts using a routine that we developed in Bassi, Muoio, Porzio, Sen, and

[63] Computing the moments does not pose any complications either in the model or in the data. Only a few simple decisions are to be made. First, we need to define what complex tasks are in the data. We assume, following the evidence discussed, that non-production tasks are more complex. Second, we need to decide whether to purge the data of some variation. Here, again, we closely follow the empirical section and use the same set of controls. Finally, when calculating the distribution of firm revenue and workers' earnings, we trim the top and bottom 5% to get rid of excessive variation plausibly driven by measurement error.

Tugume (2022b). Details are in the Supplemental Appendix, where we also show that the parameters are well-identified: the likelihood function is single-peaked around the estimates, and we verify that our estimation procedure recovers the true parameters when we run it on a synthetic set of moments generated by the model itself.

Identification. While all the parameters are jointly estimated, we can provide a heuristic identification argument, which we verify by computing the Jacobian matrix that traces out how each moment is affected by each parameter. The matrix is in the Supplemental Appendix, but in the last column of Table 4, we include the *key parameters* that are linked to each moment.

As Lemma 2 highlights, the within-firm allocation of time is tightly linked to the share of complex tasks in production (α), the overhead time (α_o), and the parameters of the unbundling cost (κ_0, κ_1). Our unique data on the relationship between firm size and the entrepreneur's time spent on complex tasks identifies κ_0. This same relationship, but estimated for employees, helps to pin down κ_1. In equilibrium, larger firms are managed by more skilled entrepreneurs; hence, if κ_1 is large, workers in large firms should spend less time on complex tasks. In the data, however, the relationship is flat, suggesting that κ_1 is small.

The second block of moments includes statistics on workers' earnings to pin down λ and ω, which modulate the pass-through of entrepreneurial ability to worker productivity and the relationship between productivity and wages.[64] In the data, we do not directly observe worker productivity, hence we have to rely on the fact that workers' earnings are increasing in productivity due to how wages are set in this context. Targeting the distribution of wages within and between firms allows us to separately identify λ and ω. The intuition is as follows. When ω is high, the variance of wages is high overall, both across and within firms. With a high λ, on the other hand, only the variance across firms is high, but the one within firms is low.[65] In practice, since earnings are measured with error, we do not directly match their variance. Rather, we target the relationships between worker earnings and firm characteristics, as well as the average earnings gap across more versus less productive firms, normalized by their standard deviation.

The third block of moments includes the distribution of firm revenue and its relationship with the managerial ability index. These moments discipline the variance of managerial

[64]It is important to emphasize that a low κ_0 has the same effect on pass-through as a high λ. Therefore, our identification strategy relies on the assumed relationship between time spent on complex tasks and productivity.

[65]For example, if $\omega = 0$ all workers are paid the same, while if $\lambda = 1$ all workers within a firm are paid the same, but there is large wage heterogeneity across firms.

talent (σ_z) and the noise term (σ_ϵ) in our empirical proxy. σ_z increases the variance of productivity and revenues. Given σ_z, a large σ_ϵ flattens the relationship between revenue per worker and the managerial index due to attenuation bias.

Finally, the last block of moments pin down the parameters of the hiring cost (χ_0, χ_1, σ_χ) since—as shown in Lemma 4— these parameters directly map to the firm size distribution and its relationship with managerial ability.

Importance of Time Use Data. Time use data is crucial for identifying barriers to labor specialization within the firm separately from any other constraints that keep firms small. Even observing aggregate measures of specialization would not be enough, since it would be impossible to distinguish whether firms are small because they are not specialized or whether they are not specialized because there is not enough scope for specialization given their small size. Our data shows that specialization increases weakly with firm size, which implies a large barrier to specialization (κ_0). Then, given κ_0 and the other parameters, the distribution of χ_0 is chosen to match the firm size distribution.

Figure 9: Model Fit for Firm Heterogeneity and Time Allocation

(a) Firm Sizes (b) Firm Revenues (c) Time in Complex Tasks

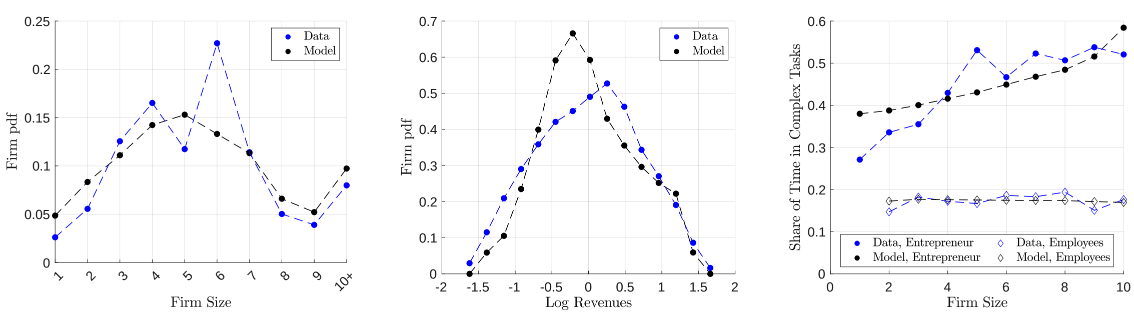

Notes: The figure compares empirical moments, in blue, with their model-generated counterparts.

5.3 Estimation Results and Model Fit

The model matches the data well, as Table 4 shows. Figure 9 illustrates the fit for some key moments: the model matches the heterogeneity between firms in terms of size and revenue, as well as the time allocation within firms.[66]

Table 5 includes the estimated values of all parameters. A few are worthwhile to discuss. First, the value of λ implies that the entrepreneur is able to pass through approximately

[66] We describe the model fit for all 150 moments in the Supplemental Appendix.

Table 4: Summary of Targeted Moments and Model Fit

Moments	Data	Model	Param.
A. Allocation of Time to Complex Tasks			
(i) Average Time on Complex Tasks	0.234	0.235	α
(ii) Average of Entrepreneurs	0.457	0.449	α_o, α
(iii) Average for Self-Employed	0.341	0.353	α_o, α
(iv) Average of Low-Skilled Workers	0.137	0.172	α
(v) Average of High-Skilled Workers	0.217	0.177	α
(vi) Slope w/ Size (Entrepreneur)	0.021	0.022	κ_0, κ_1
(vii) Slope w/ Size (Low-Skilled Workers)	0.002	-0.001	κ_1, ω, χ_1
(viii) Slope w/ Size (High-Skilled Workers)	0	-0.001	κ_1, ω, χ_1
(ix) Slope w/ Log(Earn) (All Workers)	0.033	0.009	κ_1
B. Distribution of Earnings w/i and b/w Firms			
(i) Log(Earn) on Man. Ability (Normalized)	0.187	0.199	$\sigma_\epsilon, \lambda, \omega$
(ii) Log(Earn) on Log(Rev p.w.)	0.191	0.200	$\lambda, \omega, \sigma_\chi$
(iii) Norm. Earn Gap by Rev p.w.	0.389	0.655	$\lambda, \chi_1, \overline{\chi_0}$
(iv) Norm. Earn Gap by Man. Ability	0.137	0.330	$\omega, \chi_1, \sigma_\epsilon$
C. Distribution of Firm Revenues			
(i) Std of Log(Rev)	0.726	0.632	ω, χ_1, σ_z
(ii) Log(Rev p.w) on Man. Ability	0.145	0.145	$\sigma_\epsilon, \sigma_z, \overline{\chi_0}$
(iii) Log(Rev) Gap by Man. Ability	0.305	0.382	$\sigma_\epsilon, \omega, \overline{\chi_0}, \sigma_z$
D. Firm Size Distribution			
(i) Average Size	5.701	5.559	$\omega, \chi_1, \overline{\chi_0}$
(ii) Std of Log(Size)	0.489	0.571	$\sigma_\chi, \overline{\chi_0}$
(iii) Std of Size	2.263	2.522	σ_χ, χ_1
(iv) Log(Size) on Man. Ability	0.100	0.086	$\chi_1, \omega, \sigma_\epsilon$
(v) Size Gap by Man. Ability	0.275	0.310	$\omega, \chi_1, \sigma_\epsilon$

Notes: Empirical moments used in estimation and corresponding values in the model, together with the key parameters relating to each moment. For details of the computation of the empirical moments, see Appendix C.2.

Table 5: List of Parameters and their Estimated Values

Param.	Value	Param.	Value	Param.	Value
χ_f	$0.1\overline{\pi}(z)$	κ_0	0.081^{-1}	κ_1	0.707
λ	0.201	ω	0.472	$\overline{\chi_0}$	13.931
χ_1	0.601	σ_χ	4.780	σ_z	0.912
α	0.190	α_0	0.185	σ_ϵ	2.144

20% of her ability to her workers. To put this number in perspective, we can compare it with the productivity pass-through due to vertical specialization. In the estimated model, we find that, on average, the typical worker completes approximately 90% of her complex tasks. Given the value of λ, this means that the productivity pass-through due to vertical specialization is $\sim 6\%$—that is, a bit less than a third of the direct pass-through due to λ.

Second, the value of ω implies a prominent role for the piece-rate component of workers' earnings, consistent with the evidence described in Section 3: more productive workers are compensated for around half of their higher output.[67]

Third, we estimate large heterogeneity in managerial ability at the top of the distribution. The estimated value of the shape parameter σ_z implies that ability at the 98^{th} percentiles is approximately 8 times that at the 80^{th} percentiles.[68]

Finally, we find that κ_1 is small, which shows that the estimated model is close to the case of $\kappa_1 = 0$ analyzed in Proposition 1.[69]

6 Quantification

We now use the estimated model to quantify the importance of firms in shaping aggregate productivity in our setting.

We do so through two exercises. First, we calibrate combinations of entrepreneurial pass-through and hiring costs that hold constant either total net output or average firm size. These exercises are conceptual tools to assess how far the economy is from the benchmark case of self-employment within the firm—and, more broadly, to quantify how much the presence of firms contributes to aggregate productivity.

Second, we conduct counterfactuals to shed light on the potential effects of changing unbundling or hiring costs within a range that is relevant for policy interventions. To do so, we change these costs to replicate observed sectoral differences in specialization and firm size respectively.

[67] The estimated value of ω satisfies Assumption 2 and hence the single-crossing holds.

[68] The 80^{th} percentiles of the ability distribution correspond roughly to the marginal entrepreneur, given an average firm size ~ 6.

[69] One way to assess the magnitude of κ_1 is to compare the share of complex tasks completed by low- and high-skilled workers. We find that, on average, a worker at the 10^{th} percentile of the distribution completes $\sim 91\%$ of his complex tasks, whereas a workers at the 90^{th} percentile completes $\sim 94\%$.

6.1 Quantifying Self-Employment within the Firm

A key result from Section 5 is that entrepreneurs pass through approximately 25% of their ability to workers. This means that if entrepreneurs are twice as skilled as workers, having workers employed in firms instead of self-employed would raise their productivity by about 25%. However, the level of pass-through is not sufficient to quantify the full contribution of firms to aggregate productivity, as the total effect also depends on how pass-through determines equilibrium decisions on firm entry and size.

We therefore turn to the full quantitative model to answer two related questions. First, how different would aggregate output be if the economy consisted of firms with the same average size as we observe in the data, but different internal organizations? This exercise allows us to isolate the role of within-firm specialization in shaping aggregate productivity and assess how close the current equilibrium is to the two polar cases described in Lemma 5: "self-employment within the firm" and "scalable entrepreneurial talent." Second, we ask: what is the "true" size of firms in our context? That is, what firm size distribution would be required to generate the same aggregate output in a world where firms are fully specialized and thus serve as perfect vehicles for leveraging entrepreneurial ability?

To answer these questions, we vary the within-firm pass-through of entrepreneurial ability from 0% to 100%—thus ranging from one extreme in which there is no specialization and $\lambda = 0$ to the opposite in which there is either full specialization or $\lambda = 1$—and recalibrate the average hiring cost $\overline{\chi_0}$ to either (i) hold average firm size constant or (ii) hold aggregate net output constant.[70] Net output refers to total output net of hiring and entry costs (i.e., aggregate consumption).

Quasi-*Self-employment within the Firm.* Figure 10a presents the results of the first exercise, which holds average firm size constant at the benchmark (~ 5.6 workers) as we vary the pass-through. In an economy with zero pass-through, output would be just 20% lower. By contrast, increasing pass-through to 100%, while still holding average firm size constant, would yield nearly three times more output. This suggests that the economy is much closer to self-employment within the firm than to a world of scalable entrepreneurial talent.

By construction, these output gains occur without any change in selection into entrepreneurship: since average firm size is held constant, so are the number of entrepreneurs

[70] In both exercises, we recalibrate the mean of the hiring cost distribution while allowing its standard deviation to scale proportionally—multiplying both $\overline{\chi_0}$ and σ_χ by the same factor. In the second exercise, we start from a pass-through value of 15% because at zero pass-through, it is not possible to generate the same level of aggregate net output.

and their average productivity. Instead, two mechanisms drive the gains: (i) workers become more productive as they inherit more of their entrepreneur's ability, and (ii) the allocation of talent improves, as workers increasingly sort towards higher-ability entrepreneurs. Although average firm size is fixed, *dispersion* in firm size increases. This occurs because higher pass-through increases the returns to working for more talented entrepreneurs.[71]

Firms Are Smaller than They Seem. The second exercise, which holds net output constant tells a complementary story. Figure 10b shows that an economy with full pass-through could generate the same level of net output with an average firm size of just 1.6 workers.

To further illustrate this point, Figure 10c compares the firm size distributions in three economies that all achieve the same aggregate net output: the benchmark, one with double the estimated pass-through, and one with full pass-through. In the benchmark economy, more than half of firms have four or more workers.[72] Yet in this setting with low specialization, limited pass-through, and small fixed costs, firm size is a poor measure of the allocation of talent and productivity. The same net output could be produced in an economy where firms employ fewer than one employee on average, provided those employees fully inherit their entrepreneur's productivity. In short, the economy functions almost as if most workers were self-employed: the effective size of firms is far smaller than headcounts suggest.

Appendix Figure D.1 shows how the estimated parameters drive this result. Low pass-through—due to both limited specialization and a low λ—keeps worker productivity low in the benchmark. In addition, the small fixed costs we measure imply that the entry of many small firms is not particularly wasteful. Finally, the large heterogeneity in entrepreneurial ability z amplifies the productivity cost of low pass-through in the benchmark. Although marginal entrants reduce average entrepreneurial quality and raise total entry costs, these losses are modest and are more than offset by the gains in worker productivity that arise from higher pass-through in the counterfactual.

Together, these two exercises highlight the *main message of the paper*: even in sectors such as furniture and metal products in urban Uganda—where firms appear relatively large—their internal organization makes them effectively small and close to the benchmark

[71] Appendix Figure D.1 illustrates in more detail how higher pass-through reshapes the economy. Entrepreneurial talent becomes more strongly rewarded, while workers become more interchangeable since their productivity depends increasingly on who they work for. As a result, wage differences compress, while the dispersion of both entrepreneurial returns and firm sizes increases.

[72] Recall that the entrepreneur is included in the firm's worker count, so a self-employed individual operates a firm of size one.

of *self-employment within the firm*. The limited ability of entrepreneurs to scale their talent depresses aggregate productivity and constrains the potential of these types of firms as engines of development.

External Barriers May Be Smaller than We Think. A corollary of internal barriers to firm size is that the external hiring costs needed to rationalize the observed average firm size are smaller than they might appear. As Lemma 4 shows, low labor specialization, and hence low pass-through, generates diminishing returns within the firm. The stronger this internal friction, the lower the external costs needed to explain firm size, given entrepreneurial ability. Indeed, Appendix Figure D.1 shows that rationalizing the same average firm size under full pass-through would require hiring costs roughly five times higher.

6.2 Aggregate Gains of Reducing Unbundling and Hiring Costs

We now estimate the aggregate effects of reducing unbundling and hiring costs. Relative to the previous exercises, we consider smaller changes in these costs—within a range that we observe empirically—and perform standard counterfactuals. That is, we do not re-calibrate other parameters to hold size or output fixed.

The objective of these exercises is to measure the gains from potential policies that reduce unbundling costs, hiring costs, or both. To be clear, we do not take a stand on how such changes might be implemented in practice. In our model, unbundling and hiring costs are a reduced-form representation of any barriers that limit specialization and firm size, respectively. Still, the counterfactuals are policy-relevant in three ways: (i) they show that reductions in internal frictions, which is the focus of our paper, can have quantitatively large effects; (ii) they highlight that the gains from removing external barriers, which have been the focus of the literature, depend on the magnitude of internal barriers; and (iii) they show that distinguishing between internal and external barriers matters, as each type has qualitatively and quantitatively distinct effects on aggregate outcomes.

The benchmark model is calibrated to the furniture and metal sectors. For the counterfactuals, we reduce internal barriers to match the extent of vertical specialization observed in grain milling, which we interpret as an "attainable" benchmark in our setting. Specifically, we recalibrate κ_0 to target alternative relationships between firm size and the time entrepreneurs spend on complex tasks, holding all other parameters fixed. Figure 11a shows how varying κ_0 produces different slopes, ranging from 0 (no vertical specialization), to 0.21 (the value estimated for furniture and metal), to 0.63 (the value estimated for grain milling).

Reducing the unbundling cost leads to an increase in firm size, as predicted by Proposition

Figure 10: Quantifying Self-employment within the Firm

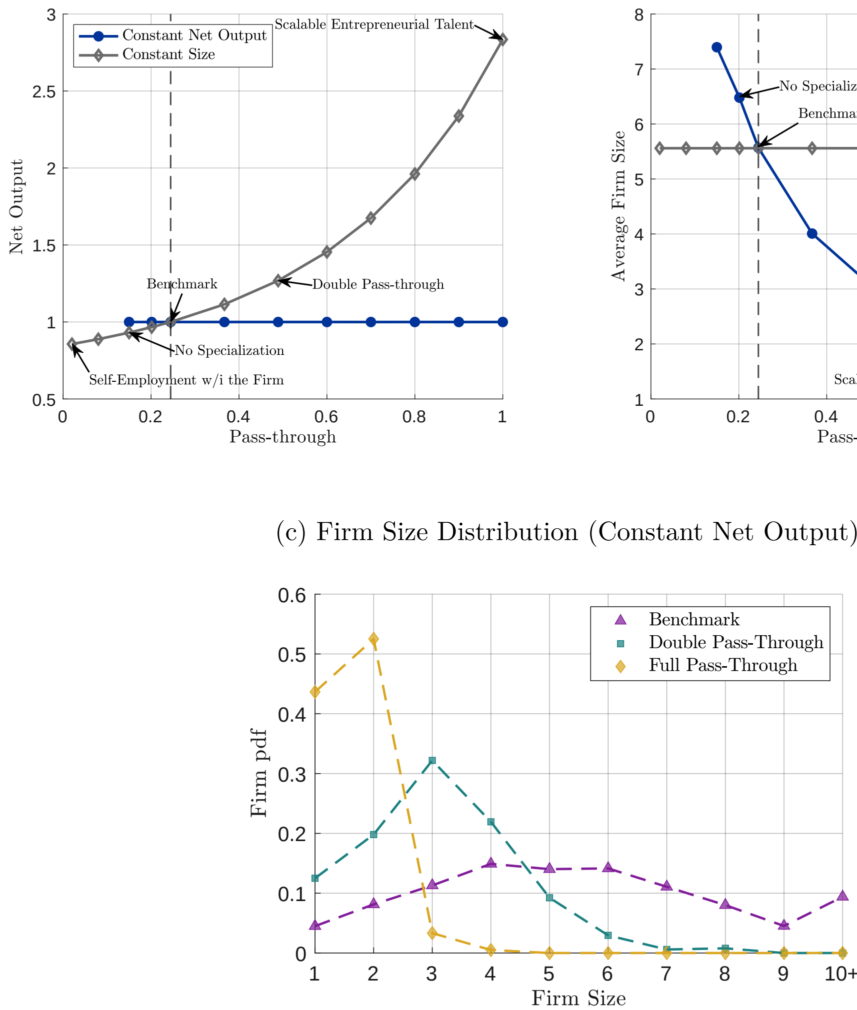

(a) Net Output (Constant Firm Size)

(b) Firm Size (Constant Net Output)

(c) Firm Size Distribution (Constant Net Output)

Notes: The top two panels show net output and average firm size across economies in which we vary the pass-through of entrepreneurial ability, recalibrating the hiring cost to hold either average firm size or net output constant. The bottom panel compares the firm size distribution in the benchmark economy with two alternative economies—one with full pass-through and one with double the estimated pass-through—both calibrated to achieve the same level of net output.

1. To contrast the effects of reducing internal barriers with a more conventional counterfactual that lowers external barriers, we also recalibrate the hiring cost to generate the same increase in firm size (as shown in Figure 11b).[73] This allows us to compare how identical

[73] In practice we solve the model for a vector of unbundling costs spanning from no specialization to the specialization of grain-milling. For each value of this unbundling cost, we pick a hiring cost to match the same average firm size while keeping all other parameters fixed.

changes in firm size—driven by reductions in either internal or external frictions—differ in their effects on aggregate outcomes.

Figure 12 plots key aggregate outcomes against firm size for both sets of counterfactuals: those varying unbundling costs (in green) and those varying hiring costs (in orange).

Figure 11: Calibration of Unbundling and Hiring Costs

(a) Unbundling Cost, κ_0 (b) Hiring Cost, $\overline{\chi}_0$

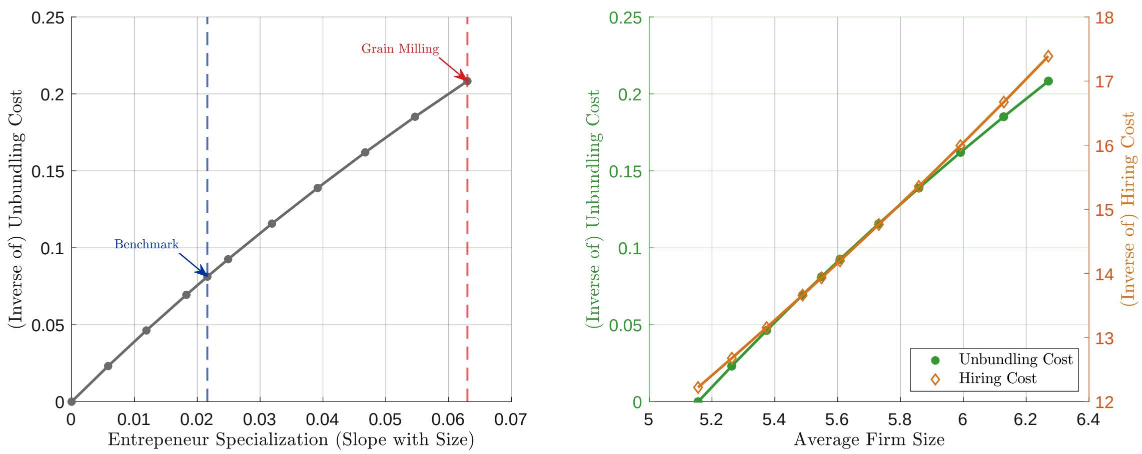

Notes: The left panel shows the relationship between the inverse of the unbundling cost (κ_0^{-1}) and the slope of the relationship between firm size and the share of time entrepreneurs spend on complex tasks—a key measure of entrepreneurial specialization. The right panel shows how we target the same average firm sizes by varying either the unbundling cost (left y-axis) or the hiring cost (right y-axis).

Large Gains from Reducing Unbundling Costs. Lowering the unbundling cost to the level of grain milling leads firms to grow significantly—by roughly 2/3 of a worker on average, or about a 13% increase (see Figure 11b). As shown in Figure 12a, this cost reduction increases net output by 10%. The gains arise through three main channels. First, selection into entrepreneurship improves: fewer, more able entrepreneurs start firms (Figure 12b). Second, worker productivity rises, due both to improved sorting and higher pass-through of entrepreneurial ability (Figure 12c). This second channel is new to our model and turns out to be quantitatively more important than the classic selection effect. Third, there is an improvement in the allocation of talent across firms since more productive entrepreneurs increase their firm size relatively more (Figure 12d).

Reducing Internal Frictions Leads to Larger Gains. We compare these results with those obtained by lowering the hiring cost to match the same average firm size. While both counterfactuals exhibit similar qualitative patterns—larger firms, higher output, and more

Figure 12: Aggregate Effects of Unbundling and Hiring Costs

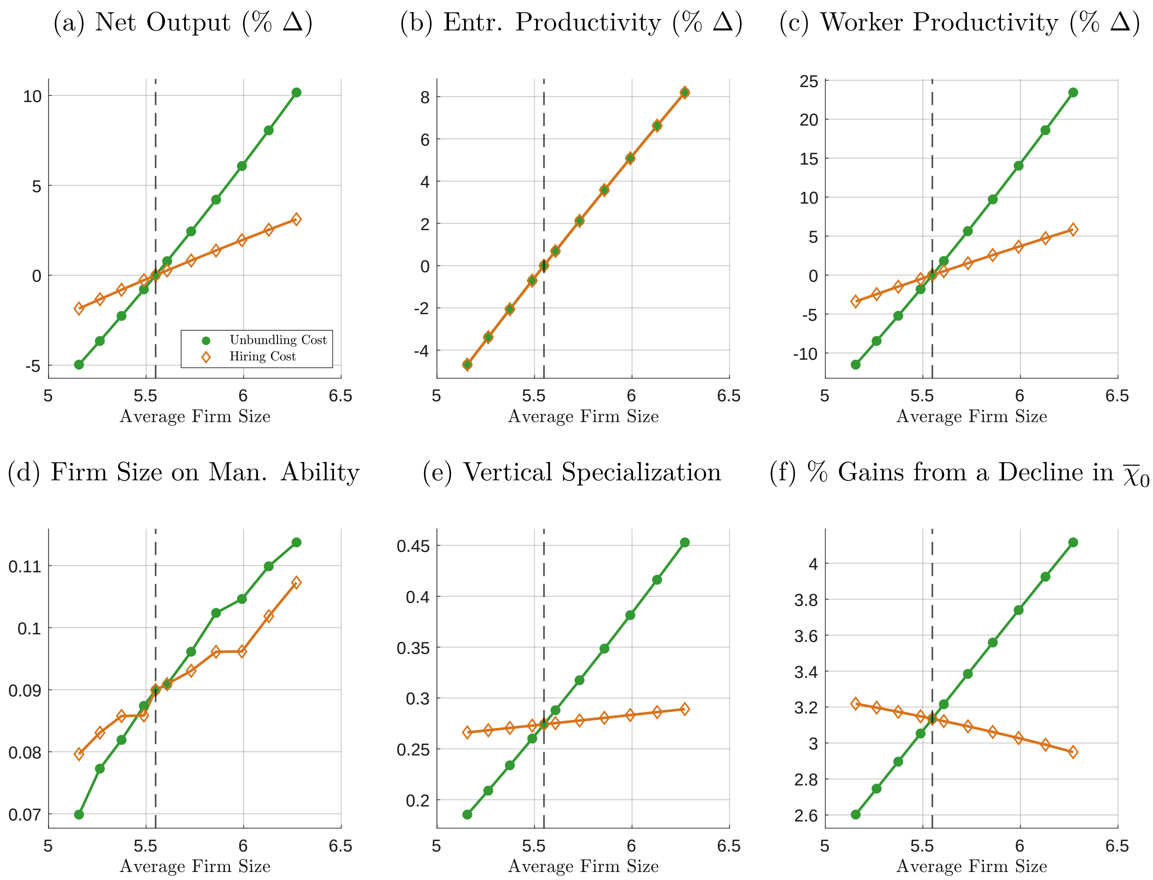

Notes: This figure presents key aggregate statistics for economies that differ in their average firm sizes. Differences in firm size are generated either by changing unbundling costs (green line) or hiring costs (orange line). The vertical dashed line represents the benchmark economy.

specialization—the gains are consistently larger when firm size increases due to lower internal barriers. For example, the same increase in firm size leads to only a 3% rise in output when driven by a reduction in hiring costs, compared to 10% in the case of lower unbundling costs (Figure 12a).

Figure 12e helps explain why. When firm size increases due to a reduction in hiring costs, specialization rises only modestly. This is because external frictions affect scale, but not the internal organization of production. The resulting larger firms remain relatively unspecialized and unproductive, since entrepreneurial productivity is diluted across many low-productivity workers.[74]

[74]This result is informative in itself: while our model allows for causality to run both from labor specialization to firm size and from firm size to specialization, the former channel is quantitatively stronger. In this

Understanding Why Firms Are Small Matters for Policy. Figure 12f illustrates how the returns to policies aimed at increasing firm size depend critically on the within-firm organization of labor. In our framework, such policies are captured by a reduction in the reduced-form "hiring cost" χ. We simulate a 20% reduction in this cost, starting from economies with different baseline levels of unbundling costs, and thus different degrees of specialization and average firm size.

The results show that internal and external barriers are complementary: while lowering hiring costs raises output in all cases, the gains are substantially larger when unbundling costs are low. The intuition is simple: greater specialization allows firms to expand without quickly hitting diminishing returns.

Last, Figure 12f compares the effects of reducing hiring costs across economies that differ in their baseline level of hiring costs. In this case, the pattern reverses: net output gains *decline* with average firm size. The reason is that larger firms that remain poorly organized struggle to translate scale into productivity, so the marginal return to expanding them further is low. In contrast, economies or sectors with more specialized firms can scale more effectively and thus respond more strongly to external interventions.

7 Conclusion

This paper offers a new perspective on how firms operate in low-income countries by focusing on their internal organization. Combining a novel time-use survey of manufacturing firms in Uganda with an equilibrium model, we show that firms in this context are largely self-employment in disguise, despite an average firm size of close to six workers. Organizing labor into this type of firms yields only modest productivity gains over a world of literal self-employment.

Our findings suggest that, to achieve economic growth, economies need different firms altogether: scalable organizations with specialized labor that can enable talented entrepreneurs to amplify their skills and become engines of development. Future research and policy should focus on understanding why such firms are missing and what it would take to foster their emergence. Our results suggest that the prevalence of customized production is one important factor limiting labor specialization. One promising direction is therefore to study the role of demand-side interventions in facilitating the transition to scalable organizations.

sense, our analysis is more consistent with the notion that firms are small because they are not specialized than with the idea that they are not specialized because they are small.

References

Abebe, G., S. A. Caria, M. Fafchamps, P. Falco, S. Franklin, S. Quinn, and F. J. Shilpi (2023): "Matching frictions and distorted beliefs: Evidence from a job fair experiment," Tech. rep., Working Paper. A.2

Adhvaryu, A., V. Bassi, A. Nyshadham, J. Tamayo, and N. Torres (2023): "Organizational Responses to Product Cycles," *Available at SSRN 4403515*. 9

Adhvaryu, A., V. Bassi, A. Nyshadham, and J. A. Tamayo (2024): "No line left behind: Assortative matching inside the firm," Tech. rep. 9

Agness, D., T. Baseler, S. Chassang, P. Dupas, and E. Snowberg (2025): "Valuing the time of the self-employed," *Review of Economic Studies*, rdaf003. 1

Akcigit, U., H. Alp, and M. Peters (2021): "Lack of selection and limits to delegation: firm dynamics in developing countries," *American Economic Review*, 111, 231–275. 3.3.1

Alchian, A. A. and H. Demsetz (1972): "Production, information costs, and economic organization," *The American Economic Review*, 62, 777–795. 10

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, 2369–2414. A.2

Anderson, S. J. and D. McKenzie (2022): "Improving business practices and the boundary of the entrepreneur: A randomized experiment comparing training, consulting, insourcing, and outsourcing," *Journal of Political Economy*, 130, 157–209. 13, A.2, A.2

Atack, J., R. A. Margo, and P. Rhode (2023): "De-skilling: Evidence from Late Nineteenth Century American Manufacturing," Tech. rep., National Bureau of Economic Research. 44

Atack, J., R. A. Margo, and P. W. Rhode (2019): ""Automation" of manufacturing in the late nineteenth century: The hand and machine labor study," *Journal of Economic Perspectives*, 33, 51–70. 44

Atencio-De-Leon, A., M. Lee, and C. Macaluso (2023): "Does Turnover Inhibit Specialization? Evidence from a Skill Survey in Peru," . 3.3.1

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017a): "Organizational barriers to technology adoption: Evidence from soccer-ball producers in Pakistan," *The Quarterly Journal of Economics*, 132, 1101–1164. 9

Atkin, D., A. Costinot, and M. Fukui (2021): "Globalization and the Ladder of Development: Pushed to the Top or Held at the Bottom?" Tech. rep., National Bureau of Economic Research. 51

Atkin, D., A. K. Khandelwal, and A. Osman (2017b): "Exporting and Firm Performance: Evidence from a Randomized Experiment*," *The Quarterly Journal of Economics*, 132, 551–615. A.2, A.3

Bandiera, O., A. Elsayed, A. Heil, and A. Smurra (2022): "Economic Development and the Organisation of Labour: Evidence from the Jobs of the World Project," *Journal of the European Economic Association*, 20, 2226–2270. 1, 1, 9

Bandiera, O., A. Prat, S. Hansen, and R. Sadun (2020): "CEO behavior and firm performance," *Journal of Political Economy*, 128, 1325–1369. 11

Banerjee, A., E. Breza, E. Duflo, and C. Kinnan (2024): "Can Microfinance Unlock a Poverty Trap for Some Entrepreneurs?" Working Paper 26346, National Bureau of Economic Research. 1

Bassi, V., M. E. Kahn, N. Lozano Gracia, T. Porzio, and J. Sorin (2022a): "Searching for Customers, Finding Pollution," Tech. rep., mimeo. 24

Bassi, V., R. Muoio, T. Porzio, R. Sen, and E. Tugume (2022b): "Achieving scale collectively," *Econometrica*, 90, 2937–2978. 8, 1, 20, 24, 3.3.1, 5.2, E.3, E.6

Bassi, V. and A. Nansamba (2022): "Screening and Signalling Non-Cognitive Skills: Experimental Evidence from Uganda," *The Economic Journal*, 132, 471–511. A.2

Bassi, V., T. Porzio, R. Sen, and E. Tugume (2021): "The Impact of the COVID-19 Lockdown on SMEs and Employment Relationships in Uganda," Tech. rep., International Growth Centre Policy Brief UGA-20112. 100

Becker, G. S. and K. M. Murphy (1992): "The division of labor, coordination costs, and knowledge," *The Quarterly journal of economics*, 107, 1137–1160. 10

Bennedsen, M., K. M. Nielsen, F. Pérez-González, and D. Wolfenzon (2007): "Inside the family firm: The role of families in succession decisions and performance," *The Quarterly Journal of Economics*, 122, 647–691. 3.3.1

Berge, L. I. O., K. Bjorvatn, and B. Tungodden (2015): "Human and Financial Capital for Microenterprise Development: Evidence from a Field and Lab Experiment," *Management Science*, 61, 707–722. A.2

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2013): "Does Management Matter? Evidence from India," *The Quarterly Journal of Economics*, 128, 1–51. 3.3.1, A.2

Bloom, N., L. Iacovone, M. Pereira-López, and J. Van Reenen (2022): "Management and misallocation in Mexico," Tech. rep., National Bureau of Economic Research. 1

Bloom, N., R. Sadun, and J. Van Reenen (2012): "The organization of firms across countries," *The Quarterly Journal of Economics*, 127, 1663–1705. 9

Boehm, J. and E. Oberfield (2023): "Growth and the Fragmentation of Production," Tech. rep., mimeo. 9

Bold, T., S. Ghisolfi, F. Nsonzi, and J. Svensson (2022): "Market Access and Quality Upgrading: Evidence from Four Field Experiments," *American Economic Review*, 112, 2518–52. 1

Bolton, P. and M. Dewatripont (1994): "The firm as a communication network," *The Quarterly Journal of Economics*, 109, 809–839. 10

Breza, E., S. Kaur, and Y. Shamdasani (2021): "Labor Rationing," *American Economic Review*, 111, 3184–3224. 1

Brooks, W., K. Donovan, and T. R. Johnson (2018): "Mentors or Teachers? Microenterprise Training in Kenya," *American Economic Journal: Applied Economics*, 10, 196–221. A.2

Buera, F. J., J. P. Kaboski, and Y. Shin (2011): "Finance and development: A tale of two sectors," *American economic review*, 101, 1964–2002. 1, 1

Caliendo, L., F. Monte, and E. Rossi-Hansberg (2015): "The anatomy of French production hierarchies," *Journal of Political Economy*, 123, 809–852. 1

Campos, F., M. Frese, M. Goldstein, L. Iacovone, H. C. Johnson, D. McKenzie, and M. Mensmann (2017): "Teaching personal initiative beats traditional training in boosting small business in West Africa," *Science*, 357, 1287–1290. A.2

Carranza, E. and D. McKenzie (2024): "Job Training and Job Search Assistance Policies in Developing Countries," *Journal of Economic Perspectives*, 38, 221–244. A.2

Chandler, A. D. (1990): *Scale and scope: The dynamics of industrial capitalism*, Harvard University Press. 10

Coase, R. H. (1937): "The Nature of the Firm," *Economica*, 4, 386–405. 10

Crépon, B. and P. Premand (2019): "Direct and Indirect Effects of Subsidized Dual Apprenticeships," *SSRN Electronic Journal*. A.2

De Mel, S., D. McKenzie, and C. Woodruff (2008): "Returns to capital in microenterprises: evidence from a field experiment," *The Quarterly Journal of Economics*, 123, 1329–1372. A.2

——— (2019): "Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises," *American Economic Journal: Applied Economics*, 11, 202–235. A.2

Dessein, W. and T. Santos (2006): "Adaptive organizations," *Journal of political Economy*, 114, 956–995. 48

Dupas, P. and J. Robinson (2013): "Savings constraints and microenterprise development: Evidence from a field experiment in Kenya," *American Economic Journal: Applied Economics*, 5, 163–92. 101

Freund, L. (2022): "Superstar Teams: The Micro Origins and Macro Implications of Coworker Complementarities," *Available at SSRN 4312245*. 1

Garicano, L. (2000): "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108, 874–904. 3.2, 52

GARICANO, L. AND E. ROSSI-HANSBERG (2006): “Organization and inequality in a knowledge economy,” *The Quarterly journal of economics*, 121, 1383–1435. 10, 3.2, 52

GHOSH, A. (2022): “Religious divisions and production technology: Experimental evidence from India,” *Available at SSRN 4188354*. 9

GOLDBERG, P. K. AND T. REED (2022): “Demand-side constraints in development: The role of market size, trade, and (in) equality,” Tech. rep., Yale University Working Paper, New Haven, CT. 1

GOLLIN, D. (2002): “Getting Income Shares Right,” *jpe*, 110, 458–474. 1, 1

HARDY, M., S. KIM, J. MCCASLAND, A. MENZEL, AND M. WITTE (2024): “Labor Reallocation between Small Firms: Experimental Evidence on Information Constraints,” *Working Paper*. 23

HARDY, M. AND J. MCCASLAND (2023): “Are Small Firms Labor Constrained? Experimental Evidence from Ghana,” *American Economic Journal: Applied Economics*. 3.3.1, A.2

HJORT, J. (2014): “Ethnic divisions and production in firms,” *The Quarterly Journal of Economics*, 129, 1899–1946. 9

HJORT, J., V. IYER, AND G. DE ROCHAMBEAU (2020): “Informational Barriers to Market Access: Experimental Evidence from Liberian Firms,” Working Paper 27662, National Bureau of Economic Research. 1

HJORT, J., H. MALMBERG, AND T. SCHOELLMAN (2025): “The missing middle managers: Labor costs, firm structure, and development,” Tech. rep. 9

HOLMES, T. J. AND J. J. STEVENS (2014): “An alternative theory of the plant size distribution, with geography and intra-and international trade,” *Journal of Political Economy*, 122, 369–421. 3.3.1, 48

HOLMSTROM, B. AND P. MILGROM (1991): “Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design,” *The Journal of Law, Economics, and Organization*, 7, 24–52. 3.3.1

——— (1994): “The firm as an incentive system,” *The American economic review*, 972–991. 3.3.1

HSIEH, C.-T. AND B. A. OLKEN (2014): “The Missing "Missing Middle",” *The Journal of Economic Perspectives*, 28, 89–108. 1, 2.3

HUSSAM, R., N. RIGOL, AND B. N. ROTH (2022): “Targeting high ability entrepreneurs using community information: Mechanism design in the field,” *American Economic Review*, 112, 861–898. A.2

JENSEN, R. AND N. H. MILLER (2018): “Market Integration, Demand, and the Growth of Firms: Evidence from a Natural Experiment in India,” *American Economic Review*, 108, 3583–3625. 1, 12

JUHÁSZ, R. AND C. STEINWENDER (2018): “Spinning the web: Codifiability, information frictions and trade,” *NBER working paper*. 51

KOHLHEPP, J. (2023): “The Inner Beauty of Firms,” . 1

LAGAKOS, D. (2016): “Explaining cross-country productivity differences in retail trade,” *journal of political economy*, 124, 579–620. 1

LUCAS, R. (1978): “On the Size Distribution of Business Firms,” *bje*, 9, 508–523. 1, 4.4

MACCHI, E. AND J. STALDER (2023): “Work Over Just Cash: Informal Redistribution among Employers and Workers in Kampala, Uganda,” *Working Paper*. 23

MCCASLAND, J., M. HARDY, AND J. ZHANG (2024): “Returns to Capital for Whom? Experimental Evidence from Small Firm Owners and Workers in Ghana,” *Experimental Evidence from Small Firm Owners and Workers in Ghana (July 03, 2024)*. 13

MCKENZIE, D. (2017a): “How Effective Are Active Labor Market Policies in Developing Countries? A Critical Review of Recent Evidence,” *The World Bank Research Observer*, 32, 127–154. A.2

——— (2017b): “Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition,” *American Economic Review*, 107, 2278–2307. A.2

MCKENZIE, D., N. ASSAF, AND A. P. CUSOLITO (2017): “The additionality impact of a matching grant programme for small firms: experimental evidence from Yemen,” *Journal of Development Effectiveness*, 9, 1–14. 101

MCKENZIE, D. AND S. PUERTO (2021): “Growing Markets through Business Training for Female Entrepreneurs: A Market-Level Randomized Experiment in Kenya,” *American Economic Journal: Applied Economics*, 13, 297–332. A.2

McKenzie, D. and C. Woodruff (2008): "Experimental evidence on returns to capital and access to finance in Mexico," *The World Bank Economic Review*, 22, 457–482. 101

——— (2014): "What are we learning from business training and entrepreneurship evaluations around the developing world?" *The World Bank Research Observer*, 29, 48–82. A.2, 101

——— (2017): "Business Practices in Small Firms in Developing Countries," *Management Science*, 63. 2.2

Minni, V. M. L. (2023): "Making the invisible hand visible: Managers and the allocation of workers to jobs," . 9

Moll, B. (2014): "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?" *American Economic Review*, 104, 3186–3221. 1

Piore, M. J. and C. F. Sabel (1984): "The second industrial divide: possibilities for prosperity," . 3.3.1, 48

Porzio, T. (2017): "Cross-Country Differences in the Optimal Allocation of Talent and Technology," . 9

Quinn, S. and C. Woodruff (2019): "Experiments and Entrepreneurship in Developing Countries," *Annual Review of Economics*, 11, 225–248. 1, 1, A.2, 76

Söderbom, M. and F. Teal (2004): "Size and efficiency in African manufacturing firms: evidence from firm-level panel data," *Journal of Development Economics*, 73, 369–394. 3.3.1

Startz, M. (2019): "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade," *mimeo*. 1, 24

Stole, L. A. and J. Zwiebel (1996): "Intra-firm bargaining under non-binding contracts," *The Review of Economic Studies*, 63, 375–410. 4.2

UBOS (2011): *Census of Business Establishments, 2010/11*, Uganda Bureau of Statistics. 15, 2.4

Vickery, S., C. Dröge, and R. Germain (1999): "The relationship between product customization and organizational structure," *Journal of Operations Management*, 17, 377–391. 48

Vitali, A. (2022): "Consumer Search and Firm Location: Theory and Evidence from the Garment Sector in Uganda," *mimeo*. 1, 24

Walker, M. W., N. Shah, E. Miguel, D. Egger, F. Samy Soliman, and T. Graff (2024): "Slack and Economic Development," Working Paper 33055, National Bureau of Economic Research. 1

Williamson, O. E. (1971): "The vertical integration of production: market failure considerations," *The American Economic Review*, 61, 112–123. 10

——— (1973): "Markets and hierarchies: some elementary considerations," *The American Economic Review*, 63, 316–325. 10

——— (1979): "Transaction-cost economics: the governance of contractual relations," *The Journal of Law and Economics*, 22, 233–261. 10

Yang, X. and J. Borland (1991): "A microeconomic mechanism for economic growth," *Journal of political economy*, 99, 460–482. 10

A Online Appendix - Empirical Evidence

A.1 Measuring Labor Specialization: Details

Panel A of Table A.1 lists the 17 tasks elicited in our time use module, together with the share of time spent in each task by the average firm. Panel B shows the production steps for the core product in the three sectors with the share of production time accounted for by each step.

Table A.1: Measuring Time Use

Panel A: All Tasks					
(i) Production	**(58.9%)**	Book-keeping	(0.5%)	Other non-prod. tasks	(0.8%)
Producing Core prod.	(17.6%)	Maintanence	(0.4%)		
Producing other prod.	(41.3%)	Organizing stock	(4.6%)	**(iii) Idle**	**(25.6%)**
		Procuring inputs	(2.0%)	Eating/Resting	(13.5%)
(ii) Non-prod. Tasks	**(15.5%)**	Looking for input supp.	(0.6%)	Waiting for customers	(11.4%)
Interacting with customers	(3.5%)	Looking for new mach.	(0.1%)	Away not for business	(0.7%)
Supervising	(2.2%)	Looking for workers	(0.0%)		
Training	(1.0%)	Managing loans	(0.0%)		
Panel B: Production Steps					
(i) Furniture		**(ii) Metal Products**		**(iii) Grain milling**	
Design	(3.7%)	Design	(7.0%)	Cob shelling	(0.5%)
Drying (before prod.)	(3.0%)	Cutting	(17.9%)	Drying	(1.6%)
Cutting	(13.3%)	Bending	(10.8%)	Cleaning/Destoning	(14.1%)
Planing	(14.0%)	Grinding	(12.9%)	Conditioning	(12.1%)
Thicknessing	(6.8%)	Welding	(28.0%)	De-hulling	(23.5%)
Edging	(10.3%)	Polishing	(11.5%)	Milling	(40.4%)
Sanding	(16.3%)	Painting	(11.9%)	Sealing	(7.8%)
Mortising	(15.4%)				
Finishing	(12.5%)				
Drying (after painting)	(4.8%)				

Notes: The table reports the average share of firm-level time in each task, computed by summing the time spent by the entrepreneur and all employees within a firm on a given task. Panel A uses information from the time use module, asking about time spent hour by hour on the last day worked. For firms not producing the core product, the category "Producing Core prod." corresponds to the production of their main product. Panel B breaks down production time on the core product into time in each step. Steps are listed in typical order of implementation. The statistics in Panel B are conditional on doing a given step for the core product. The data from Panel B comes from the survey module asking the entrepreneur and each employee whether they usually work on each step. This information is available for entrepreneurs and employees but only in firms with at least one employee.

A.2 Firm Characteristics and External Validity: Details

Firm Size Distribution. Figure A.1 plots the size distribution in the three sectors (see the Supplemental Appendix for a version without top coding at 10 workers).

Figure A.1: Firm Size Distribution in the Three Sectors

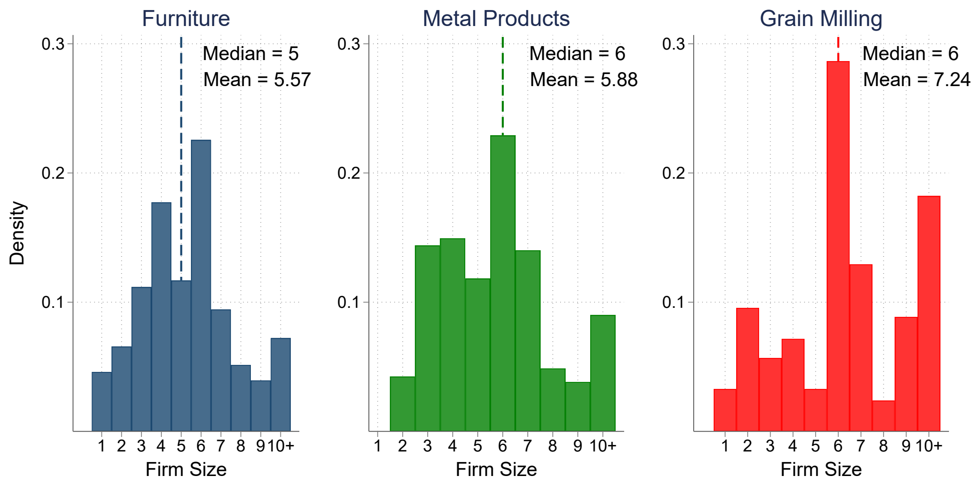

Notes: Sample: all surveyed firms. Firm size is defined as the entrepreneur plus all employees. Firms with more than 10 workers are grouped together in the category "10+". Vertical lines represent the median.

Prevalence of Piece-rate Contracts. The data in Table A.2 provides a breakdown of payment modes reported by workers in the furniture, metal product, and grain milling sectors. Piece-rate-only payments are dominant across all three sectors, accounting for 84.1% in furniture, 81.4% in metal products, and 64.8% in grain milling. Some workers receive a combination of piece-rate payments with additional components, such as in-kind payments, which range from 7.4% to 10.3% across sectors.

Salary-only payments are more prevalent in grain milling (13.6%) compared to furniture (2.3%) and metal products (2.4%), indicating a more formalized payment structure in some grain milling firms. Moreover, the share of unpaid workers is negligible, with 0.7% in both furniture and metal products, and none in grain milling. This highlights that unpaid labor is not a significant feature in these sectors.

Table A.2: Mode of Payment by Sector

Payment Mode	Furniture	Metal Products	Grain Milling
Unpaid	0.007	0.007	0.000
Salary Only	0.023	0.024	0.136
Piece-Rate Only	0.841	0.814	0.648
In-Kind Only	0.009	0.018	0.013
Salary + Piece-Rate	0.015	0.020	0.051
Salary + In-Kind	0.020	0.015	0.046
Piece-Rate + In-Kind	0.074	0.098	0.103
Salary + Piece-Rate + In-Kind	0.005	0.000	0.003
The Worker is Paying the Owner	0.006	0.004	0.000

Notes: Sample: all surveyed firms. This table reports the modes of payment reported by workers interviewed in the furniture (column 1), metal products (column 2) and grain milling (column 3) sectors.

External Validity: Interventions Targeting Firm Growth. To assess the external validity and relevance of our sample for development policy, we compare the average firm size to that in studies evaluating interventions aimed at increasing firm size and productivity in developing countries. To do so, we selected papers that implement Randomized Controlled Trials (RCTs) related to market access, labor market assistance, grants, consulting, and business training, meeting the following criteria: (i) published or forthcoming in top general interest journals; (ii) published in the last twenty years; and (iii) set in a low-income or lower-middle-income country according to the World Bank's classification of country income thresholds.[75] To identify the studies meeting the above criteria, we started from the following review papers: McKenzie and Woodruff (2014), McKenzie (2017a), Quinn and Woodruff (2019), and Carranza and McKenzie (2024), then complemented this with our own literature search in these journals.[76] We identified a total of 16 papers, focusing on market access (Atkin et al., 2017b), labor market assistance (De Mel et al., 2019; Crépon and Premand, 2019; Alfonsi et al., 2020; Bassi and Nansamba, 2022; Abebe et al., 2023; Hardy and McCasland, 2023), capital grants (De Mel et al., 2008; McKenzie, 2017b; Hussam et al., 2022), consulting

[75]The list of journals we considered includes: American Economic Review, American Economic Journal: Applied Economics, Economic Journal, Journal of the European Economic Association, Journal of Political Economy, Management Science, Quarterly Journal of Economics, Review of Economic Studies, Review of Economics and Statistics, and Science.

[76]We exclude microfinance studies and only include capital grants studies because the latter have been shown to have larger returns on average (Quinn and Woodruff, 2019). We also restrict our literature search to papers that implement interventions directly targeting firms, rather than households or individuals. For example, under the "labor market assistance" category we exclude evaluations of vocational training programs only administered to job seekers and without a firm-side randomization.

(Bloom et al., 2013; Anderson and McKenzie, 2022) and business training (Berge et al., 2015; Campos et al., 2017; Brooks et al., 2018; McKenzie and Puerto, 2021; Anderson and McKenzie, 2022).

In Figure A.2, we report the average firm size in each of the 16 studies returned by this literature search. The figure shows that 14 of these 16 studies targeted samples with average firm size less than 10 workers. This shows that we are focusing on the segment of the size distribution that attracts the great majority of efforts by policy-makers and practitioners interested in fostering firm growth.

Figure A.2: Firm Size in Development Interventions

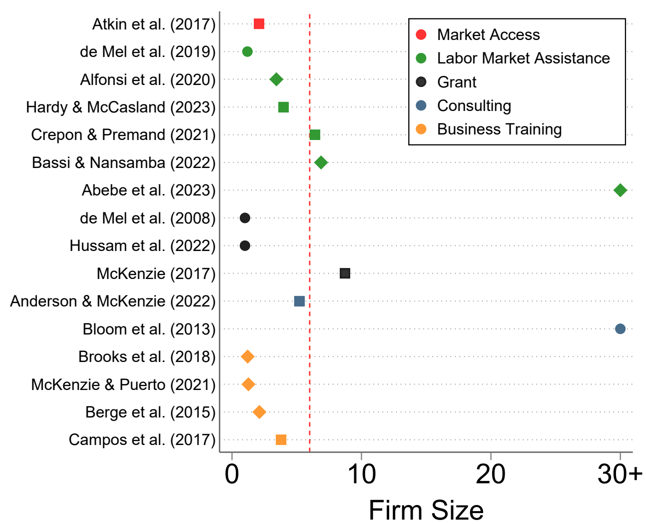

Notes: Figure A.2 illustrates the mean size (x-axis) of firms examined in various studies (y-axis), based on inclusion criteria detailed in the main text. Studies are categorized by type of intervention (using marker colors) and geographical area (using marker symbols). Some studies cover multiple interventions, e.g., Anderson and McKenzie (2022), but are reported only once for brevity. Diamond markers represent studies in East Africa, square markers represent studies in North and West Africa, and circle markers represent studies in South Asia. The red dotted line indicates the median firm size observed in our dataset

Appendix Figure A.3 complements this summary by presenting an overview of the intent-to-treat effects of these interventions on firm size. We report intent-to-treat (ITT) estimates along with their corresponding 95% confidence intervals for all follow-up rounds conducted after the completion of each intervention.[77] In instances where studies implement multiple

[77]In cases where authors report the treatment effects on the total number of employees (rather than firm

(and distinct) experimental arms corresponding to our study topics (e.g., business training and consulting in Anderson and McKenzie (2022)), we report estimates for both sets of interventions. Moreover, some authors report intent-to-treat estimates that either condition upon firm survival and/or do not condition upon firm survival, adding 0s for firms that exit. We report both sets of estimates where available and only one otherwise.

Figure A.3: ITT Effects on Firm Size from the Literature

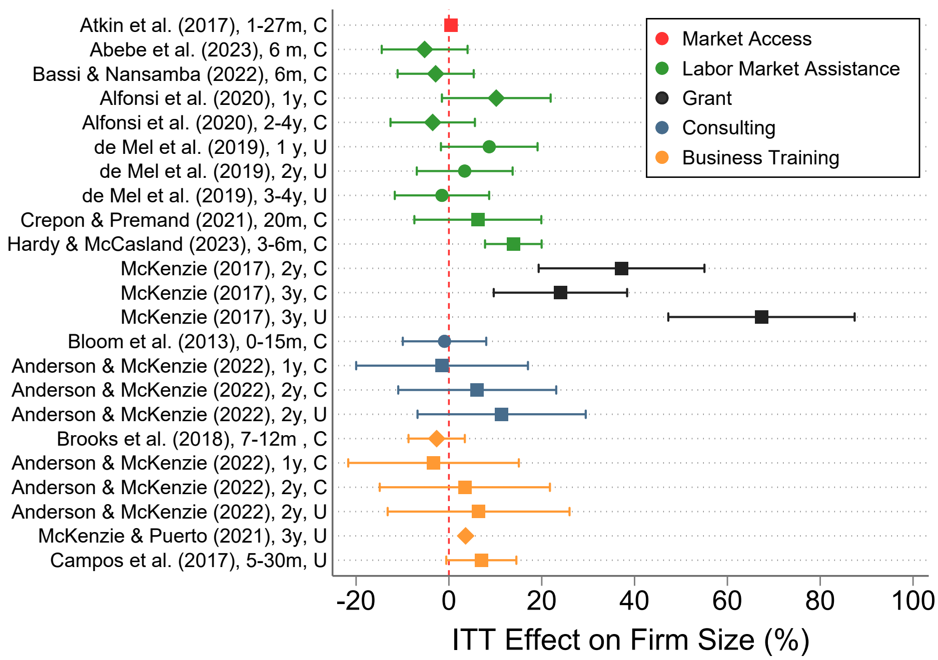

Notes: reports the impact of various interventions on the total number of employees in a firm. We selected papers that implemented Randomized Controlled Trials (RCTs) focused on market access, labor market assistance, grants, consulting, and business training aimed at enhancing firm size and productivity. Our criteria included studies published or forthcoming in top-five or field journals and conducted in low-income or lower-middle-income countries. We report the percentage change in firm size, along with 95% confidence intervals on the x-axis. For y-axis interpretation, refer to the format: "Atkin et al. (2017), 1-27 m, C." This denotes that Atkin et al. (2017b) reports effects pooled over surveys spanning 1 to 27 months post their intervention. C=ITT estimates conditional on firm survival. U=ITT estimates unconditional on firm survival.

Figure A.3 shows that development interventions (e.g. management consulting or business training programs) have typically had small to moderate effects on firm size. As most of these studies focus on firms with 6 or fewer individuals at baseline (Figure 1), this implies that even a 20% increment in firm size represents a modest effect in absolute terms. Table

size), we apply the following rule to recover the percentage change in firm size: average firm size = average no. of employees +1 owner. We further assume that the interventions only impact the number of persons employed, and not the number of firm owners (although hours worked by the owner may change). Two of the sixteen studies selected by our literature search do not report the effect of interventions on employment or firm size, and are omitted from Figure A.3.

A.3 further illustrates the primary components of interventions reported under "consulting" and "business training" categories in the above figure. The table highlights that these interventions frequently address topics such as financial management, employee performance appraisals, and marketing/sales. In contrast, there has been considerably less focus on organizational aspects of the production process, such as task assignment and coordination.

Table A.3: Overview of Consulting & Business Training Interventions

Study	Intervention Type	Main Topics
Anderson & McKenzie (2022)	Business Training	financial management, operations, business plan development, marketing, & HR (e.g., recruiting managers, performance appraisals, compensation & job design).
Anderson & McKenzie (2022)	Consulting	business planning, marketing strategy, goal-setting, financial planning
Berge et al. (2015)	Business Training	entrepreneurship, marketing strategies, customer service, & personnel management (e.g. recruitment, allocation of responsibilities and performance appraisals),
Bloom et al. (2013)	Consulting	factory operations, quality control, inventory, sales and order management, HR (e.g. performance-based incentives and defined job descriptions for personnel)
Brooks et al. (2018)	Business Training	accounting, marketing, cost control, and business planning
Campos et al. (2017)	Business Training	personal initiative, self-starting behavior, innovation, goal-setting, planning, feedback and overcoming obstacles
McKenzie & Puerto (2021)	Business Training	record keeping, financial literacy and management, sales strategies, entrepreneurship, support groups, gender-specific constraints, social norms

Notes: reports the components of management consulting and business training programs shown in Figure A.3.

A.3 Additional Evidence on Task Composition

In this section, we report robustness checks on task composition across the firm size distribution. The main results corresponding to this section can be found in Section 3.1 in the main text.

Figure A.4: Task Composition across the Size Distribution: Robustness

(a) Average Prod. Step

(b) Prod., Non-prod., Idle Time

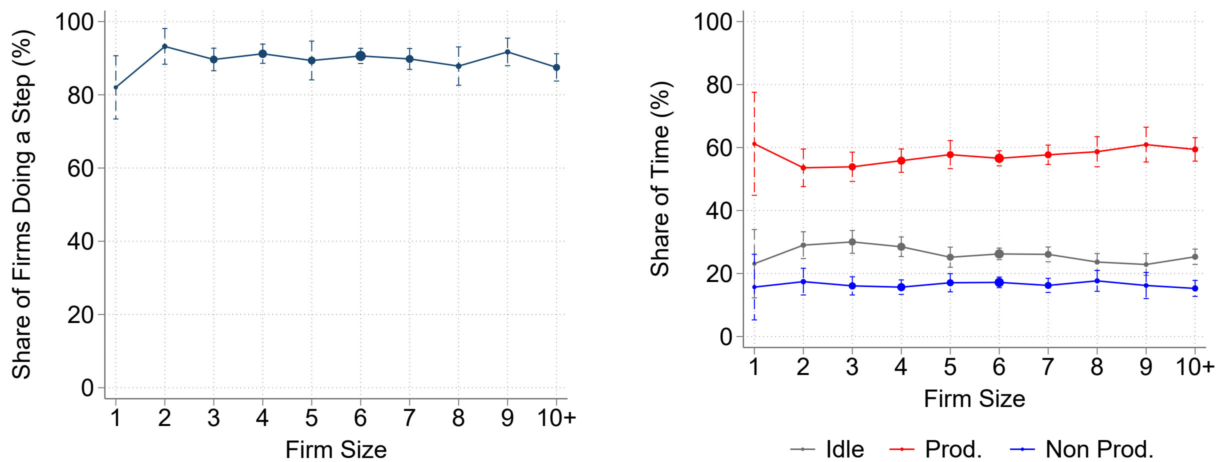

Notes: Sample: furniture and metal products. Panel (a): share of firms doing the representative step, computed as described in the text. Panel (b): share of firm-level time in Production, Non-Production, and Idle tasks. 95% confidence intervals are reported in both panels.

In the main paper, we present results on the share of firms performing the representative production step to assemble the core product(s) in furniture and metal products, shown in Panel (a) of Figure 2. Additionally, Panel (b) displays, for each firm size category, the share of firm time allocated to idle time, production tasks, and non-production tasks. Figure A.4 provides versions of Panels (a) and (b) with 95% confidence intervals. These figures confirm that we cannot reject the hypothesis that task composition is uniform across the firm size distribution.

A.4 Additional Evidence on Horizontal Labor Specialization

In this section, we report additional results and robustness checks on horizontal labor specialization, expanding on the discussion in Section 3.2 of the main text.

Limited Specialization within Production Tasks: Robustness. In the main paper, we present results on the share of employees working on the representative production step required to assemble the core product(s) in furniture, metal products, and grain milling (see Figures 3 and 7(a)). Figure A.5 presents a version of these figures with 95% confidence intervals for robustness. The results confirm that two key patterns discussed in the main text are robust to sampling uncertainty. First, horizontal specialization shows no significant variation between mid-sized and larger firms in the furniture and metal product sectors. Second, there

Figure A.5: Task Allocation Within Production Across the Size Distribution: Robustness

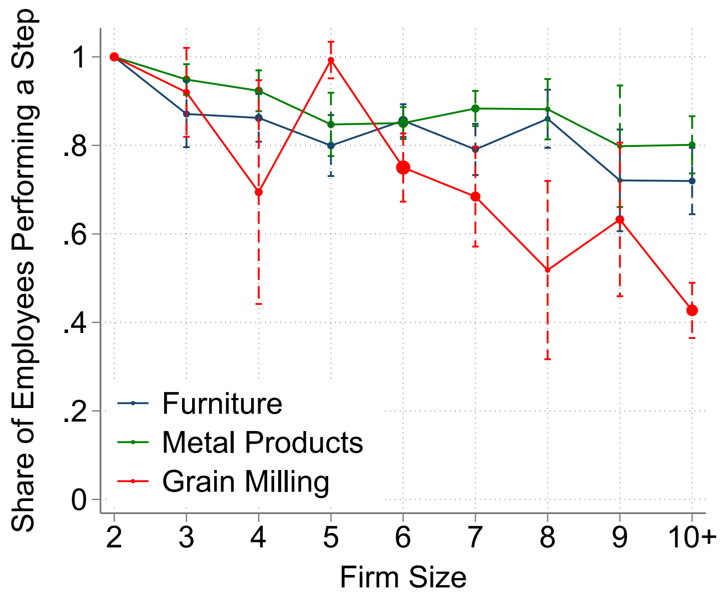

Notes: Sample: all sectors. The figure reports the share of employees performing the representative step in furniture (blue), metal products (green) and grain milling (red) respectively. 95% confidence intervals are reported.

are substantial inter-sectoral differences: firms in the grain milling sector, particularly larger firms with 6 to 10+ employees, demonstrate significantly higher horizontal specialization of workers across production tasks compared to firms in furniture and metal products. Additionally, within the grain milling sector, firms with 10+ employees are significantly more specialized than mid-sized firms with 5–6 employees.

Labor Specialization of Entrepreneurs within Production Tasks. Figure A.6 shows the share of entrepreneurs working on the representative step, as defined in Figure 3. As we have shown that entrepreneurs are less likely to work on production (Figure 5), the share of entrepreneurs working on the typical step is naturally lower than for employees. However, we again find no strong evidence of specialization increasing with firm size: comparing Figures 3 and A.6, we see that the gap between the share of employees and entrepreneurs performing the typical step is roughly constant.

Product-level Evidence on Labor Specialization within Production. We complement the analysis of horizontal specialization from Section 3 by studying: (i) how the number of different employees working on each finished product unit varies with firm size, and (ii)

Figure A.6: Task Allocation within Production: Entrepreneurs

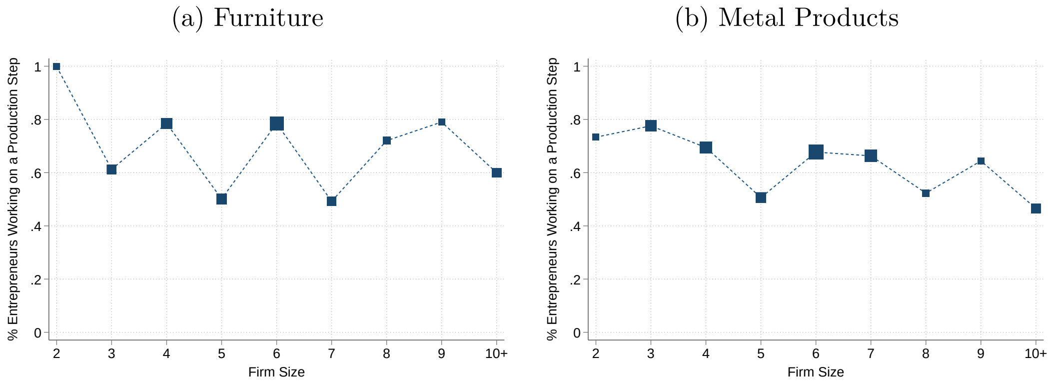

Notes: Replication of Figure 3 for entrepreneurs, focusing on the representative production step (see Figure 3).

how the share of product types worked on by the average employee varies with firm size. To do so, we rely on a new survey of 1,034 firms that we implemented in the same sectors and sub-counties in 2023-23, as part of a new panel data collection designed to study productivity dispersion. The sampling strategy for this new survey was exactly the same as the initial survey described in Section 2: a new listing of firms was conducted in late 2022, and a random sample of firms was extracted for this new survey (again oversampling firms with more than five employees). Reassuringly, basic summary statistics look similar across the two surveys.[78]

This new survey included a similar time-use module to that of the initial survey. In addition, we followed a "product-level" approach, and collected information on the number of different employees doing work on one typical unit of the core products, which are again doors, windows, and maize. Finally, we also recorded all the different product types produced by the firm in the last three months (e.g., in furniture these would be doors, tables, chairs, beds, sofas etc) and added questions on the number of different product types that each employee worked on in the last three months.[79]

Figure A.7, Panel (a), shows the average number of employees working on one unit of the core product, by firm size. The figure shows two key findings: first, in furniture and metal products the number of employees working on each product unit rises slowly hovering at around 2 employees in firms with 4 to 8 employees, and then increases again marginally to about 2.5 employees in firms with 9 or more employees. Second, there is more labor

[78]The average firm size in furniture, metal products and grain milling in the initial (new) survey is 5.57 (4.92), 5.88 (5.35) and 7.24 (9.65) employees, respectively.

[79]As in the initial survey described in Section 2, our new survey targeted all the employees working on the core products.

specialization in grain milling across the entire size distribution, and particularly in larger firms.

Panels (b) and (c) of Figure A.7 study the allocation of labor across product types. Panel (b) shows that carpenters and welders produce about 3-5 different product types, and while larger firms do produce more products, the gradient is small: going from a firm of size 2 to one of size 10 workers is associated with the introduction of about 0.5-1.5 additional products. Grain millers instead tend to produce only 1-2 products, and the relationship with firm size is flat. Panel (c) shows that, in furniture and metal products, the share of product types worked on by each employee decreases only marginally with firm size: even in firms of size 8-10 workers, the average employee still works on about 75-85% of product types. In grain milling the share of product types worked on by each employee is close to 1 across the size distribution, but this is not surprising given that grain millers make only 1-2 products, so that, naturally, any scope for labor specialization across products is limited in this sector.

Taken together, this additional product-level analysis confirms that labor specialization of employees is limited both across tasks within each product unit and across different product types.

A.5 Additional Evidence on Vertical Labor Specialization

In this section, we report additional results and robustness checks on vertical labor specialization, expanding on the discussion in Section 3.2 of the main text.

Limited Specialization in Non-production Tasks Between Employees. Figure A.8 replicates Figure 4 but comparing employees above and below median earnings. The Figure shows much higher overlap and substantially less evidence of specialization in non-production tasks for more skilled employees: when focusing on the headline summary categories of "production" and "managerial tasks", we clearly see that the time allocation of higher and lower skilled employees is more similar than the time allocation of entrepreneurs and the average employee in Figure 4.

Robustness to Focusing on Non-production Sub-categories. In Table A.4 we test the robustness of Figure 5 by focusing only on the share of time spent in the three most complex non-production tasks: supervision/training, customer interactions, and input procurement (see Table A.5 and related discussion on the complexity of different non-production tasks). For comparison, the first row reports the average share of time in all non-production tasks combined (using the same definition as in Figure 5) and the correlation with firm size for both entrepreneurs (columns 1 and 2) and employees (columns 3 and 4). In the second row, we

Figure A.7: Task Allocation within Production: Product-Level Evidence

(a) Employees per Product Unit

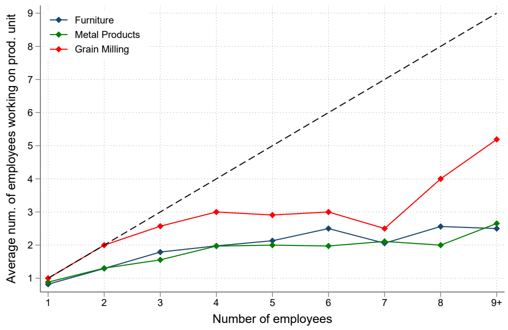

(b) Number of Product Types

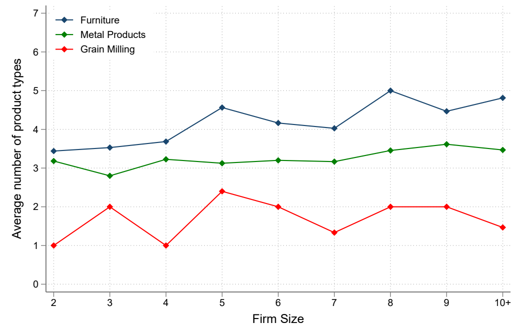

(c) Share of Product Types per Employee

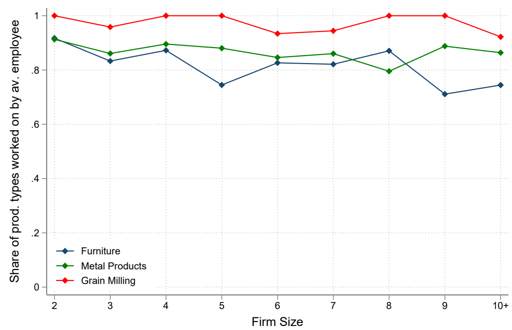

Notes: Panel (a) shows the average number of different employees working on one typical unit of the core product, together with the 45 degree line; Panel (b) shows the number of different product types made by the average firm in the last three months; Panel (c) shows the share of product types that the average employee worked on in the last three months. Metal product firms are in green, furniture in blue, and grain millers in red. Data is from the new survey conducted in 2022-23 and described in the text.

focus instead only on the share of time in any of supervision/training, customer interactions, and input procurement. Finally, in the last three rows, we focus on each of these sub-categories individually. Three key findings emerge. First, comparing columns 1 and 3, we confirm that entrepreneurs are specialized in all complex non-production tasks, regardless of their exact definition. Second, the slope with firm size is always positive for entrepreneurs and close to zero for employees (comparing columns 2 and 4). Third, comparing across rows in column 2, we confirm that the positive relationship with firm size for entrepreneurs is very similar when focusing on all non-production tasks (row 1) or only on the most complex non-production tasks (row 2), and that this relationship is driven primarily by customer interactions and supervision/training (rows 3-5).[80]

[80] In the Supplemental Appendix, we conduct this analysis separately by sector, including also grain milling.

Figure A.8: Task Allocation Between Production and Non-production: Employees

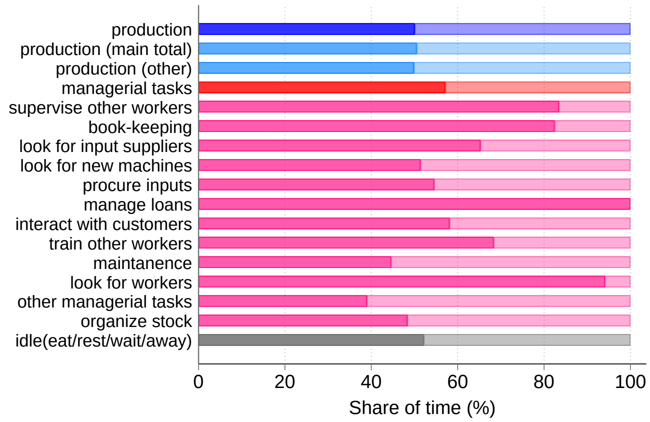

Notes: Replication of Figure 4 but comparing skilled and unskilled employees. Dark bars: skilled employees. Light bars: unskilled employees. Sample: furniture and metal products. The classification between skilled and unskilled employees is based on whether an employee's monthly earnings are above the median among the employees in each firm.

Table A.4: Task Allocation and Firm Size: Robustness

	Entrepreneur		Employees	
	Mean	Slope	Mean	Slope
	(1)	(2)	(3)	(4)
All Non-Production Tasks	0.457	0.021	0.178	0.000
		(0.005)		(0.003)
Cust. Int. + Superv. + Input Proc.	0.322	0.021	0.088	0.004
		(0.005)		(0.002)
Customer Interactions	0.091	0.004	0.036	0.001
		(0.002)		(0.001)
Supervision/Training	0.151	0.014	0.028	0.003
		(0.004)		(0.001)
Input Procurement	0.081	0.003	0.024	-0.000
		(0.002)		(0.001)

Notes: Sample: furniture and metal products. Non-Production Tasks refer to non-production tasks (same definition as in Figure 5). Columns 1 and 3: means. Columns 2 and 4: OLS regression results of the share of time in various non-production categories on firm size (top coded at 10 workers), controlling for region and sector dummies. Robust standard errors clustered at the firm level in parentheses. Superv./Train. represents supervising or training other workers; Cust. Int. represents interacting with customers; Inp. Proc. represents looking for input suppliers, looking for new machines, looking for workers or procuring inputs; Cust. Int. + Superv. + Input Proc. represents all three categories together (Customer Interaction, Supervision/Training, Input Procurement).

Limited Heterogeneity Across Workers Not Varying with Firm Size. We explore whether the findings in panel (a) of Figure 5 that there is little heterogeneity in specialization among employees can mask the fact that some workers may spend significant time in non-production tasks, while others very little. Figure A.9 reports the distribution of time shares in non-production tasks among workers (left panel) and entrepreneurs (right panel). Since our measurement of time use refers to the last day worked, naturally we expect some variation across workers in the share of time spent in non-production tasks. Despite this, the Figure shows that there is very limited heterogeneity among workers, with most workers spending little time in non-production.[81] We farther validate this in Figure A.10, where we split employees by whether their share of time in non-production tasks is above or below the median within each firm-size group. The figure shows that: (i) even workers that spend above median time in non-production activities spend just over 20% of their time in non-production; (ii) the gap between the two types of workers does not increase sharply with firm size. This confirms that heterogeneity across workers is limited, especially considering that our measurement of time use refers to the last day worked, and that there is no organizational change with firm size with respect to employees' time use.

Figure A.9: Distribution of Time Allocation of Employees and Entrepreneurs

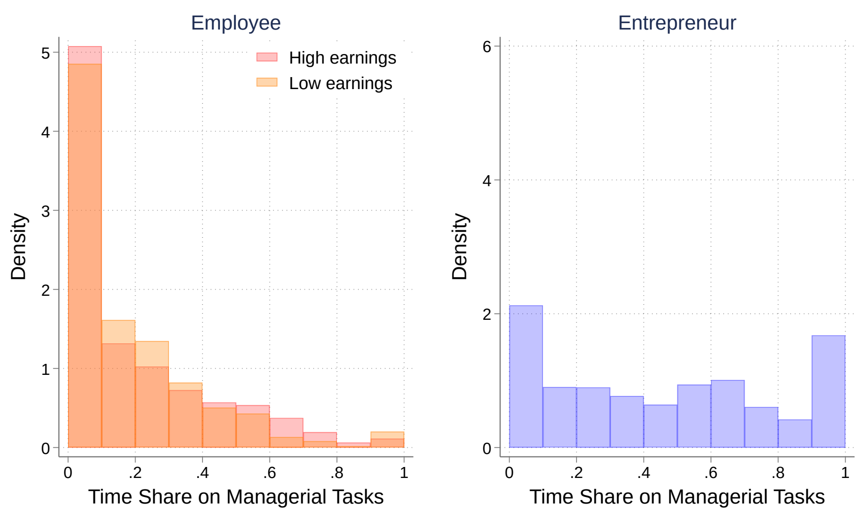

Notes: Distribution of share of time spent on non-production tasks. Left panel: Employees, classified as high and low earnings by above/below median earnings within each firm. Right panel: Entrepreneurs. Sample: furniture and metal products.

Non-Production Tasks Are More Complex. In Appendix Table A.5 we provide evidence to support the claim that non-production tasks are more complex. We do so by studying whether employees more involved in non-production tasks earn more, controlling for firm fixed effects and other worker characteristics. Excluding entrepreneurs from this analysis is important to verify that non-production tasks are indeed more complex tasks,

[81]The right panel shows that there is more variation among entrepreneurs, and this is consistent with panel (b) of Figure 5, where we see that entrepreneurs in larger firms spend more time in non-production tasks.

Figure A.10: Limited Increase in Specialization Across Employees with Firm Size

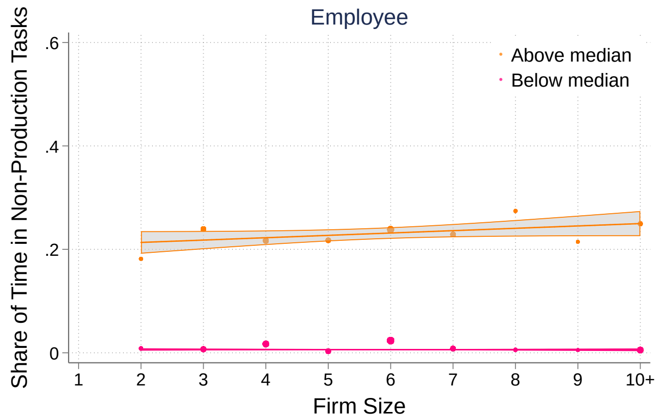

Notes: Share of time employees spend on non-production tasks by firm size. Orange lines correspond to employees that spend an above-median share of time on non-production tasks within each firm-size group; pink lines correspond to those that spend a below-median share. Sample: furniture and metal products.

but not a different kind of task altogether, which may be specific to entrepreneurs. The inclusion of firm fixed effects is critical as it allows us to compare employees within the same firm, thus perfectly controlling for other firm-level determinants of employee earnings or involvement of workers in different types of tasks. In addition, we also control for worker characteristics including age, years of education, tenure at the firm, and whether the worker received vocational training, to narrow in the comparison between workers with similar observables, but who differ in their involvement in non-production activities.

The results in column 1 show that those employees spending a higher share of time on non-production tasks earn substantially more: going from no involvement in non-production to spending all working time in non-production tasks is associated with an increase in earnings of 30%. As this regression controls for firm fixed effects and worker characteristics, this result shows that there are sizeable returns from involvement in non-production tasks, thus suggesting that they are more complex: the higher earnings are consistent with the idea that workers are compensated to be able to complete more challenging tasks that not everyone is able to perform well.[82] In column 2 we then unpack which specific non-production tasks are correlated with higher earnings, by including separate dummies for whether the employee is involved in the different non-production categories. We find that supervision/training, interaction with customers and input procurement drive the earnings gains (and so are

[82]Note also that in Figure 5, Panel (a), we have shown that higher skilled employees (as measured by earnings) spend a larger share of time in non-production tasks, which is again consistent with non-production tasks being more complex.

particularly complex).

Table A.5: Heterogeneity in Skill Intensity of Tasks

	(Log) Employee Earnings		
	(1)	(2)	(3)
Time Share Non-prod. Tasks	0.302		0.270
	(0.107)		(0.114)
Supervise/Train (0/1)		0.213	
		(0.086)	
Customer Int. (0/1)		0.088	
		(0.065)	
Input Procurement (0/1)		0.102	
		(0.047)	
Org. Stock (0/1)		0.005	
		(0.056)	
Other Managerial Tasks (0/1)		-0.047	
		(0.094)	
Avg. Difficulty of Prod. Steps Performed			0.308
			(0.086)
Firm FE	Yes	Yes	Yes
Demographic Ctrl.	Yes	Yes	Yes
Adjusted R^2	0.518	0.525	0.530
Observations	1,976	1,976	1,677

Notes: OLS regression coefficients, standard errors clustered at the firm level in parentheses. Sample: employees in furniture and metal products. Dependent variable: log of monthly employee earnings. Column 1: we include the employee share of time on non-production tasks as a continuous variable). Column 2: we include dummy variables taking value one if the employee performs each task (the reference group are employees who do not perform any non-production tasks). Supervise/Train represents supervising or training other workers; Customer Int. represents interacting with customers; Input Procurement represents looking for input suppliers/new machines/workers or procuring inputs; Org. Stock represents organizing stock; Other Managerial Tasks represents book-keeping, looking for new loans, maintenance or managing loans. Column 3: the variable "Avg. Difficulty of Prod. Steps Performed" is computed as the weighted average of the difficulty levels of the steps performed by each employee (as described in the text), where the weights are the time spent on each step. Demographic controls: age, years of education, tenure at the firm, and a dummy for whether the worker received vocational training.

Variation in Task Difficulty Within Production. In column 3 of Table A.5 we show that even within production there is evidence of vertical differentiation in terms of task difficulty. We exploit a survey question where each employee working on the core product was asked to state their ability to perform each production step (regardless of whether they work on the step), using a 1 to 5 scale. Using this information, we rank steps in each sector by average reported difficulty, and then create a variable that for each employee captures the average difficulty of the steps they perform. We find that employees working on more difficult steps earn more, even controlling for firm fixed effects and other worker characteristics.[83]

Table A.6: Heterogeneity in Skill Distribution within the Firm

	Yrs. schooling		Age		Tenure	
	(1)	(2)	(3)	(4)	(5)	(6)
Entrepreneur (0/1)	0.626		10.658		6.348	
	(0.182)		(0.501)		(0.355)	
Skilled (0/1)		0.739		3.892		1.562
		(0.276)		(0.802)		(0.328)
Sample	Ent.+Emp.	Emp.	Ent.+Emp.	Emp.	Ent.+Emp.	Emp.
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.228	0.318	0.440	0.335	0.491	0.402
Observations	3,237	2,299	3,220	2,281	3,280	2,316

Notes: OLS regression coefficients, standard errors clustered at the firm level in parentheses. Sample: entrepreneurs and employees (odd columns); employees (even columns). Furniture and metal products sectors only. The classification between skilled and unskilled employees is based on whether an employee's salary is above the median among employees in each firm. The variable "Tenure" measures the years of experience of the individual in the firm.

Entrepreneurs Are More Skilled than Employees. In Appendix Table A.6 we show that entrepreneurs are on average more skilled than employees. In columns 1, 3, and 5 we regress years of schooling, age and experience in the firm on a dummy for whether the individual is the entrepreneur or an employee, with firm fixed effects. Entrepreneurs on average have 0.6 more years of education, are 10.7 years older, and have 6.3 more years of experience than employees in their firm, thus confirming that entrepreneurs are significantly more skilled. For comparison, in columns 2, 4 and 6 we limit the sample to employees and create a dummy for whether the employee has above median salary within the firm. We find that more skilled employees within the firm (as proxied by salary) also have more schooling, are older and have longer tenure; however, differences between employees are overall less pronounced, compared to differences between entrepreneurs and employees, apart from education, where the gaps are similar in columns 1 and 2.[84]

Labor Specialization Between Difficult and Simple Steps. We study specialization across production steps of different difficulty by exploiting a survey question where each employee ranked (on a scale 1 to 5) their ability to perform each production step conducted

[83]Consistent with this result, in Appendix Figure A.11 we show that higher-skilled employees (as measured by earnings) spend a larger share of their time on difficult production steps.

[84]Note however that since entrepreneurs are on average more than 10 years older than employees (column 3) there are large cohort effects at play, and controlling for such trends in education would increase farther the gap between entrepreneurs and employees in column 1.

by the firm (regardless of whether the particular employee performs that step). We use this to rank production steps and then we split them by above/below median difficulty. In Figure A.11, we study how employees and entrepreneurs allocate their production time to simple and difficult steps. If an individual only works on difficult steps, the share of time in difficult steps would be 100%. The figure shows that: (i) high skilled employees are more likely to work on difficult steps than low skilled employees, but the gap between the two groups is small and does not vary with firm size; (ii) entrepreneurs spend slightly more time than employees on difficult steps, but again their share of time in difficult steps is close to 50% and there is no gradient with firm size. We conclude that while there is some evidence of entrepreneurs and more highly skilled employees specializing in more difficult steps, this is limited and there is no organizational change with firm size in this dimension.

Figure A.11: Specialization between Simple and Difficult Production Steps

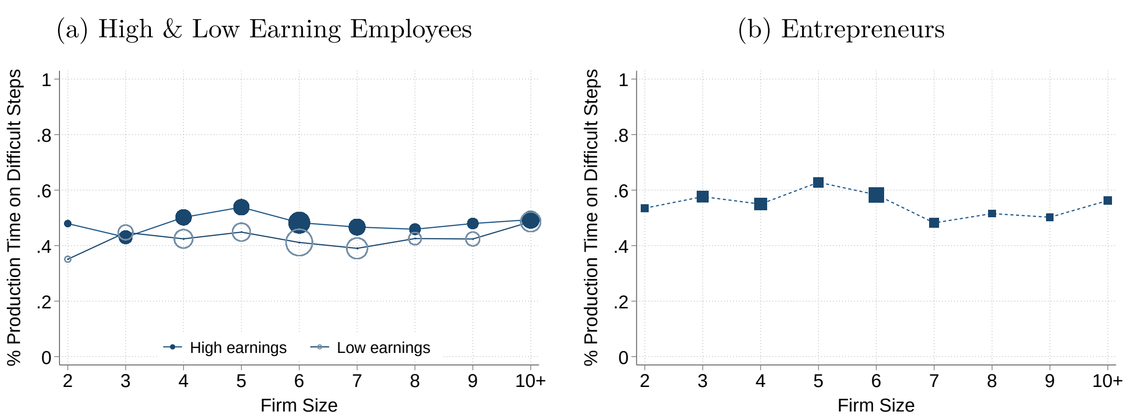

Notes: Sample: furniture and metal products. Panels (a) and (b) represent the share of worker and entrepreneur production time spent on difficult production steps. Definition of difficult steps: see main text.

A.6 Additional Evidence on Heterogeneity in Specialization Within and Across Sectors

In this section, we report additional results and robustness checks on heterogeneity in labor specialization within and across sectors, expanding on the discussion in Section 3.3 of the main text.

Correlates of Specialization Within Sector: Robustness. In the main text, Figure 6 illustrates the correlation between our measures of horizontal and vertical specialization and various firm-level characteristics. Horizontal specialization is defined as the share of employees engaged in the representative production step, while vertical specialization is measured by the difference between the share of time spent by the entrepreneur and the average employee in non-production tasks. These measures are regressed on the relevant firm-level characteristics, controlling for sector and region fixed effects, as well as firm size. In Appendix Figure A.12, we reproduce each of these regressions without controlling for firm size, and find that the results remain very similar to those in the main text figure.

Figure A.12: Correlates of Specialization: Robustness

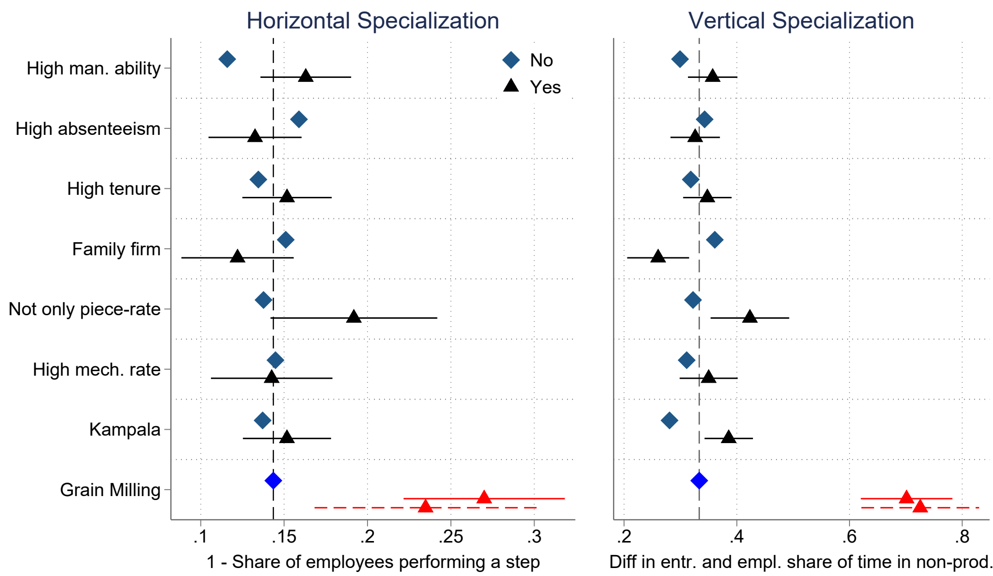

Notes: Sample: rows 1-7: furniture and metal products; rows 8: all sectors. Vertical dotted line: mean horizontal (left panel) and vertical (right panel) specialization in furniture and metal products. Results of OLS estimation of equation 3.1 with one important difference from the main text Figure 6 - we exclude controls for firm size from all regression specifications. The y-axis lists the independent variables of interest in each regression (see main text for definitions). For managerial ability, absenteeism, tenure and mechanization rate we split firms by below/above median. Diamonds: predicted mean in the comparison group (e.g., below median managerial ability firms). To predict this mean, we subtract from the average specialization in furniture and metal products the coefficient on the relevant characteristic of interest (e.g. dummy for above median managerial ability), weighted by the share of observations with that characteristic. Triangles: predicted mean in the comparison group plus estimated coefficient on the characteristic of interest. Bars: 90% confidence intervals.

Evidence on Coordination Costs from Idle Time Data. In Figure A.13 we compare the distribution of idle time across hours of the day, by sector and size. We do so by reporting for each time slot, the share of firms where at least one worker is idle, splitting the sample by below and above median firm size. The figure shows two main results. First, there is significantly more idle time in furniture and metal products. Second, while in grain milling employees in larger firms are significantly less idle (apart from around lunch time), this is not the case in furniture and metal products, where there is effectively no relationship between idleness and firm size.

The results in this figure relate to our main empirical results in two ways. First, firms with higher labor specialization likely exhibit lower idle time as a result of better coordination of work. Therefore, if we take idle time as another proxy of labor specialization, this evidence is again in line with grain milling firms being more specialized, and with the relationship between specialization and firm size being steeper in grain milling. Second, as idle time could reflect the presence of coordination and communication costs, the larger idle time in furniture and metal products is consistent with the prevalence of customization creating sizable communication and coordination costs.[85]

Figure A.13: Heterogeneity in the Share of Idle Workers in a Time Slot

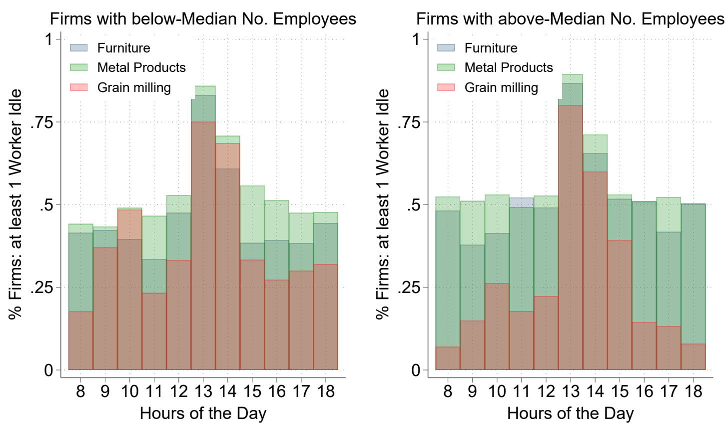

Notes: In both panels, the bars depict the share of firms with at least 1 worker being idle in a time slot. The navy, green, and red colors correspond to the furniture, metal products, and grain milling sectors, respectively. Left panel: sample is firms with below median firm size. Right panel: sample is firms with above median firm size.

Customized Production and Specialization Within Furniture and Metal Products. In Figure A.14, we analyze the correlation between specialization and customization in the furniture and metal products sectors. Our measures of involvement in customized production are taken from Panel B of Table 1 and include: (i) the share of sales made to order in the three months preceding the baseline survey; (ii) a dummy variable indicating whether customization is the primary reason customers buy on order; (iii) a dummy variable for above-median price dispersion for the same product; and (iv) a dummy variable indicating whether customization is the primary driver of price dispersion. We also construct an aggregate index combining these variables, dividing firms into below- and above-median groups based on index values.

In each sub-figure, we regress horizontal (Panel A) and vertical (Panel B) specialization

[85] In the Supplemental Appendix, we show robustness of this Figure A.13 by considering the share of firms where more than 50% of workers are idle at the same time.

measures on the corresponding firm-level customization indicators, controlling for sector and region fixed effects, as well as firm size. The results show that firms with above-median customization levels, as proxied by the overall index, exhibit significantly lower vertical specialization, even within the same sector. However, no distinct patterns are observed for horizontal specialization.

Figure A.14: Correlation between Specialization and Customization

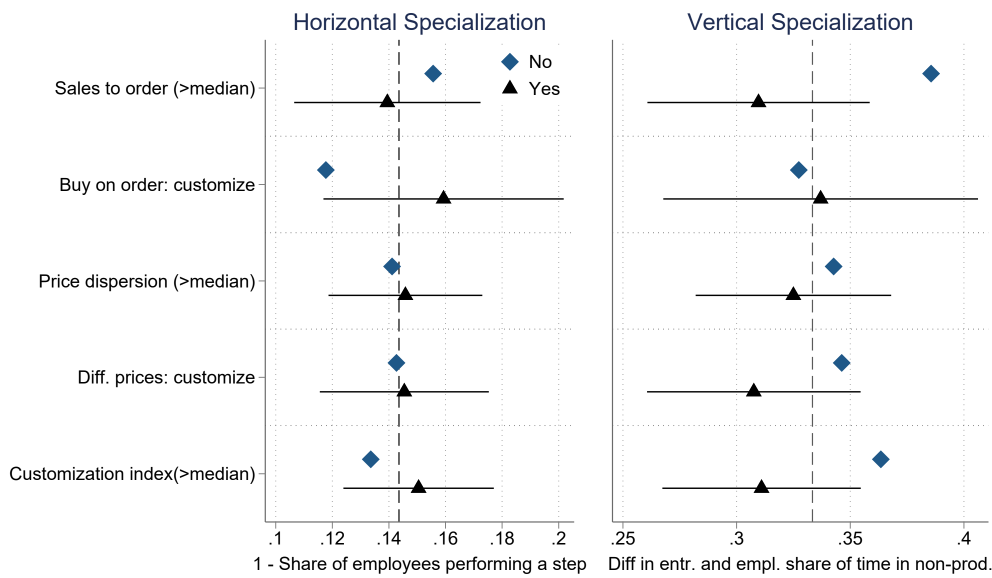

Notes: Sample: furniture and metal products. Results are based on OLS estimation of equation 3.1. The y-axis displays independent measures of customization (see accompanying text for definitions). Diamonds represent the predicted mean for the comparison group (e.g., firms with below median customization). This predicted mean is calculated by subtracting the coefficient on the characteristic of interest (e.g., the dummy for above-median customization), weighted by the share of observations with that characteristic, from the average specialization in furniture and metal products. Triangles represent the predicted mean for the comparison group plus the estimated coefficient on the characteristic of interest. Bars indicate 90% confidence intervals.

B Online Appendix - Model

B.1 Firm Problem - General Setup

We describe firm output in the general case and prove that it simplifies to the setup in the main text under Assumptions (1) and (2). Let $\mathcal{W}$ denote the set of workers, so that $\{\mathcal{W} \cup \hat{z}\}$ is the set of all individuals in the firm.

Task assignment is summarized by $\boldsymbol{\mu}$, which consists of two functions. For all pairs of employees $\{z, z'\} \in \{\mathcal{W} \times \mathcal{W}\}$, $\mu_c(z, z')$ and $\mu_s(z, z')$ specify the fraction of z's complex (c) and simple (s) tasks that are performed by z'. In order to economize on notation, we use the same function to denote the share of complex and simple tasks delegated to and by the entrepreneur $\hat{z}$.[86] For example, $\mu_c(z,\hat{z})$ is the share of an employee z's complex tasks that are performed by the entrepreneur, while $\mu_c(\hat{z},z')$ is the share of the entrepreneur's complex tasks performed by z'. Output is given by:

$$
\begin{aligned}
Y(\hat{z},n,\boldsymbol{\mu}) &= y(\hat{z},\hat{z},\boldsymbol{\mu}) + (n-1)\int y(z,\hat{z},\boldsymbol{\mu})\frac{dF_w(z)}{F_w(z_{\max})} \qquad \text{(B.1)}\\
&\text{where,} \quad \forall z \in \{\mathcal{W}\cup\hat{z}\}\\
y(z,\hat{z},\boldsymbol{\mu}) &= \hat{z}^{\lambda} \quad \tilde{z}(z,\hat{z},\boldsymbol{\mu})^{1-\lambda} \quad \mathbb{I}_{y(z,\hat{z},\boldsymbol{\mu})}\\
\tilde{z}(z,\hat{z},\boldsymbol{\mu}) &= \exp\left\{\mu_C(z,\hat{z})\log(\hat{z}) + (n-1)\int \mu_C(z,z')\log(z')\frac{dF_w(z')}{F_w(z_{\max})}\right\}(1-\kappa(1-\mu_C(z,z)))\\
\mathbb{I}_{y(z,\hat{z},\boldsymbol{\mu})} &= \mathbb{I}_{[\mu_C(z,\hat{z})+(n-1)\int \mu_C(z,z')\frac{dF_w(z')}{F_w(z_{\max})}\geq 1]} \quad \mathbb{I}_{[\mu_S(z,\hat{z})+(n-1)\int \mu_S(z,z')\frac{dF_w(z')}{F_w(z_{\max})}\geq 1]},
\end{aligned}
$$

where the indicator function guarantees that, for each output line $y(.)$, all simple and complex tasks are performed by *someone* in the firm.

In order for an assignment $\boldsymbol{\mu}$ to be *feasible*, no individual in the firm can spend more than their one unit of time across all tasks. Formally, $\boldsymbol{\mu}$ is feasible if and only if

$$
\begin{aligned}
\forall z \in \{\mathcal{W}\cup\hat{z}\}: \quad & \alpha\mu_C(\hat{z},z) + (1-\alpha)\mu_S(\hat{z},z)\\
&+ (n-1)\int (\alpha\mu_C(z',z) + (1-\alpha)\mu_S(z',z))\frac{dF_w(z')}{F_w(z_{\max})} \leq 1 \qquad \text{(B.2)}
\end{aligned}
$$

In its general form, the problem of assigning workers to tasks is highly complex. The entrepreneur needs to choose, for all possible pairs of workers as well as combinations of herself and a worker, what fraction of one individuals' simple and complex tasks are performed by the other, and vice-versa. We now show that, under assumptions (1) and (2), the general problem (B.1) simplifies to choosing only the fraction of complex tasks each worker delegates to the entrepreneur.[87] This necessitates two conditions: (i) the entrepreneur performs all her complex tasks, and (ii) all tasks that are delegated are given to the entrepreneur.

Proposition 2. *Under Assumptions (1) and (2), the task assignment that maximizes (B.1) satisfies (i) $\forall z' : \mu_c(\hat{z},z') = 0$ and (ii) $\forall z' \neq \hat{z} : \mu_c(z,z') = 0$. Hence $\mu_c(z,\hat{z}) = 1-\mu(z,\hat{z})$ where $\mu(z,\hat{z})$ maximizes (4.3).*

[86] In terms of the notation in the main text, $\mu(z,\hat{z}) = 1-\mu_c(z,\hat{z})$.

[87] Problem (4.3) further suppresses $\mu_S(.)$, the share of simple tasks delegated. This is w.l.o.g: since output is independent of who performs the simple tasks and the total number of tasks is fixed, there exist a—possibly large—set of $\mu_s(.)$ for any set of $\mu_c(.)$ such that the assignment is feasible and all tasks are completed.

Proof of Proposition 2.

Assumption (1) guarantees that the feasibility condition is slack, that is, a local increase in $\mu_c(z,\hat{z})$ is feasible. Assumption (2) guarantees that $\forall z \in \mathcal{W} : \hat{z} \geq z$. We prove parts (i) and (ii) by contradiction.

(i) Suppose the optimal assignment has $\mu_c(\hat{z}, z') > 0$ for some z'. Consider an alternative assignment $\boldsymbol{\mu})^*$ with $\mu_c^*(\hat{z}, z') = \mu_c(\hat{z}, z') - \varepsilon$ and $\mu_c^*(\hat{z}, \hat{z}) = \mu_c(\hat{z}, \hat{z}) + \varepsilon$. Note that under $\boldsymbol{\mu}^*$, all complex tasks are performed. Then, $\tilde{z}(z, \hat{z}, \boldsymbol{\mu})^*) > \tilde{z}(z, \hat{z}, \boldsymbol{\mu})$ since both the weight on $\hat{z}$ relative to z' has increased, and the higher $\mu_c^*(\hat{z}, \hat{z})$ reduces the unbundling cost. All other production lines are unaffected. Since $\boldsymbol{\mu}^*$ yields higher firm-level output, $\boldsymbol{\mu}$ could not have been optimal.

(ii) Suppose the optimal assignment has $\mu_c(z, z') > 0$ for some $z' \neq \hat{z}$. Consider an alternative assignment $\boldsymbol{\mu}^*$ with $\mu_c^*(z, z') = \mu_c(z, z') - \varepsilon$ and $\tilde{\mu}_c^*(z, \hat{z}) = \mu_c(z, \hat{z}) + \varepsilon$. Again, all complex tasks are performed under $\boldsymbol{\mu}^*$. Then, $\tilde{z}(z, \hat{z}, \boldsymbol{\mu})^*) > \tilde{z}(z, \hat{z}, \boldsymbol{\mu})$ since the weight on $\hat{z}$ relative to z' has increased and the unbundling cost is unchanged $(\mu_c^*(z, z) = \mu_c(z, z))$. All other production lines are unaffected. Since $\boldsymbol{\mu}^*$ yields higher firm-level output, $\boldsymbol{\mu}$ could not have been optimal.

■

B.2 Proofs

Proof of Lemma 1. We want to show that $\exists z_0 \geq 0$ such that

- $\forall z < z_0$, $\pi(z) < \int w(z, \hat{z}, \boldsymbol{\mu}) \frac{dF_o(\hat{z})}{F_o(z_{\max})}$
- $\forall z \geq z_0$, $\pi(z) \geq \int w(z, \hat{z}, \boldsymbol{\mu}) \frac{dF_o(\hat{z})}{F_o(z_{\max})}$

The proof proceeds in three steps

1. We show that $\frac{\partial \pi(z)}{\partial z} > \frac{\partial}{\partial z} \int w(z, \hat{z}, \boldsymbol{\mu}) \frac{dF_o(\hat{z})}{F_o(z_{\max})} \equiv \frac{\partial \mathbb{E}_{\hat{z}}(w(z,\hat{z},\boldsymbol{\mu}))}{\partial z}$
2. $\pi(0) \leq \int w(0, \hat{z}, \boldsymbol{\mu}) \frac{dF_o(\hat{z})}{F_o(z_{\max})}$
3. $\pi(z_{\max}) \geq \int w(z_{\max}, \hat{z}, \boldsymbol{\mu}) \frac{dF_o(\hat{z})}{F_o(z_{\max})}$

Together, (1)-(3) guarantee the existence of such a threshold.

1. From Equation (4.3) combined with the solution to the wage bargaining in Equation (4.4), we can write the derivative of profits wrt to owner ability as

$$\frac{\partial \pi(x)}{\partial x} = \frac{\partial y(\hat{z}, x, \boldsymbol{\mu})}{\partial x} + (n-1)(1-\omega) \int \frac{\partial y(z, x, \boldsymbol{\mu})}{\partial x} \frac{dF_w(z)}{F_w(z_{\max})} \tag{B.3}$$

   Here, we used the fact that n and $\boldsymbol{\mu}$ are optimal choices and hence the envelope theorem applies. Since the owner performs all her own complex tasks, $\frac{\partial y(\hat{z},x,\boldsymbol{\mu})}{\partial x} = 1$ and therefore $\frac{\partial \pi(x)}{\partial x} \geq 1$.

   Turning to expected wages,

$$\frac{\partial \mathbb{E}_{\hat{z}}(w(z, \hat{z}, \boldsymbol{\mu})}{\partial z} = \omega \int \frac{\partial y(z, \hat{z}, \boldsymbol{\mu})}{\partial z} \frac{dF_o(\hat{z})}{F_o(z_{\max})} \leq 1 \tag{B.4}$$

   by Assumption 2.

2. Suppose instead that $\pi(0) > \int w(0, \hat{z}, \boldsymbol{\mu}) \frac{dF_o(\hat{z})}{F_o(z_{\max})}$. From above, the derivative of profits is always larger than the derivative of expected wages. Then, the set of workers would be the empty set and the labor market would not clear.

3. Suppose instead that $\pi(z_{\max}) < \int w(z_{\max}, \hat{z}, \boldsymbol{\mu}) \frac{dF_o(\hat{z})}{F_o(z_{\max})}$. From above, the derivative of profits is always larger than the derivative of expected wages. Then the set of entrepreneurs would be the empty set and the labor market would not clear.

■

Proof of Lemma **2.**

1. The share of time each worker spends on complex tasks is equal to α – the total amount of complex tasks in his production line – minus the share of tasks delegated to the entrepreneur. Using Equation (4.6), this can easily be rewritten as $\theta(z, \hat{z}) = \alpha \left(1 - \frac{1}{\kappa_0} (\log \hat{z} - \log z)^{\kappa_1}\right)$. The share of time the entrepreneur spends on complex tasks is equal to D – the time it takes her to complete her own complex tasks – plus the time to complete all her $n-1$ workers' complex tasks that were delegated to her. $\hat{\theta}(\hat{z}) = \alpha \left(1 + \frac{n-1}{\kappa_0} \int (\log \hat{z} - \log z)^{\kappa_1} \frac{dF(z)}{F_w(z_{max})}\right)$

■

Proof of Lemma **3.** Rearranging Equation (4.1) gives the result.

■

PROOF OF LEMMA 4.

Let $\overline{w}(\hat{z},\boldsymbol{\mu}) \equiv \int w(z,\hat{z},\boldsymbol{\mu})\frac{dF_w(z)}{F_w(z_{\max})}$. The equation in Lemma (4) follows from taking the first-order condition of (4.3) with respect to n. Further,

$$\frac{\partial \mathbb{Z}(\hat{z},n,\boldsymbol{\mu})}{\partial n} = \hat{z}^{\lambda}\frac{1}{n^2}\left[-\hat{z}^{1-\lambda} + \int \tilde{z}(z,\hat{z},\boldsymbol{\mu})^{1-\lambda}\frac{dF_w(z)}{F_w(z_{\max})}\right] \leq 0$$

where the last inequality follows from the definition of $\tilde{z}(z,\hat{z},\boldsymbol{\mu})$.

Solving for n,

$$n = \frac{1}{\chi_0}\left[\hat{z}^{\lambda} + \int \tilde{z}(z,\hat{z},\boldsymbol{\mu})^{1-\lambda}\frac{dF_w(z)}{F_w(z_{\max})}\right]^{\chi_1}$$

which is declining in χ_0.

Using the envelope theorem,

$$\frac{\partial \tilde{z}(z,\hat{z},\boldsymbol{\mu})}{\partial \kappa_0} = z^{\mu(z,\hat{z})}\hat{z}^{1-\mu(z,\hat{z})}\frac{\partial(1-\kappa(\mu(z,\hat{z})))}{\partial \kappa_0} \leq 0$$

and hence $\frac{\partial n}{\partial \kappa_0} < 0$ as long as $\lambda < 1$.

■

PROOF OF LEMMA 5.

1. When $\lambda = 1$, $Y(\hat{z},n,\boldsymbol{\mu})$ directly collapses to $\hat{z}\,n$. When $\kappa_0 = 0$, then $\mu(z,\hat{z}) = 1\ \forall z$. Note that Assumption (1) guarantees that the entrepreneur has capacity to take on all complex tasks of her workers. With $\mu(z,\hat{z}) = 1$, we have again that $Y(\hat{z},n,\boldsymbol{\mu}) = \hat{z}\,n$. Optimal firm size is then simply given by

$$\hat{z} = \overline{w}(\hat{z},\boldsymbol{\mu}) + \chi'(n)$$

and is increasing in $\hat{z}$ since $\chi'(n) = (\chi_0 n)^{\frac{1}{\chi_1}}$ is increasing in n

2. When $\kappa_0 \to \infty$, no tasks are unbundled and $\mu(z,\hat{z}) = 0\ \forall z$. Hence $\tilde{z}(z,\hat{z},\boldsymbol{\mu}) = z$. Moreover, if $\lambda = 0$, we get that $Y(\hat{z},n,\boldsymbol{\mu}) = \hat{z} + (n-1)\int z\frac{dF_w(z)}{F_w(w_{max})}$ and optimal firm size solves

$$\chi'(n) + (1-\omega)\overline{w} + \omega\int z\frac{dF_w(z)}{F_w(w_{max})} = \int z\frac{dF_w(z)}{F_w(w_{max})}$$

which is independent of $\hat{z}$.

■

Proof of Proposition 1.

Consider an increase in $1/\kappa_0$ (decrease in κ_0). With $\kappa_1 = 0$, firm-level output simplifies to

$$Y(\hat{z}, n) = \hat{z} + (n-1)\,\hat{z}^{\lambda + \frac{1-\lambda}{\kappa_0}} \int_0^{z_0} z^{(1-\lambda)(1-\frac{1}{\kappa_0})} \frac{dF(z)}{F(z_0)} \tag{B.5}$$

To simplify notation, let average output per worker in a firm owned by an individual with ability $\hat{z}$, when the marginal entrepreneur in the economy is given by z_0, be denoted $\mathbb{Y}(\hat{z}, z_0)$. That is,

$$\mathbb{Y}(\hat{z}, z_0) \equiv \hat{z}^{\lambda + \frac{1-\lambda}{\kappa_0}} \int_0^{z_0} z^{(1-\lambda)(1-\frac{1}{\kappa_0})} \frac{dF(z)}{F(z_0)} \tag{B.6}$$

Profits of an entrepreneur with ability $\hat{z}$ can then be written as:

$$\pi(\hat{z}; z_0) = \hat{z} + (n-1)(1-\omega)\left[\mathbb{Y}(\hat{z}, z_0) - \overline{w}\right] - \chi(n) \tag{B.7}$$

where n is equal to

$$n = \frac{1}{\chi_0}\left[(1-\omega)\left(\mathbb{Y}(\hat{z}, z_0) - \overline{w}\right)\right]^{\chi_1} \tag{B.8}$$

The expected wage of a worker z is equal to

$$\mathbb{E}\left(w(z; \boldsymbol{\mu})\right) = (1-\omega)\bar{w} + \omega\mathbb{Y}_w(z; z_0) \tag{B.9}$$

where $\mathbb{Y}_w(z; z_0)$ is, analogously to $\mathbb{Y}(\hat{z}, z_0)$, the average output that a worker of ability z would get given the equilibrium distribution of entrepreneurs in the economy:

$$\mathbb{Y}_w(z_0) \equiv z^{(1-\lambda)(1-\frac{1}{\kappa_0})} \int_{z_0}^{z_{\max}} \hat{z}^{\lambda + \frac{1-\lambda}{\kappa_0}} \frac{n(z)dF(z)}{\int_{z_0}^{z_{\max}} n(z)dF(z)} \tag{B.10}$$

The two equations that pin down the aggregate equilibrium objects—z_0 and $\overline{w}$—are given by

$$z_0 + (n-1)(1-\omega)\left[\square(\hat{z}, z_0) - \overline{w}\right] - \chi(n) = (1-\omega)\bar{w} + \omega \square_w(z_0), \tag{B.11}$$

$$\int_{z_0}^{\overline{z}} n(z) f(z) dz = 1. \tag{B.12}$$

The structure of the proof then is as follows: We find the level of the wage $\overline{w}^*$ such that, given a marginal increase from $1/\kappa_0$ to $1/\kappa_0^*$, the marginal entrepreneur z_0 is unchanged. For small enough ω, as we assumed, $\overline{w}^* > \overline{w}$.

We then show that at this wage level, aggregate labor demand exceeds aggregate supply. Thus, the new equilibrium wage level must be bigger than $\overline{w}^*$, implying that z_0 is higher in the new equilibrium as well. The last part of the argument follows from our assumption on the slope of aggregate labor demand wrt the wage.

Let $n^*(z_0)$ be the level of employment of the cut-off type z_0 under κ_0^* and $\overline{w}^*$.

$$z_0 + (n-1)(1-\omega)\left[\square(z_0, z_0, \kappa_0^*) - \overline{w}^*\right] - \chi(n) = (1-\omega)\bar{w}^* + \omega \square_w(z_0, \kappa_0^*) \tag{B.13}$$

$$n^*(z_0) = \frac{1}{\chi_0}\left[(1-\omega)\left(\square(z_0, z_0, \kappa_0^*) - \overline{w}^*\right)\right]^{\chi_1} \tag{B.14}$$

Combining the two equations:

$$n^* = \frac{1}{\chi_0}\left[\frac{(1-\omega)\bar{w}^* + \omega \square_w(z_0, \kappa_0^*) - z_0 + \chi(n^*)}{n^*-1}\right]^{\chi_1} \tag{B.15}$$

We want to show that $\frac{\partial n^*}{\partial w^*} > 0$. Totally differentiating, we get

$$\begin{aligned}\chi_0 dn = \chi_1 &\left[\frac{(1-\omega)\bar{w}^* + \omega \square_w(z_0, \kappa_0^*) - z_0 + \chi(n^*)}{n^*-1}\right]^{(\chi_1 - 1)} \\ \times &\left[\frac{(1-\omega)d\bar{w}^* + \omega \frac{\partial \square_w(z_0, \kappa_0^*)}{\partial 1/\kappa_0} d(1/\kappa_0) + \chi'(n)dn}{n-1} - dn\frac{(1-\omega)\bar{w}^* + \omega \square_w(z_0, \kappa_0^*) - z_0 + \chi(n^*)}{(n-1)^2}\right]\end{aligned} \tag{B.16}$$

$$\chi_0 dn = \chi_1 \left(\chi_0 n\right)^{\frac{\chi_1 - 1}{\chi_1}} \frac{(1-\omega)d\bar{w}^* + \omega \frac{\partial \square_w(z_0, \kappa_0^*)}{\partial 1/\kappa_0} d(1/\kappa_0)}{n-1} \tag{B.17}$$

For small enough ω, implies that $\frac{\partial n^*}{\partial w^*} > 0$, that is, the cut-off entrepreneur z_0 chooses to run a larger firm under $\overline{w}^*$ and κ_0^*. Note that if firm size increases for the cut-off entrepreneur, it also increases for all entrepreneurs with higher ability. Therefore, the labor market cannot clear.

In equilibrium therefore, we must have that $\overline{w}$ increase to a *higher* level than $\overline{w}^*$. Together with the fact that aggregate labor demand declines in the wage level, it must be that z_0 increases.

1. With $\kappa_1 = 0$, the share of time the entrepreneur spends on complex tasks simplifies to $\hat{\theta}(\hat{z}, n) = \alpha\left(1 + \frac{n}{\kappa_0}\right)$, which is increasing in $1/\kappa_0$.

2. $\frac{\partial \hat{\theta}(\hat{z},n)}{\partial n} = \frac{\alpha}{\kappa_0}$ which is increasing in $1/\kappa_0$.

3. Shown above.

4. Implied by the fact that z_0 increases and the labor market clears.

5. The output of the production line associated to each individual either stays constant (for entrepreneurs who stay entrepreneurs under κ_0^*) or increases. To see this, recall that

$$y(z, \hat{z}, \boldsymbol{\mu}) = \hat{z}^{\lambda}\left(z^{\mu(z,\hat{z})}\hat{z}^{1-\mu(z,\hat{z})}[1 - \mu(z, \hat{z})]\right) \tag{B.18}$$

   and consider that all individuals are matched – on average – with more skilled entrepreneurs, and also acquire more of their higher productivity due to the stronger specialization ($\mu(z, \hat{z})$ is lower). As a result, total output increases, implying that average firm productivity must increase as well since the total amount of labor is constant.

6. The wage is given by $w(z, \hat{z}, \boldsymbol{\mu}) = (1-\omega)\overline{w} + \omega z^{(1-\lambda)(1-\frac{1}{\kappa_0})}\hat{z}^{\lambda+\frac{1-\lambda}{\kappa_0}}$ which increases for all $\{z, \hat{z}\}$ since $\overline{w}$ increased and the increase in $1/\kappa_0$ increases the wage as long as $\hat{z} > z$. Further, the set of entrepreneurs becomes more productive, so in the new equilibrium, the $\hat{z}$ any worker matched with is at least as high.

■

B.3 Micro-founding the unbundling cost

The value of output produced by individual z, working for entrepreneur $\hat{z}$ with task assignment $\boldsymbol{\mu}$ is given by

$$y(z, \hat{z}, \boldsymbol{\mu}) = \hat{z}^{\lambda} \left[\hat{z}^{\mu(z,\hat{z})} z^{1-\mu(z,\hat{z})} (1 - \kappa(\mu(z, \hat{z})))\right]^{1-\lambda} \tag{B.19}$$

We first provide one possible micro-foundation of the unbundling cost $\kappa(.)$ as arising from consumers' preferences over goods of different qualities. Then we discuss how *standardized* production could arise as the optimal technology in a large enough market.

B.3.1 Quality losses of delegation

Households have linear preferences over a unit continuum of goods, indexed by quality $\xi \in [0, 1]$. For each good, they choose how much to purchase $q(\xi) \geq 0$.

$$U = \max_{\{q(\xi)\}_{\xi=0}^{1}} \int_0^1 \xi q(\xi) d \tag{B.20}$$

$$\text{s.t.} \int_0^1 p(\xi) q(\xi) = 1 \tag{B.21}$$

The higher a good's quality, the more utility households get from consuming it.

Producing high quality goods is costly. In particular, whenever a firm standardizes its production process, this comes at the expense of quality—either literally, by not using the talent of workers, or by not allowing the final output to exactly correspond to the customer's preferences. When a fraction μ of tasks is delegated inside the firm, the corresponding quality of the output is

$$\xi = (1 - \kappa(\mu))^{1-\lambda}$$

From the consumer's maximization problem, it follows that any good purchased in positive quantities must have its price equal to quality,

$$p(\xi) = \xi = (1 - \kappa(\mu))^{1-\lambda}. \tag{B.22}$$

Taken together, an individual z, working for entrepreneur $\hat{z}$ with task assignment $\boldsymbol{\mu}$ produces a *physical* quantity equal to

$$q = \hat{z}^{\lambda} \left[\hat{z}^{\mu(z,\hat{z})} z^{1-\mu(z,\hat{z})}\right]^{1-\lambda} \tag{B.23}$$

The resulting quality—and hence price—of the output is $\xi = p = (1 - \kappa(\mu(z, \hat{z})))^{1-\lambda}$. Hence the total revenues, pq are given by

$$\hat{z}^{\lambda}\left[\hat{z}^{\mu(z,\hat{z})}z^{1-\mu(z,\hat{z})}(1-\kappa(\mu(z,\hat{z})))\right]^{1-\lambda} \tag{B.24}$$

B.3.2 Choice of Technology

The simple model above can also help formalize the difference between the type of firms prevalent in our setting, which have limited specialization, and a modern firm with standardized goods and a production line. Suppose firms have a choice between two technologies: the one described above, where perceived quality (and hence price) of each unit of output depend inversely on the amount of delegation, and a "standardized technology", which we describe below.

If firms adopt the standardized technology, output of each production line is independent of the identity of the person producing, as all workers follow the same blueprint with physical productivity $\hat{z}$. Since goods are now standardized rather than produced to fit each consumer's specific tastes and needs, their quality is, on average, lower:

$$\xi = 1-\bar{\kappa} \tag{B.25}$$

Standardized technology therefore corresponds to $\lambda = 1$, but with a lower quality and hence lower value to consumers.

Assuming that adopting the standardized production comes at a fixed cost F, an entrepreneur would do so if and only if

$$\max_{n}\left[\hat{z}n-(n-1)\int w(z,\hat{z},\boldsymbol{\mu})\frac{dF_w(z)}{F_w(z_{\max})}-\chi(n)-F\right] > \tag{B.26}$$

$$\max_{n,\mu}\left[\mathbb{Z}(\hat{z},n,\boldsymbol{\mu})n-(n-1)\int w(z,\hat{z},\boldsymbol{\mu})\frac{dF_w(z)}{F_w(z_{\max})}-\chi(n)\right] \tag{B.27}$$

Next, define $n^S(\hat{z})$ and $n^C(\hat{z})$ the optimal firm sizes for the standardized and the customized problems – hence the solutions to the two problems above.

Since $\hat{z} > \mathbb{Z}(\hat{z},n,\boldsymbol{\mu})$, we get that $n^S(\hat{z}) > n^C(\hat{z})$ and entrepreneurs would thus adopt the standardized technology only if their ability and scale of operations covers the fixed costs F.

Why don't we see firms in Uganda adopting the standardized technology? Arguably, given the small market size (which would show up as a high "hiring cost" $\chi(n)$ in the equation above) and the lack of very talented entrepreneurs (low $\hat{z}$) no one has the sufficient scale of operations to justify the adoption of this organizational form.

C Online Appendix - Estimation

C.1 Empirical Validation of the Theoretical Predictions

We provide two qualitative tests to support the model predictions of Section 4.4.

Heterogeneity across Sectors. Proposition 1 is in principle testable using market-level variation in the unbundling cost κ_0. In the absence of credible exogenous variation, we rely on cross-sectoral heterogeneity. As discussed in Section 3, the degree of standardization is remarkably similar in furniture and metal products, but is larger in grain milling, suggesting a lower κ_0 in that sector.

In Table C.1, we show that the key predictions of Proposition 1 hold across sectors. Furniture and metal products are almost identical in terms of labor specialization, average size, returns to managerial ability, and selection into entrepreneurship. In grain milling, on the other hand, there is more specialization, firms are larger, and the returns from managerial ability as well as the skill gap between entrepreneurs and their employees are larger.[88]

Heterogeneity across Regions. Our model has one unique implication, shown in Lemma 3: all else equal, entrepreneurial ability is less important for firm productivity in larger firms since employees are responsible for a larger share of the "firm management". To test this prediction, we would ideally find a credible instrument for firm size. In the absence of such exogenous variation, we provide suggestive evidence exploiting heterogeneity across sub-counties. We proceed as follows. First, we drop all firms in grain milling.[89] We calculate the average firm size in each sub-county, rank them based on this statistic, and then divide them in two groups with roughly equal numbers of firms.[90] Finally, within each group of sub-counties, we estimate the return to managerial ability by regressing log revenues on sector dummies and either the managerial ability index or the years of education of the entrepreneur.

The results are shown in Table C.2. Consistent with Lemma 3, we find higher returns to managerial ability within the set of sub-counties with the smallest firms.

[88] We do not test the prediction on wages because the model-consistent wage level is not directly observable in the data. A simple comparison of average wage would not hold since employees in grain milling are (in both the model and the data) less skilled.

[89] We restrict our focus to furniture and metal products since we have shown in Table C.1 that grain milling has larger returns to managerial ability, and we want to ensure that sectoral composition across regions is not driving our estimates. The results are unaffected by the restriction, however.

[90] The "marginal" sub-county is one of the largest ones, implying that we end up with 40% of the firms in one group and 60% in the other.

Table C.1: Cross-Sectoral Heterogeneity

	Furniture (1)	Metal Products (2)	Grain Milling (3)
Panel A. Average Specialization & Firm Size			
Specialization	0.32	0.35	0.62
Firm Size	5.6	5.9	7.2
Panel B. Reg. Coeff's on Man. Ability (Std.)			
Log Revenues	0.24	0.25	0.57
Log Revenues per Worker	0.14	0.15	0.35
Log Size	0.10	0.10	0.23
Panel C. Reg. Coeff's on Entrepreneur (0/1)			
Years of Education	0.87	-0.10	3.29
Age	10.4	11.5	19.0
Log Earnings	0.72	0.94	1.00

Notes: Panel A: Sample: all firms. Average specialization: gap in the share of time in non-production tasks between the entrepreneur and her employees. Panel B: Sample: all firms. Coefficients from the regression of three dependent variables on the (standardized) index of managerial ability. Panel C: Sample: all entrepreneurs and employees. Regressions on a dummy equal to 1 if the individual is the entrepreneur, and zero if they are an employee. Regressions for Panels B and C include region fixed effects. Earnings are labor income for workers and firm profits for entrepreneurs.

Table C.2: Returns to Managerial Ability in Locations with Different Firm Size

	Dep. Var: (Log) Revenues			
	(1)	(2)	(3)	(4)
Manager Ability (Std.)	0.388	0.177		
	(0.056)	(0.041)		
Yrs. of Education			0.060	0.036
			(0.016)	(0.012)
Subcounty by Firm Size	Small	Large	Small	Large
(Average Firm Size)	(4.80)	(6.15)	(4.80)	(6.15)
Sector FE	Yes	Yes	Yes	Yes
Adjusted R^2	0.152	0.044	0.081	0.029
Observations	360	583	360	583

Notes: OLS regression coefficients. Sample: furniture and metal products. Robust standard errors are in parentheses.

C.2 Details on Empirical Moments and Calibration

We describe the computation of the 150 moments targeted in the model estimation, and of the calibrated fixed costs. We use pooled data from furniture and metal products. All moments are computed from the initial survey. The calibrated fixed cost, instead, is from the follow-up survey. We start by describing the moments, and organize the discussion by dividing them into four groups, following the four panels of Table 4.

Allocation of Time to Complex Tasks (Table 4, Panel A). The Average Time on Complex Tasks (Panel A, row (i)) is the average firm-level share of time in non-production tasks, including the entrepreneur and all employees. Rows (ii) and (iii) report, respectively, the average share of time in non-production tasks for all entrepreneurs, and for entrepreneurs in firms of size 1 (so with no employees). The statistic in row (ii) is computed exactly as in Table A.4, column 1. Rows (iv) and (v) report the average share of time in non-production tasks for employees, split by below and above median salary (we use salary as a proxy of skill).[91] The slope for entrepreneurs in row (vi) is taken from column 2 of Table A.4, where we regress the share of time of the entrepreneur in non-production tasks on firm size. Rows (vii) and (viii) report the coefficients from a similar regression for high- and low-skilled employees separately, where again we split them by below and above median salary. The results are reported in columns 2 and 3 of Table C.3. Finally, row (ix) reports the coefficient from a regression of the share of time in non-production tasks on log employee earnings, with firm fixed effects, shown in column 1 of Table C.3.

We also target several moments related to the distribution of specialization in complex tasks across the size distribution (shown in Figure 9). Specifically, we calculate the share of entrepreneurs' time in non-production tasks in each firm size group, from firms with no employees to firms with 10 or more employees (10 moments), and, similarly, the share of time in non-production tasks for employees in each firm size group (other 9 moments). Finally, we also target the share of time in non-production tasks for employees with below median earnings in each firm size group (again splitting employees by below and above median within each firm size group). This yields other 9 moments.[92]

Distribution of Earnings (Table 4, Panel B). The coefficient in row (i) of Panel B is from a regression of employee log monthly earnings on the index of managerial ability, reported in column 4 of Table C.3. We then standardize this coefficient by dividing it by the

[91] To preserve the full sample, employees with missing salary are assigned the lowest salary in the sample, and so are included in the low-skilled group. Employees are ranked by salary and split above and below median within each firm size group.

[92] The full list of moments related to the distribution of specialization in complex tasks across the size distribution is reported in the Supplemental Appendix.

standard deviation of employee log earnings.[93] The coefficient in row (ii) is from a regression of employee log monthly salary on log revenues per worker, reported in column 5 of Table C.3. Row (iii) shows the normalized average earnings gap of employees across firms below and above median revenues per worker. To compute this, we regress log employee earnings on region and sector fixed effects, keep residuals, and then normalize these residuals by their mean and standard deviation.[94] In the last row of Panel B we do the same but splitting employees by below/above median managerial ability index of their firm owner.[95]

Distribution of Firm Revenues (Table 4, Panel C). The standard deviation of log revenues reported in the first row of Panel C is after trimming revenues at the 5^{th} and 95^{th} percentiles to reduce the incidence of outliers, and after removing region and sector fixed effects. The coefficient in row (ii) comes from a regression of log revenues per worker on the index of managerial ability, shown in column 1 of Table C.4. In row (iii), we show the average gap in revenues between firms with below and above median managerial ability. To compute this, we regress log revenue on region and sector fixed effects, keep the residual, and then normalize by subtracting the weighted average of the residual.[96] In addition, we also target the pdf of residualized log firm revenues (visualized in Figure 9).[97]

Firm Size Distribution (Table 4, Panel D). Average size in the first row of Panel D is uncensored. The standard deviation of log size and of size in rows (ii) and (iii) is after top coding firm size at 10 workers. The coefficient in row (iv) is from a regression of log size on the managerial ability index, shown in column 2 of Table C.4. Finally, row (v) shows the average gap in firm size between entrepreneurs above and below median managerial ability. To compute this, we create the distribution of firm size (censored at 10 workers), separately

[93]The standard deviation is computed after trimming log earnings at the 5^{th} and 95^{th} percentiles to reduce the incidence of outliers, and after removing region and sector fixed effects.

[94]In the actual estimation, we target the 5, 15, 25, 35, 45, 55, 65, 75, 85, 95 percentiles of the distribution of (normalized) residual salary by above/below median log revenue per worker (so 20 moments, all included in the Supplemental Appendix).

[95]This produces other 20 moments (again shown in the Supplemental Appendix). To preserve the full sample, whenever we split the sample by below and above median managerial ability, we replace missing values in managerial ability by assigning them the lowest value in the sample, so that they are assigned to the low managerial ability group.

[96]In the actual estimation, we target the 5, 15, 25, 35, 45, 55, 65, 75, 85, 95 percentiles of the distribution of (normalized) log revenues by above/below median managerial ability (so 20 moments, shown in the Supplemental Appendix).

[97]To compute this, we regress log revenues per worker on region and sector fixed effects, keep the residual, and then subtract from the residual value its weighted average, and finally trim this value at the 5^{th} and 95^{th} percentile. To estimate the density, we let the program choose 15 points with default settings. So this yields other 15 moments (shown in the Supplemental Appendix).

for above and below median managerial ability firms.[98] In addition, we also target the pdf of firm size (top coded at 10 workers). This gives the final 10 moments used in the estimation (Figure 9).

Table C.3: Moments, Employee Level Regressions

Dep. Var.:	Worker Share of Time in Non. Prod.			log(Salary)	
Sample:	All	Skilled	Unskilled	All	All
	(1)	(2)	(3)	(4)	(5)
log(Salary)	0.033				
	(0.012)				
Firm Size		0.002	0.000		
		(0.003)	(0.004)		
Managerial Ability (Std.)				0.089	
				(0.028)	
log(Revenue per Worker)					0.191
					(0.037)
Firm FE	Yes	No	No	No	No
Region and Sector FE	No	Yes	Yes	Yes	Yes
Obs.	2324	1154	1170	1904	1979

Notes:: Sample: Furniture and Metal products. In Col (2) and (3), employees are classified as Skilled an Unskilled within sector by size groups. Firm size is top coded at 10 workers. Standard errors are robust in column 1, and clustered at the firm level in the other columns. The Managerial ability variable is standardized.

Table C.4: Moments, Firm Level Regressions

	log(Revenue)	log(Size)
	(1)	(2)
Managerial Ability (Std.)	0.145	0.100
	(0.030)	(0.021)
Region and Sector FE	Yes	Yes
Obs.	894	897

Notes:: Sample: furniture and metal products. Robust standard errors. The Managerial ability variable is standardized.

Calibration of Start-up Cost. In Table C.5 we show estimates of the start-up capital (column 1), and compare this with monthly profits in the first year of operation (column 2)

[98]In the actual estimation, we target the 5, 15, 25, 35, 45, 55, 65, 75, 85, 95 percentiles of the distribution of firm size by below and above median managerial ability (so 20 moments, shown in the Supplemental Appendix).

Table C.5: Start-up Capital and First Year Profit

	Start-up Capital	Monthly Profit (first year)	Monthly Profit (time of survey)
	(1)	(2)	(3)
Mean	902.996	106.606	233.749
Median	657.895	65.789	153.509
Obs.	308	303	930
Sample	Follow-up		Initial survey

Notes: Sample: furniture and metal products. All numbers are in USD. Column (1) and (2) show data from the follow-up survey, Column (3) from the initial survey. Start-up Capital definition: see text. We trimmed the top 1% and excluded all 0 values. Monthly profits are trimmed at the top 1%.

and just before the initial survey (column 3).[99] To calculate the start-up capital, we exploit a unique survey module where entrepreneurs were asked to report: (i) all personal savings and (ii) all external sources of funds (e.g., loans, gifts) used to start the business. We sum (i) and (ii) to create a measure of the start-up capital. The average and median of the start-up capital are \$903 and \$657.9. We can benchmark these values by comparing them to the the average and median monthly profit in the first year of operation, which are \$106.7 and \$65.79. Considering monthly discount rates of 1-2%, which seem appropriate for the context, and converting monthly profits to present values, the average start-up cost represents about 8-14% of the present discounted value of profits. We thus calibrate the start-up cost χ_f as 10% of average profits (as shown in Table 5).

D Online Appendix - Quantification

The figure below reports additional statistics for the two main conceptual exercises discussed in Section 6, which calibrate combinations of entrepreneurial pass-through and hiring costs while holding constant either total net output (blue) or average firm size (gray). The corresponding results are discussed in the main text.

[99] The number of observations is around 300 in columns 1 and 2 because the survey module on start-up costs was only asked to a random subset of the sample by design, to limit survey length.

Figure D.1: Constant net output or average firm size (Additional Moments)

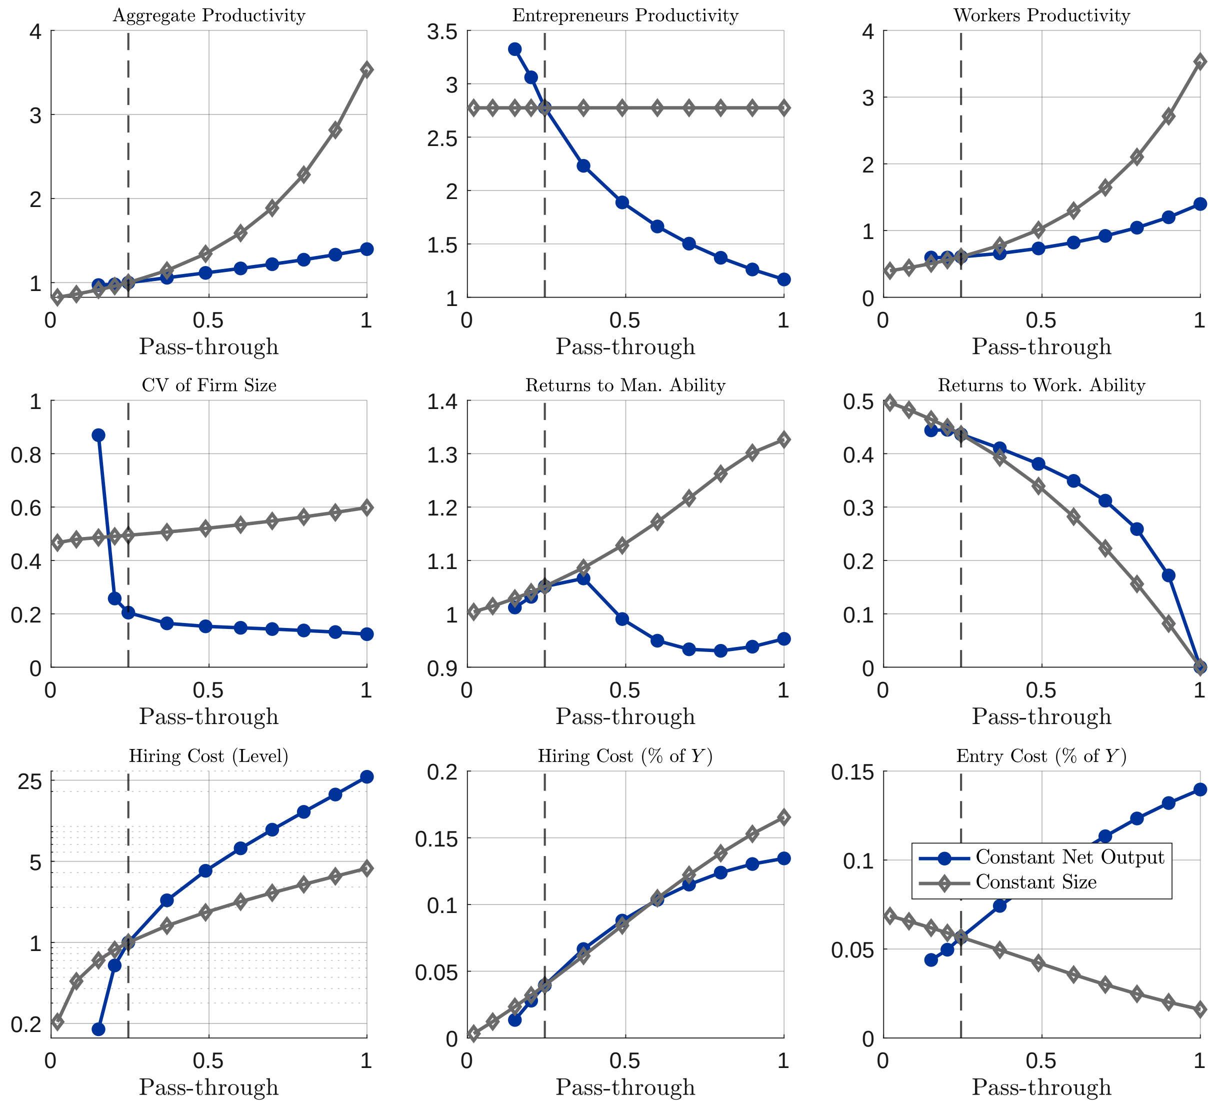

Notes: This figure reports additional aggregate moments for the economies analyzed in Figure 10. As in that figure, the blue line corresponds to the calibration in which net output is held constant while varying pass-through and recalibrating hiring costs. The gray line holds average firm size constant, also by recalibrating hiring costs as pass-through varies. The vertical dashed line represents the benchmark economy.