Self-employment within the firm *

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

September 24, 2023.

Abstract

We collect time-use data for entrepreneurs and their workers in over 1,000 manufacturing firms in urban Uganda. We document limited labor specialization within the firm for establishments of all sizes and argue that this is likely due to the prevalence of product customization. We then develop a general equilibrium model of task assignment within the firm, estimate it with our data, and find large barriers to labor specialization. This setting is close, in terms of aggregate productivity and firm scale, to an extreme benchmark in which each firm is just a collection of self-employed individuals sharing a production space. Given how firms are organized internally, the benefits from alleviating other frictions that constrain firm growth are muted: most African firms resemble artisanal workshops whose business model is not easily scalable.

JEL Codes: O11, O17, L23, L25

*We would like to thank Johannes Boehm for an insightful discussion of the paper. We also thank Oriana Bandiera, Lauren Bergquist, Bruno Crepon, Wouter Dessein, Kevin Donovan, Ben Faber, Luis Garicano, Doug Gollin, Chad Jones, Joe Kaboski, Rocco Macchiavello, Karen Macours, Andrea Prat, Simon Quinn, Imran Rasul, Roland Rathelot, Tristan Reed, Esteban Rossi-Hansberg, Todd Schoellman, John Van Reenen, Eric Verhoogen and Guo Xu as well as seminar participants at Berkeley Haas, Columbia, NYU, Yale, IIES, UCL, Stanford, Madison Wisconsin, Minnesota, Toronto, Toulouse, Notre Dame, Sciences Po, CREST, and PSE. Elena Spadini and Sai Zhang provided outstanding research assistance. All errors are our own.

†University of Southern California, BREAD, CEPR, NBER, J-PAL and IGC; vbassi@usc.edu
‡University of Southern California; leejungh@usc.edu
§New York University and NBER; alessandra.peter@nyu.edu
¶Columbia University, BREAD, CEPR and NBER; tommaso.porzio@columbia.edu
‖Northwestern University; ritwika.sen@kellogg.northwestern.edu
**BRAC NGO; tugumeesau@gmail.com

1 Introduction

Why are manufacturing firms in Sub-Saharan Africa small? The dominant view is that entrepreneurs have poor managerial skills or face external constraints, such as credit or hiring frictions, that impede their ability to scale up.[1] This perspective has guided policy action, leading to billions of dollars of investment in supply-side interventions such as capital transfers or managerial skills training.[2]

In this paper, we explore an alternative view. We show that manufacturing firms in Uganda face barriers to labor specialization, which could keep them small even in the absence of external constraints. This limited specialization is not simply a by-product of firms being small, but rather it originates from the nature of demand they face: firms sell to local consumers who ask for customized products, and when each good is unique, it is difficult to break the production process into separate tasks performed by different individuals. Overall, manufacturing firms in Sub-Saharan Africa resemble hard-to-scale artisanal workshops rather than modern factories.

Understanding why firms are small has profound implications. According to our alternative perspective, the prevailing supply-side interventions on their own may not lead to significant growth since the internal organization of firms leads to low returns to scale. Development policy should therefore prioritize promoting product standards or connecting firms with larger markets to help entrepreneurs move past the artisanal business model and allow them to leverage their talent through labor specialization.[3]

To develop our argument, we conduct a survey of three manufacturing sectors in Uganda and interpret the evidence through the lens of an equilibrium model in which entrepreneurs hire workers and assign them to tasks. The survey measures time use within the firm and shows that labor specialization is limited across the size distribution, especially so in the two sectors where product customization is more prevalent. The estimated model allows us to quantify the internal organization of firms and its implications for the aggregate economy. We find that our setting is close to *self-employment within the firm*, a benchmark in which expanding the firm

[1] See Hsieh and Olken (2014) for evidence on the prevalence of small firms, and Bloom et al. (2013), McKenzie and Woodruff (2008), Banerjee and Duflo (2014), Hardy and McCasland (2022), and McKenzie and Woodruff (2021) for evidence on external constraints.

[2] See Blattman and Ralston (2015); McKenzie and Woodruff (2021).

[3] The artisanal business model was once also prevalent in high-income countries (Atack, 1987; Sokoloff, 1984; Goldin and Katz, 1998; Katz and Margo, 2014) but is now typically confined to the high-end segment of production (Holmes and Stevens, 2014).

is akin to adding a self-employed individual whose productivity is independent of the entrepreneur's. The primary reason why firms exist is thus not to leverage the talent of entrepreneurs, but possibly simply to share fixed production costs. We then show that given this internal organization, the returns to conventional supply-side interventions are dampened.

Our sample consists of about 1,000 firms in carpentry, welding, and grain milling. These three sectors account for approximately 30% of total manufacturing employment in Uganda. The sample is representative across the entire size distribution, allowing us to document the organization of production in both small and relatively large firms—at least by East African standards.

The key innovation of our data collection is to measure time use within the firm. For both the entrepreneurs and their employees, we gather data on which tasks they perform during each hour of the workday. The set of possible tasks includes "production tasks" (e.g., specific step performed in the production process) as well as "non-production" ones (e.g., interacting with customers, supervision, input procurement). To our knowledge, this data is unique, at least in a developing country context. In addition, we collected detailed data on firm characteristics, the production process for specific products, and interactions with customers.

We begin the empirical analysis by describing our setting. The median firm has five employees, showing that there may be scope for labor specialization. At the same time, there are potential barriers to specialization, which we argue are mainly related to the *nature of demand*. Consumers typically ask for customized products with personalized dimensions, details, and finishes. As a result, each firm produces only few items of each product, making it difficult to set up a production line. Moreover, customization entails significant communication and coordination costs within the firm, making it difficult to "unbundle" the production process into separate tasks that can be performed by different individuals.[4] These "artisanal" features of production are prevalent among smaller and larger firms alike, but are more common in carpentry and welding, thus offering useful empirical variation to validate the connection between artisanality and labor specialization.

Using the time-use data, we then describe the extent of labor specialization within

[4]The link between standardization, specialization, and scale of operations has been established in the literature both empirically and theoretically (Piore and Sabel, 1984; Holmes and Stevens, 2014; Vickery et al., 1999; Dessein and Santos, 2006).

and between firms and how it varies along the firm size distribution. We pool the data for the two sectors in which artisanal production is most prevalent—carpentry and welding— and focus on these, but later study heterogeneity across all three sectors.

First, we show that there is no specialization *between* firms: most firms make the products we asked about and engage in almost all production steps. In addition, the share of time spent in different tasks does not vary along the size distribution. Larger firms simply operate as replicas of smaller ones, doing more of the same tasks.

Then, we analyze the extent of horizontal and vertical specialization *within* the firm. To measure vertical specialization, we focus on production versus non-production tasks, which we show to be more skill-intensive. Entrepreneurs, who are more skilled than their employees, do spend more time on non-production tasks; yet, the tasks performed by employees and the entrepreneur they work for display significant overlap. Moreover, while this type of labor specialization increases with firm size, it does so only in a limited way: even in firms with more than five employees, entrepreneurs spend only 50% of their time in non-production activities, even though there would be more than enough non-production tasks to fill the entrepreneur's time.

Second, we analyze horizontal, or "Smithian", specialization by measuring time spent on different tasks within production. We show that this type of specialization is almost entirely absent: on average, 85% of employees work on each production step, and this percentage varies little with firm size. If employees were fully specialized across steps, each of them would only need to do 25%–30% of the steps.

In summary, we find that labor specialization within the firm is limited, but vertical specialization is relatively more prevalent than horizontal specialization. The fact that specialization is limited in both small and large firms implies that the reason why firms do not specialize cannot simply be that they are too small to do so.

Finally, we replicate the analysis for grain milling and uncover much higher specialization across all dimensions. Linking back to the prevalence of artisanal production in the different sectors, this result validates the claim that product customization creates a significant barrier to labor specialization. It also highlights that our survey is able to properly capture labor specialization where present and that lack of managerial skills or institutional features such as limited contract enforcement cannot be the only drivers of limited specialization.

Motivated by the data, we develop a model to formalize the two-way relationship between labor specialization and firm size and to characterize and quantify how

barriers to specialization affect firm-level and aggregate productivity.

The heart of the model is an assignment problem of heterogeneous workers to tasks, which determines firm productivity. We embed this problem into an occupational choice model that is standard but for two features. First, entrepreneurs are subject to a convex hiring cost that captures any external constraints that may keep firms small (e.g., credit constraints). Second, worker earnings consist of a piece-rate component (increasing in their productivity) as well as the equilibrium wage level.

For a unit of output to be produced, a fixed set of tasks must be completed, which differ in their level of complexity. If self-employed, everyone must complete all tasks themselves. When working together in a firm instead, individuals can unbundle the production process and delegate some tasks to others.

Firm productivity has two components. The first one depends on the entrepreneur's ability and reflects the non-rival role of talent in production (e.g., a business idea). The second one is an average of the abilities of all individuals producing, weighted by the share of time each spends on the complex tasks: while everyone can complete the simple tasks equally well, high-ability individuals are better at the complex ones. This second component depends on the allocation of talent within the firm and implies that delegating complex tasks to more skilled individuals—that is, specializing labor based on comparative advantage—increases firm productivity. Unbundling tasks, however, comes at a cost, which encapsulates the barriers to labor specialization.

In the model, the extent of *artisanality* is modulated by two key parameters. The first one, λ, determines the non-rival role of talent in production; the second one, κ_0, determines the severity of barriers to specialization. When the unbundling cost is low or entrepreneurial ability is highly non-rival, firm productivity is mainly a function of the entrepreneur's ability and firms are vehicles for leveraging and scaling the talent of entrepreneurs. When the cost is large and ability is rival, each worker is essentially *self-employed within the firm* and firm productivity is equal to the average ability of all individuals. In this case, firms are mere vehicles to share fixed costs. The efficient firm size is smaller because of strong decreasing returns to scale generated by the quick dilution of the entrepreneur's talent with the one of less-skilled workers.

The way in which firms are internally organized has equilibrium effects that ripple through the economy. Higher labor specialization increases firm productivity and, thus, labor demand. As a result, wages increase, leading some marginal entrepreneurs to become workers, further increasing aggregate productivity through a

classic selection effect. Overall, when the unbundling cost is low, managerial ability is highly valued in the economy, talent can be leveraged by taking over more and more complex tasks, and only large, high-productivity firms operate.

We extend the model to accommodate additional heterogeneity and estimate it using data from carpentry and welding. We target a rich set of moments on the within-firm allocation of labor to tasks and across-firm heterogeneity in size, revenues, and worker earnings. While all parameters are jointly estimated, we offer a heuristic identification argument which we verify through model simulations.

With the estimated model, we perform three exercises. First, we vary the extent of artisanality. We find that our setting is quite close, in terms of firm size and productivity, to the polar case of self-employment within the firm, in which entrepreneurs cannot pass through any of their talent to workers. This result provides a quantitative answer to the question, Why do firms exist? In Uganda, the role of firms as vehicles for leveraging entrepreneurial talent is, at best, limited.

Second, we study the mechanism through which barriers to labor specialization affect the economy and contrast them with other constraints that keep firms small. Decreasing the unbundling cost has two effects: (i) labor specialization increases, leading to a larger pass-through of entrepreneurial talent into worker productivity; (ii) in equilibrium, labor is reallocated toward more skilled entrepreneurs. Decreasing the hiring cost has a similar reallocation effect, but a smaller impact on specialization, which is purely a by-product of having larger firms. We learn that, while the model entails a two-way relationship between firm size and specialization, the link from specialization to firm size is quantitatively stronger. In this sense, our analysis is more consistent with the notion that firms are small because they are not specialized, rather than that they are not specialized because they are small.

Third, we show that the benefits of interventions aimed at spurring firm growth hinge on the size of the unbundling cost. Relative to our benchmark, calibrating the cost to match the larger specialization observed in grain milling would increase the productivity effect of a reduction in hiring cost by 60% and the effect on firm size by 35%. Similarly, a business training program that increases the ability of the top 10% of the population would yield a much larger aggregate return when the unbundling cost is low. These results highlight the key takeaway of our work. Barriers to within-firm labor specialization make artisanal manufacturing a business model that is difficult to scale, thus limiting the returns from supply-side policy interventions.

Related Literature and Contribution. In offering an alternative perspective on why firms are small, we contribute to the large literature studying firm size and productivity in developing countries (Bloom et al. (2010); Hsieh and Olken (2014)). Our paper is closest to the work that emphasizes the role of management (Bloom et al. (2013); Bruhn et al. (2018); Anderson and McKenzie (2022)) and limits to delegation (Akcigit et al. (2021)). These studies find that poor managerial practices and contractual as well as labor market frictions impede firm expansion and lower their productivity. We highlight how barriers to labor specialization *inside the firm* prevent entrepreneurs from leveraging the talent they already possess.[5]

We build on a classic theoretical literature in organizational economics, which has long emphasized the importance of labor specialization within the firm for productivity and growth (Chandler (1990); Becker and Murphy (1992); Bolton and Dewatripont (1994); Yang and Borland (1991); Garicano and Rossi-Hansberg (2006)). Our contribution is to offer a case study that shows how the internal organization of firms could help us understand why firms are small in developing countries.

Our model is original, but builds on the seminal work on the organization of knowledge into hierarchies (Garicano (2000); Garicano and Rossi-Hansberg (2006)). Like those papers, we focus on the vertical specialization of talent based on comparative advantage. Our model, to the best of our knowledge, is unique in allowing for an overlap in the tasks performed by individuals in different layers of the organization. In previous work, the size of each layer (and even their number) is endogenous, but the assignment of tasks to layers is fixed: higher layers fully specialize in more complex tasks. In our environment, instead, the unbundling cost modulates the extent to which such specialization is possible. This departure is motivated by the ample evidence of task overlap in our data.[6]

Many before us have provided empirical evidence for and quantitative assessments of the role of labor specialization for productivity and growth.[7] In addition to our

[5]Our results also provide a plausible explanation for the low correlation between managerial skills and firm size in developing countries (Bloom et al., 2022).

[6]Studies of organizational adaptation to changes in local conditions also highlight the difficulty of coordination across unbundled production tasks (Dessein and Santos, 2006; Caliendo and Rossi-Hansberg, 2012; Adhvaryu et al., 2023).

[7]For example, Caliendo et al. (2015) use occupational data to study how French firms are organized, and Bandiera et al. (2022) to compare labor specialization across countries. A related literature exploits data on tasks within the firm to study horizontal specialization within production during the Industrial Revolution in the US (Atack et al., 2019, 2023) and in New York City hairdressers (Kohlhepp, 2023). Boehm and Oberfield (2023) use production data to study task spe-

unique focus and setting, we make a methodological contribution. We show the importance of collecting time-use data: In our 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.

Finally, a related literature studies the role of frictions in output markets as a barrier to growth (Bold et al. (2022); Jensen and Miller (2018); Hjort et al. (2020); Startz (2019); Vitali (2022)). We add to this literature by showing how a specific feature of the nature of demand—the prevalence of customization—impacts firm productivity and size by affecting the internal allocation of labor.[8] More broadly, our findings reinforce the view that demand-side constraints play a primary role for development (Goldberg and Reed, 2022).

Structure of the Paper. In Section 2, we describe the survey, and in Section 3, we document the prevalence of artisanal production. Section 4 shows evidence on labor specialization. Section 5 develops the model, Section 6 describes the estimation, and Section 7 reports our quantitative results and counterfactuals. Section 8 concludes. Additional results are in the Online Appendix.

2 Survey: Measuring within-Firm Organization

This section describes the survey we conducted in urban Uganda to study labor specialization inside the firm. We collected two waves of data: an initial survey in 2018–2019 and a follow-up survey survey in 2022. The key innovation of our data collection is to measure time-use data within the firm. Therefore, we focus the discussion on the novel survey modules designed to measure labor specialization.[9]

cialization across firms in India, while Freund (2022) uses wage data to study how labor specialization and sorting affects inequality in Germany.

[8] Jensen and Miller (2018) is a study particularly related to ours in that they show that firms specialize labor as they grow larger. While we also show that small firm size reduces specialization, our key focus is to show that barriers to specialization hinder firm size in the first place, and to isolate and quantify each channel of the two-way relationship between specialization and firm size.

[9] The other modules feature in our previous work and are described in Bassi et al. (2022b).

2.1 Sampling

Our sample consists of firms in carpentry, welding, and grain milling. We chose these sectors for two reasons: (i) they are large, employing about 30% of workers in manufacturing, and (ii) they include both smaller and—for Ugandan standards–larger firms, which allows us to study labor specialization across the size distribution.[10]

We selected a representative sample of 52 sub-counties, stratifying by population and whether the sub-county is in Kampala, the capital city.[11] Within each sub-county, we first conducted a complete listing, identifying close to 3,000 establishments in these sectors. We then randomly sampled about 1,000 establishments from the listing.[12] We interviewed the entrepreneur and all employees working on pre-specified "core" products that are common in each sector: doors in carpentry, windows in metal fabrication, and maize flour in grain milling. Our final sample includes 1,115 entrepreneurs and 2,883 employees.[13] In Appendix A.2, we compare our sample with administrative data and show that we properly cover both small and large firms.

2.2 Survey Design

Our key innovation is to collect granular measures of labor specialization inside the firm, which we describe in detail below. In addition, we collected detailed information on the production process of firms as well as the economic environment in which they operate. Specifically, we asked firms about (i) production steps and machines used to produce the core product; and (ii) features of the output market, including prices, customers, and marketing activities; (iii) characteristics of entrepreneurs and employees, including an index of managerial ability for entrepreneurs (as in McKenzie and Woodruff (2017)).

[10]The latest Census of Business Establishments from the Uganda Bureau of Statistics from 2010 shows that these three sectors comprise 32% of total manufacturing employment and 27% of manufacturing employment in firms with five or more employees.

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

[12]In practice, we over-sampled firms with more than five workers to ensure enough observations among relatively larger firms. Since in the great majority of cases these are single-establishment firms, we use the terms "firm" and "establishment" interchangeably in the rest of the paper.

[13]In the rest of the paper, we use the terms "entrepreneur" and "owner" interchangeably, since in most cases they are the same person. Compliance with the survey was over 90%, and all our results are appropriately weighted to reflect our sampling strategy. See Bassi et al. (2022b) for details.

Measuring Labor Specialization. To measure labor specialization, we exploit two survey modules, each of which were directed to both the entrepreneur and the employees. The first is a time-use module, where the respondent was asked to report all the hours worked for the firm in their last day of work. For each hour worked, we asked them to indicate the specific tasks they performed, choosing from a pre-specified list including tasks related to "production", "non-production", and "idle" time. On production, we differentiate between working on the core product or another product, and in case of the core product, we also asked about the specific production steps performed in that time slot.[14] The list of non-production tasks encompasses all other 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, or management of stock. Finally, for idle time we know the time spent eating/resting or away from the firm for non-business reasons. In Appendix A.1, we show the list of 17 tasks measured in our time-use survey, together with the overall share of time spent by the entrepreneur and all employees in the average firm in each of these tasks.

The second module asked which production steps for the core product the respondent *usually* performs (without limiting to the last day worked), as well as the hours spent on each production step in days when they work on that step. 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.

Measuring Artisanal Production. The follow-up survey collected additional details on labor specialization inside the firm, product characteristics, and interactions with customers to shed light on the prevalence of artisanal production in this context and how customization may create a barrier to labor specialization.[15]

[14]For firms not producing the core product, we have information on time spent producing their main product, without the breakdown in production steps.

[15]The follow-up survey targeted the entire sample of entrepreneurs and employees. It was conducted through phone surveys, and the attrition rate is about 32% for entrepreneurs and 41% for employees (see the Supplemental Appendix for details). This survey is used to provide qualitative evidence on labor specialization and prevalence of artisanal production. As described in Appendix C, none of the moments used for estimation come from this survey; we rely on this follow-up survey only for one calibrated parameter.

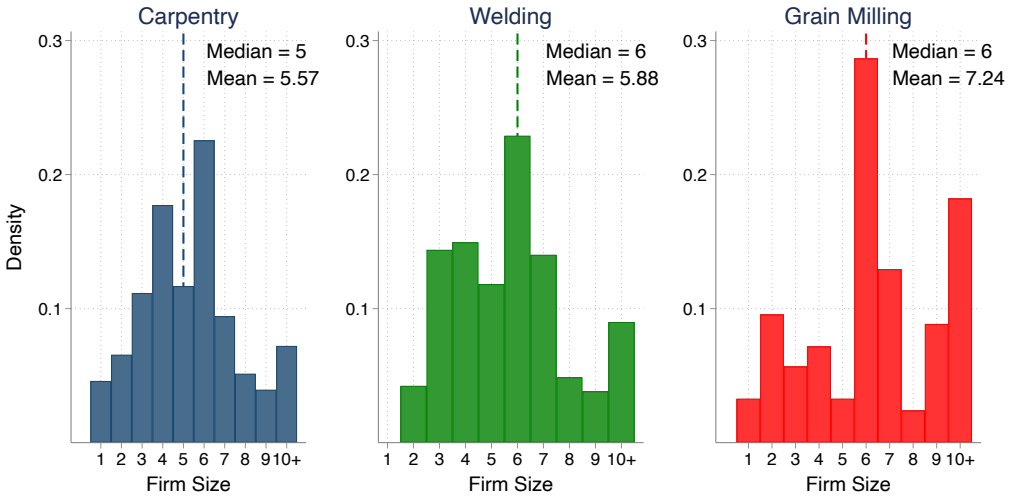
3 Background: Artisanal Production

We use our survey data to show why the carpentry, welding, and grain-milling sectors in Uganda provide an ideal setting to study labor specialization within the firm. First, firms are large enough for there to be potential gains from specialization. Second, the nature of demand is such that artisanal production entailing customized orders is prevalent, which we argue can plausibly create a barrier to specialization.

3.1 Basic Firm Descriptives

Figure 1 plots the firm size distribution in the three sectors. In line with previous findings of the development literature (Hsieh and Olken (2014)), most firms employ less than 10 workers. However, they are *not* micro-enterprises: the median firm employs 6 workers, which implies that there is, in principle, scope for labor specialization.

Figure 1: Firm Size Distribution



Notes: Sample: all surveyed firms. Firms with more than 10 workers are grouped together in the category "10+". Vertical lines represent the median. A version of this figure without top coding at 10 workers is in the Supplemental Appendix.

In Appendix A.2, we also show that (i) there is substantial dispersion in revenues per worker which is systematically correlated with managerial ability, suggesting that there could be gains from reallocating resources to higher-ability managers; (ii) entrepreneurs earn substantially more than employees; (iii) employees are paid primarily piece-rate and are less educated than entrepreneurs; and (iv) output markets are local-

ized, informal face-to-face interactions with customers are prevalent, and marketing investments are limited, which is in line with the literature.[16] Importantly, these key features of the output market do not vary with firm size.

Representativeness of Our Sample. In Appendix A.2 we compare our sample to two sources of administrative data: the 2010 Census of Business Establishments, which should include all establishments, and the Corporate Income Tax (CIT) data for 2018/19, which is supposed to cover all large firms with more than $40,000 in yearly revenues. We find that in the CIT data there are handful of very large firms that we are not able to cover in our sample. Nonetheless, we show that our sample is well equipped to describe the typical large formal firm as there is substantial overlap between the distribution of firms registered in the CIT data and those in our data.

3.2 Prevalence and Implications of Artisanal Production

Manufacturing in advanced economies is typically characterized by a large "modern" segment making *standardized* products with production line techniques (Holmes and Stevens, 2014; Piore and Sabel, 1984). An alternative is what we define as "artisanal" manufacturing, where products are *customized* for individual consumers and quality heavily depends on the skills of the craftsman working on production. In this section, we show that the nature of demand is such that artisanal manufacturing is prevalent in Ugandan manufacturing, especially in carpentry and welding. Then, we use our data to argue that the prevalence of artisanal manufacturing may limit the extent of labor specialization. Finally, we discuss possible reasons why artisanal production is widespread in carpentry and welding but less so in grain milling.[17]

Prevalence of Artisanal Production in Carpentry and Welding. Panel A of Table 1 presents several pieces of evidence in line with the notion that customization is widespread in carpentry and welding, but limited in grain milling. While the majority of firms in the three sectors sell on order, the underlying reasons are different: in carpentry and welding, customers buy on order to customize products, whereas in

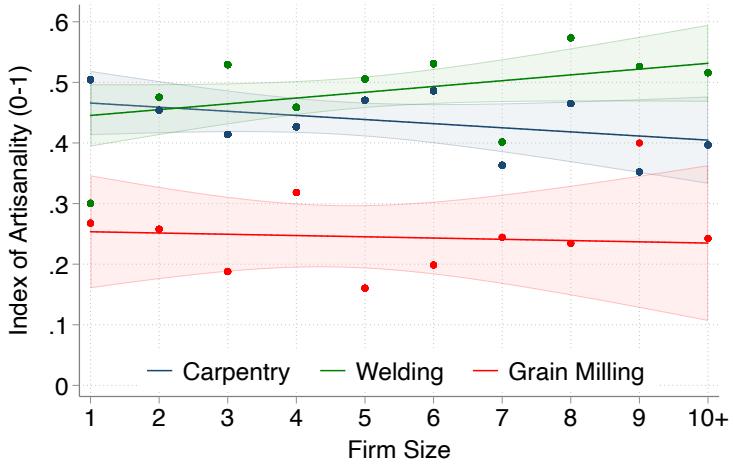
[16] See, for instance, Startz (2019), Bassi et al. (2022b), Bassi et al. (2022a) and Vitali (2022).

[17] Holmes and Stevens (2014) show that firms producing customized products in the US represent a small share of the market serving high-end consumers. In contrast, we show the broad prevalence of customization in Ugandan manufacturing.

grain milling, they do so in order to bring their own maize to be processed into flour. In addition, only in carpentry and welding do firms cite product customization as a key reason for charging different prices for the same type of product and for not having ready-made products in stock.

To summarize this cross-sectoral heterogeneity, we aggregate the information from Panel A into one standardized index of "artisanality". Figure 2 shows that: (i) artisanal production is much more common in carpentry and welding, and (ii) small and large firms alike engage in artisanal production.

Figure 2: Relationship between Artisanal Production and Firm Size



The figure shows the prevalence of artisanal manufacturing by firm size in each of our study sectors. The index of artisanality (range: 0-1) is an average of four binary variables from Panel A of Table 1 that reflect the extent of customized production in a firm: (i) sales are made to order; (ii) the firm does not have finished products in stock because of customization; (iii) customers buy on order to customize products; and, (iv) the firm charges different prices for the same product because of customization. Sample: all sectors.

Implications of Artisanal Production. We document two reasons why the prevalence of customization in artisanal manufacturing may limit labor specialization.

First, quality uncertainty in customized production leads to sizable communication costs within the firm: making a customized order requires agreeing on various product features such as design and materials, and the quality of the product critically depends on whether the craftsman producing has understood the specific product features each consumer asked for and is able to create them. This may make it difficult for the firm to "unbundle" the production process into separate tasks that can be performed by different specialized individuals, as doing so would require substantial

Table 1: Artisanal Production across Sectors

	Carpentry	Welding	Grain milling
	(1)	(2)	(3)
Panel A. Prevalence of Artisanal Production			
Share of sales made to order	75%	89%	69%
Why customers buy on order: Customization	65%	65%	26%
Why customers buy on order: Bring own inputs	5%	5%	52%
Reason for charging different prices: Bargaining	64%	60%	59%
Reason for charging different prices: Customization	45%	54%	18%
Reason for charging different prices: Quantity discounts	21%	28%	53%
Why no products in stock: No money for inputs	46%	45%	38%
Why no products in stock: Customization	20%	14%	0%
Why no products in stock: Customers bring own inputs	0%	1%	26%
Panel B. Consequences of Artisanal Production			
80-20 price dispersion within sub-county for main product	1.46	1.61	1.15
Within-firm ratio of highest to lowest price for same product	1.43	1.31	1.13
Customers pay fully upfront	26%	29%	54%
Reason for orders: Discuss details 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 C. Drivers of Artisanal Production			
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 are reported. Panel A, rows 2–3: dummies for main reason why customers buy on order (we label as "Customization" the two answer options "Customers want to choose the materials/inputs" or "Each customer wants a different product"). Panel A, rows 4–6: dummies if reason listed among top three for charging different prices for the same product. Panel A, rows 7–9: dummies for main reason for not having products in stock (30% of firms do not have products in stock). Panel B, row 1: dummy if discussing details with person producing is listed as top 3 reasons why customers buy on order. Panel B, row 3: dummy if employees perform independent orders. Panel B, row 6: price dispersion is for doors, windows, and flour, after removing sub-county fixed effects. Panel B, row 7: number of different varieties of doors, windows, and flour in the sample. Rows 4–6 of Panel B are conditional on the firm having at least one employee. Panel C, row 1: number of distinct machine types used to produce doors, windows, and flour.

communication and coordination costs.

Panel B shows evidence on both quality uncertainty and communication costs driven by customized orders. Consistent with quality variation and uncertainty being

higher in carpentry and welding, price dispersion for the same product both across and within firms is higher in these sectors, and paying at delivery (another proxy of quality uncertainty) is also more common. Panel B further shows that carpenters and welders are twice as likely to report that customers buy on order because they want to discuss the details with the person producing, and customers communicating directly via phone with the person producing is also twice as common in grain milling. In the extreme, the communication costs related to artisanal production may favor the emergence of a production arrangement whereby a single employee manages the entire production of the order and the relationship with the customer all by themselves through "independent orders"—as though they were self-employed. This arrangement has emerged in practice, and is particularly common in carpentry and welding.

Second, as each door or window can have different features, effectively customers can choose from a very large menu of possible products and each firm makes a small quantity of each product. This may make it difficult to set up a standardized production line for each possible product whereby workers specialize in different tasks. To get a sense of this, the last row of Panel B shows that in our data, we document 13 different types of doors and 7 types of windows being made. 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 shape and size of the two-panel door), and thus provide a substantial underestimate of the size of the menu that customers can choose from. By contrast, in grain milling there are only 4 types of flour being made in our data, so that setting up standardized processes is likely simpler as firms produce a larger quantity of each product.

The communication and coordination costs arising from product customization plausibly lead to idle time in production as workers might be idle while waiting for others to finish their tasks. In line we this, in Appendix A.3 we show that idle time is much larger in carpentry and welding, and this is true in both small and large firms.

Why is Artisanal Production Common in Carpentry and Welding? Finally, we discuss possible reasons why artisanal production is widespread in carpentry and welding but limited in grain milling. One key difference between these sectors is the product inherent complexity. In carpentry and welding, products can be made with a multitude of features and finishes: as shown in Panel C, 24 different types of machines can be used to make the doors in our sample, and 20 different machines are used to

make windows. Doors and windows also take several hours to make, usually over multiple days. The scope for customization and quality variation is therefore high, as products are complex. This is not the case in grain milling, as flour is by nature a much simpler and more standard product: grain millers use half the number of machines, and production takes one-sixth of the time or less.[18]

One important question, of course, is the role that the institutional environment plays. In principle, building codes could facilitate product standardization in sectors such as carpentry and welding (e.g., by specifying standard sizes for door and window frames in buildings). While building codes are present in Uganda, anecdotal evidence from our field visits confirms that they are loosely followed or enforced.

For example, one relatively large carpenter at the top of the size distribution reported that even when they get orders of doors for large formal buildings (e.g., hospitals), the size of the door frames usually vary from building to building. According to his report, this uncertainty related to the size of the doors is a primary reason why they have not been able to set up a production line with standardized products and processes. When asked about how the firm is organized, he said: "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 immediately replied: "I would set up a production line. In fact, I also have a snack factory, and there we have a production line".[19]

Based on these results, in the next section we begin by focusing on labor specialization in carpentry and welding, as we expect limits to labor specialization to be more significant there because of the prevalence of artisanal production. We then validate our results by contrasting them with grain milling.

4 Evidence: Limited Specialization of Labor

In this section, we describe the organization of labor inside the firm and how it varies across the size distribution. As a preliminary step, we describe the tasks firms do.

[18] 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).

[19] During this interview, we discussed several potential challenges in operating the business. The respondent was informed that we were conducting research on productivity in the carpentry sector but not of our focus on labor specialization. Importantly, the interview was executed in English.

We then turn to our core analysis and document the extent of labor specialization *within the firm* across different tasks. We start by using only data for carpentry and welding, and at the end of the section, we contrast the results with grain milling.

4.1 Task Composition: What Do Firms Do?

We document which tasks firms do and how the composition of tasks varies across the size distribution. We find two results. First, firms perform the great majority of the production steps for the core product in-house. Second, firms of all sizes spend similar time shares on each task.

Which Production Steps Do Firms Do? Since we collected data on production steps for the core product, 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.[20] Panel (a) of Figure 3 shows that: (i) each step is done by most firms, and (ii) this does not vary across the size distribution.

How Do Firms Allocate Time Across Tasks? Panels (b)-(d) of Figure 3 plot, for each firm size, the share of time spent on different tasks.[21] 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 if we zoom in within production or within managerial activities (Panels (c) and (d)), we note very little variation in task composition across the firm size distribution.[22]

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

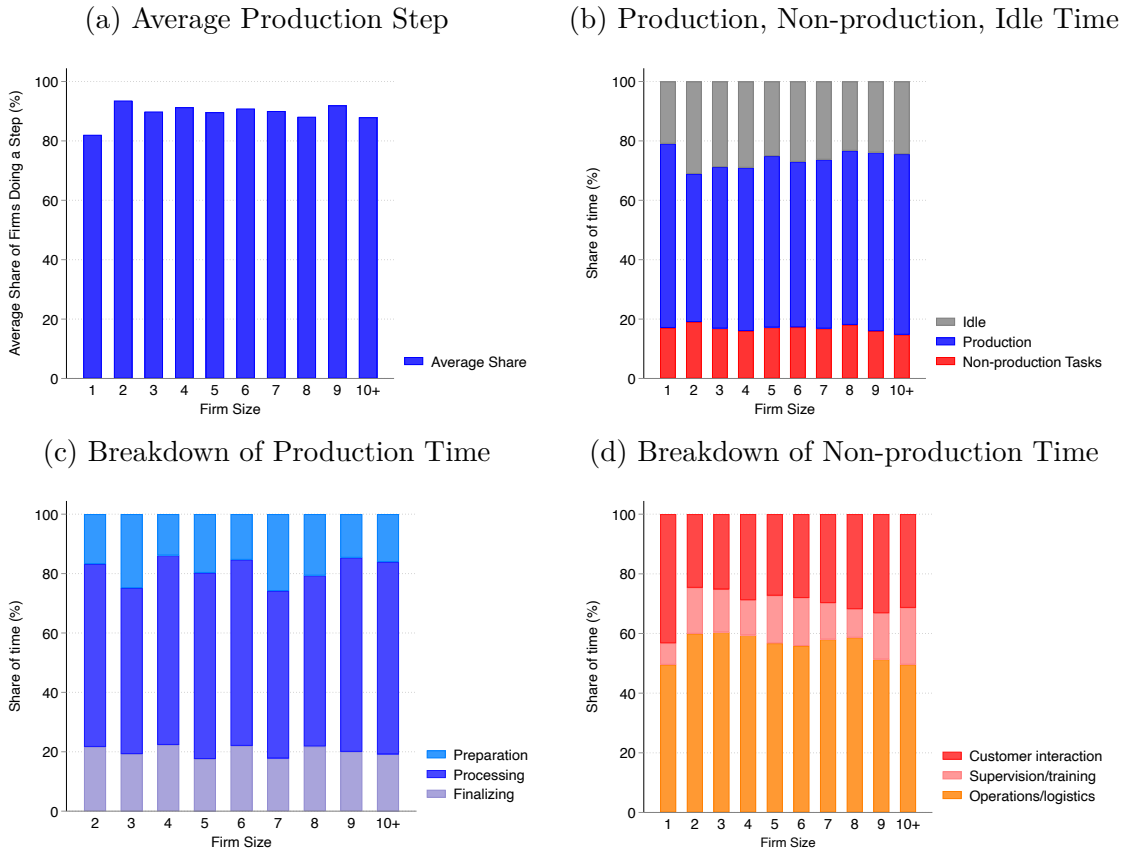
[20] The core product has 10 production steps in carpentry and 7 steps in welding. See Appendix A.1 for details. We average across steps, weighting by the average share of time that production step accounts for in the data, so that steps that represent a larger fraction of total production time get a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

[21] Panel (c) uses 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.

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

customer interactions. 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 operate as replicas of smaller ones, simply doing more of the same tasks.[23]

Figure 3: Task Composition across the Size Distribution



Notes: Sample: carpentry and welding sectors. Panel (a): we compute the average share of firms doing each step and then average across steps as described in the text. Panel (b): share of firm-level time in Production, Non-Production, and Idle tasks (see Appendix Table A.1 for details). Panel (c): breakdown of the production time of the core product into Preparation, Processing, and Finalizing, which include the following steps. "Preparation": (i) Carpentry: Design-Drying (before production), (ii) Welding: Design; "Processing": (i) Carpentry: Cutting-Mortising, (ii) Welding: Cutting-Welding; "Finalizing": (i) Carpentry: Finishing-Drying (after painting), (ii) Welding: Polishing-Painting. Panel (d): 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. In Panels (a) and (c), the sample is restricted to firms making the core product.

[23] In the Supplemental Appendix, we show that the results in Figure 3 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.

4.2 Task Allocation: Who Does What Within the Firm?

In this section, we study the division of labor inside the firm. Measuring labor specialization could be a very complex exercise as there are multiple possible margins of specialization. We thus organize the analysis by focusing on two main margins of specialization: (i) between production and non-production tasks and (ii) within production across steps. (i) is motivated by the fact that non-production tasks are more skill-intensive and entrepreneurs are more skilled than employees, which we verify in our data. (ii) is motivated by the classic "Smithian" specialization: as in the pin factory described by Adam Smith, individuals can increase their proficiency by specializing in a narrow production task, leading to an overall increase in productivity. On both margins, we allow for specialization between employees as well as between the entrepreneur and her employees. We first document the patterns of labor specialization and then summarize the implications of these findings.

4.2.1 Labor Specialization Between Production and Non-production Tasks

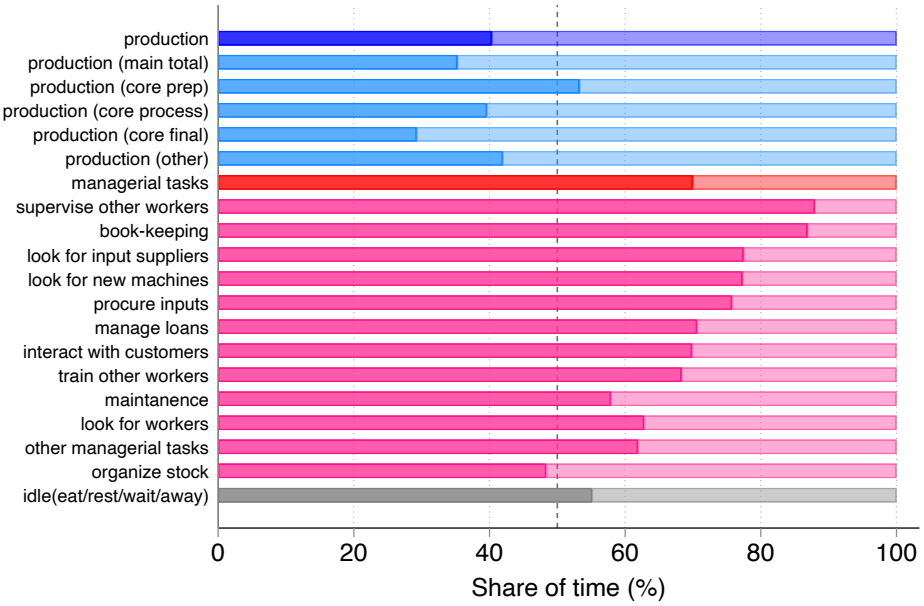
In Figure 4, we compare the time that the entrepreneur and the average employee spend on each task. The y-axis shows the different tasks: blue ones are related to production, red ones to non-production, and grey ones to idle time. Each bar reports the share (normalized to 100%) of that task done by the entrepreneur (the dark portion of the bar) and the average employee (the light portion). If the entrepreneur and the average employee spend the same amount of time on a given task, the dark and light bars would each amount to 50%.

Figure 4 offers two takeaways: (i) entrepreneurs specialize in non-production tasks; and (ii) still, there is substantial overlap between entrepreneurs and the average employee in terms of time allocation, so that specialization is present, but limited.[24]

In Appendix A.3, we compare instead the time allocations of higher-skilled versus lower-skilled employees. In this case, we find an almost complete overlap across all tasks: any specialization between employees along this margin is essentially muted.

[24]The fact that entrepreneurs specialize in all non-production tasks (apart from organizing stock where the entrepreneur and the average employee spend a similar amount of time) and in none of the production tasks (apart from preparation of the core product where again the time share of the entrepreneur and the average employee is similar) also confirms that production and non-production tasks are different, thus justifying our partitioning of tasks into production and non-production.

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

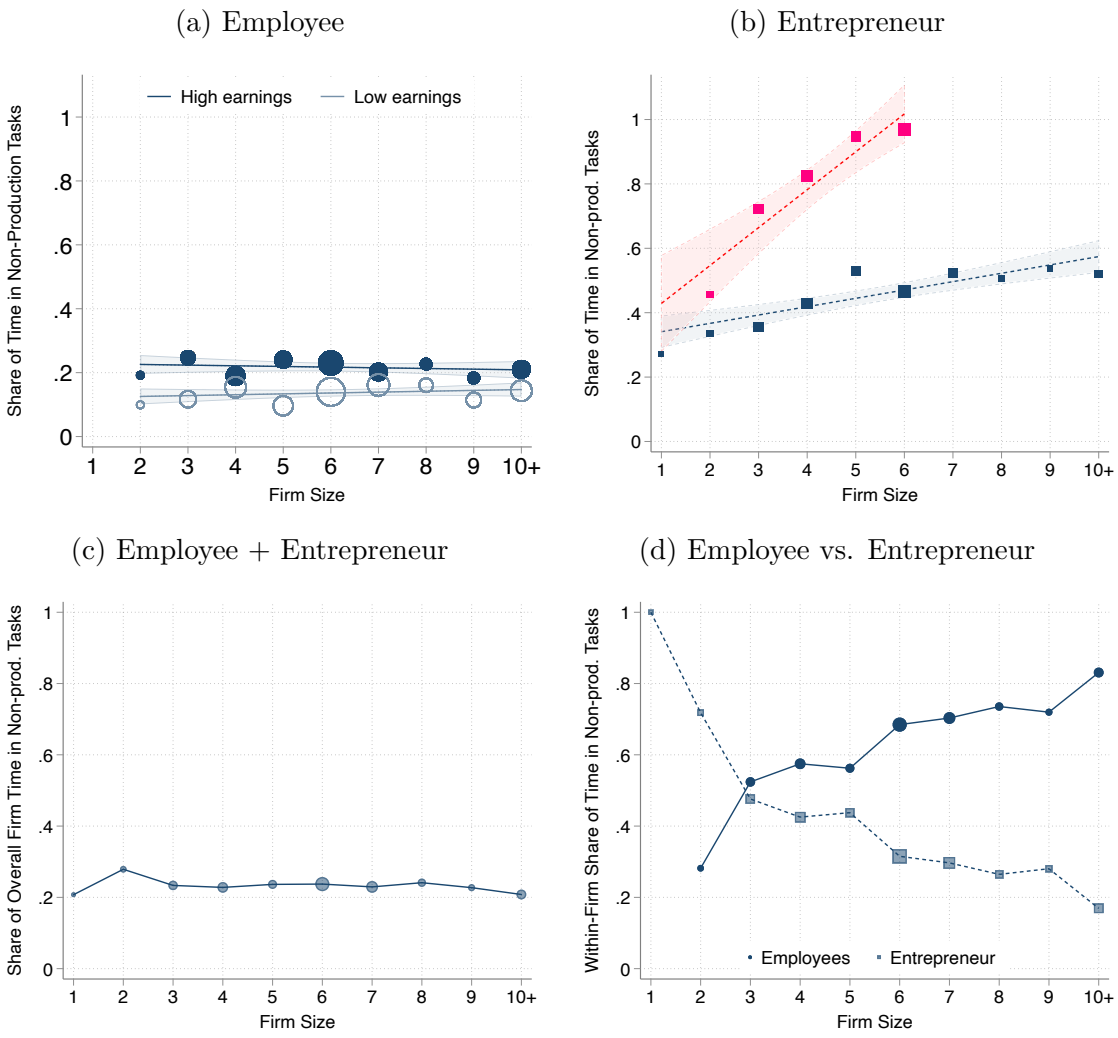


Notes: The figure compares the time spent on each task by the entrepreneur (dark bars) and the average employee (light bars). Blue bars: production tasks. Red bars: Non-production tasks. Grey bars: Idle time. "Production (core prep)", "Production (core process)" an "Production (core final)" refer to the following production stages of the core product: "Preparation", "Processing" and "Finalizing". See Figure 2 for more details on which production steps map to these production stages. Sample: all surveyed firms in carpentry and welding sectors. Time use reported by interviewed entrepreneurs and employees. All figures are weighted by sampling weights within each sector and the relative number of surveyed firms per sector.

More Specialization of Entrepreneurs in Larger Firms. In Figure 5, we study how the specialization of entrepreneurs in non-production tasks documented in Figure 4 varies across the firm size distribution. To do so, we plot the average individual's share of time spent in non-production tasks as a function of firm size, for both employees (Panel (a)) and entrepreneurs (Panel (b)).[25] The figure confirms that specialization among employees is limited and it also shows that it does not vary with firm size: employees spend about 20% of their time in non-production activities. Furthermore, while high-skilled employees (as measured by earnings) spend a little more time on non-production tasks, the gap relative to low-skilled employees is small and does not vary across the size distribution. On the other hand, entrepreneurs do specialize in non-production tasks, and the gap relative to employees increases in firm size: larger firms are more specialized.

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

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



Notes: Sample: all surveyed firms in the carpentry and welding sectors. 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, where all available non-production tasks done by anyone within the firm are reassigned to the entrepreneur. Panel (c): The total (entrepreneur + employee) share of firm time in non-production tasks. Panel (d): Breakdown of the total firm-level non-production time between the share supplied by the entrepreneur and that supplied by all employees combined.

Even in Large Firms, Entrepreneurs Are Not Fully Specialized. Panel (b) of Figure 5 further shows that while entrepreneurs in larger firms take on more non-production tasks, specialization only weakly increases with firm size: the coefficient from a regression of the share of entrepreneur's time spent on non-production tasks

on firm size is 0.022. This implies that when going from a firm of size one to a firm of with five workers, the share of time in non-production activities only increases from about 34% to 45%. So, even in firms with five or more employees, the entrepreneur spends only about *half* of her time on non-production activities.

One possibility could be that there are simply not enough non-production tasks to keep entrepreneurs busy. To show that this is not the case, we compute, for each firm, the (counterfactual) share of time that the entrepreneur would spend in non-production tasks if she had fully specialized in these tasks.[26] This empirical full-specialization benchmark is depicted in pink. We see that the observed relationship between specialization and firm size is closer to a flat line than to the empirical benchmark. This highlights that limited specialization is not merely an artifact of firms being small.

We present additional robustness checks in Appendix A.3, where we show that: (i) the results in Panels (a) and (b) of Figure 5 are robust to disaggregating non-production tasks in their different sub-categories; (ii) the results in Panel (a) are not driven by some workers within the firm specializing in production while others specialize in non-production tasks; (iii) our measurement of non-production tasks is consistent across the size distribution; and (iv) the substantial involvement of entrepreneurs in production tasks is not driven by apprenticeship motives.[27]

4.2.2 Labor Specialization Between Production Steps

We examine specialization within production of the core product, across steps, to study whether this margin of division of labor is important.[28] In Figure 6, we plot the share of employees and entrepreneurs performing a production step by firm size. To do so, we compute the share of employees performing each production step in each firm.[29] We then aggregate across steps, weighing by the time intensity of the

[26] To do so, we reassign the time spent by all employees in a firm on non-production tasks to the entrepreneur. The counterfactual share of time in non-production tasks stays at 100% in firms with more than six workers.

[27] We also report suggestive evidence that more specialized firms are more productive.

[28] Since we sampled employees working on the core/main product, our sampling strategy is appropriate for studying labor specialization across steps within the core product, rather than specialization across products. However, we note that despite this sampling restriction, we still interviewed more than 50% of all employees in our sampled firms. This confirms that employee specialization across products is also limited.

[29] As explained in Section 2, for this we use information on which production steps individuals in the firm "usually" perform, rather than information from the time-use diary for the last day worked.

step, following a procedure similar to Figure 3, Panel (a). This creates a measure of the average share of employees performing a representative step. We do the same for entrepreneurs to create the share of entrepreneurs performing a production step.

Starting from employees (Panel (a)), we see that the share of employees working on a 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.

To better interpret this magnitude, we build an empirical benchmark corresponding to the share of employees that would work on a production step under full specialization.[30] We report this benchmark for firms of size 6 and 10. Comparing the actual allocation with the full specialization benchmark highlights that this type of specialization is indeed very limited and does *not* increase with firm size.

Panel (b) reveals a similar pattern for entrepreneurs: as we have shown that they 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 significant evidence of specialization increasing with firm size: the gap between the share of employees and entrepreneurs performing the typical step is constant.[31]

4.2.3 Interpreting the Evidence on Labor Specialization

To summarize, we find that labor specialization within the firm is limited overall, but specialization of entrepreneurs in non-production tasks is relatively more prevalent than specialization of workers within production across steps. We conclude this section by highlighting three key implications of these findings.

Specialization in Non-production Corresponds to Vertical Specialization. In Appendix A.3, we show that (i) non-production tasks are more complex, and (ii) entrepreneurs are more skilled than employees. This implies that the time allocation patterns in Figure 4, Panels (a) and (b), are evidence for "vertical" specialization between employees and entrepreneurs.[32]

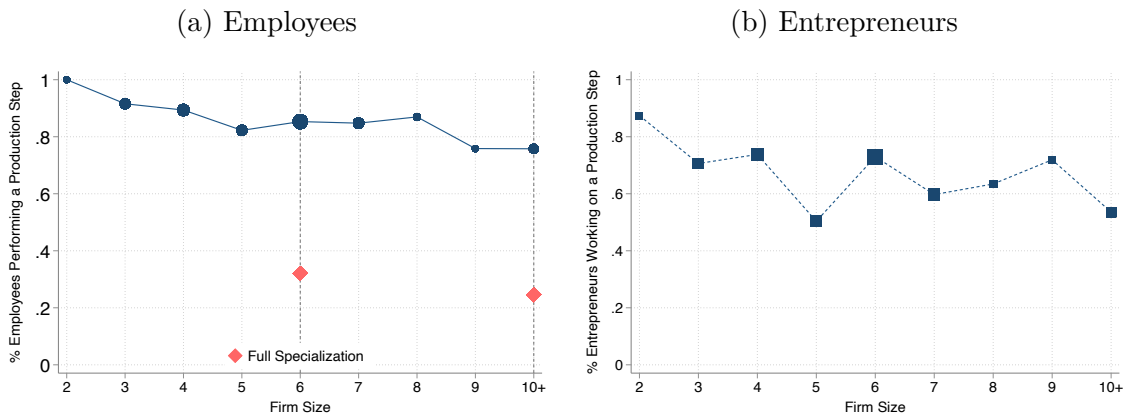
We do this because not all production steps for one product may be completed on the same day.

[30] To do so, we simply reassign employees across production steps to minimize the overlap in steps between employees while keeping the total amount of time worked by employees in production within the firm constant.

[31] In the Supplemental Appendix, we report this analysis separately for each step by sector.

[32] In Appendix A.3, we also show that entrepreneurs specialize in the more difficult steps within production, although again only to a limited extent.

Figure 6: Task Allocation Within Production Across the Size Distribution



Notes: Sample: all surveyed firms in the carpentry and welding sectors. Panels (a) and (b): share of employees and entrepreneurs (respectively) working on the representative production step (see main text for definition). The red diamond markers in Panel (a) represent the share of employees that would work on a production step under full specialization (reported for firms of size 6 and 10 only) (again, see main text for definition).

Is Labor Specialization an Important Reason Why Firms Exist? The evidence in this section suggests that the "Smithian", or horizontal, division of production tasks between individuals within the firm is not an important reason, in our context, for why firms exist: specialization across production steps is limited, both among employees and between employees and entrepreneurs, and does not vary with firm size. On the other hand, we find more evidence of vertical specialization between production and non-production tasks: larger firms allow entrepreneurs to leverage their talent by specializing in more complex non-production tasks, which is in line with the theoretical literature on hierarchies in organizations (Garicano and Rossi-Hansberg, 2006). This type of vertical specialization therefore seems to be a more important reason why individuals get together in a firm. Still, even this type of specialization is far from perfect. Using the model will allow us to more precisely quantify the limits to specialization, as well as the resulting implications for aggregate productivity and the size distribution of firms.

In Larger Firms, Most Non-production Tasks Are Done by Employees. A direct implication of the limited vertical specialization is that most non-production activities in larger firms are done by employees, not the entrepreneur. This is shown in Panels (c) and (d) of Figure 5. Panel (c) confirms that the share of firm-level time

in non-production tasks is constant across the size distribution at around 20% (in line with Figure 3, Panel (b)). However, Panel (d) shows that *who does* the non-production tasks varies dramatically with firm size: in firms with no employees or just one employee, naturally most of the non-production time in the firm is supplied by the entrepreneur. However, in larger firms, most of the non-production tasks are in fact done by employees: for instance, in firms of size eight, 70% of non-production activities are done by employees.

This result is striking: even though the entrepreneur is usually the most skilled individual in the firm and non-production tasks are more complex, most non-production tasks in larger firms are done by employees, not the entrepreneur.[33] It is a direct consequence of the fact that the entrepreneurs do not fully specialize in non-production activities, even as the firm grows large.

4.3 Heterogeneity Across Sectors

Finally, we study whether labor specialization varies across the three sectors.

In Table 2, we summarize the main statistics and relationships from Figures 3, 5, and 6, by sector. Panel A reports the average share of firms doing the typical production step, as well as its slope with firm size. While carpentry and welding are similar, we find evidence of higher specialization across firms in grain milling, where smaller firms specialize in fewer steps.

In Panel B, we look at specialization between production and non-production, and here we find a striking pattern. While carpentry and welding are almost identical, in grain milling there is significantly more specialization: the difference in the share of time spent in non-production activities between employees and the entrepreneur is twice as large in grain milling as in carpentry and welding, and the slope of the entrepreneur's specialization with firm size is also significantly larger. We find a similar pattern in Panel C, which looks at specialization within production across steps: the share of employees performing the typical step is lower in grain milling, and it decreases more steeply with firm size; the same is true when looking at entrepreneurs.[34]

[33] In line with this, in Appendix A.3 use a separate set of questions from the follow-up survey to show that in larger firms, employees play a more prominent role in customer generation and interaction, a key non-production activity.

[34] In Appendix A.3, we report the equivalent of Figures 3, 5, and 6 by sector. There we also show that grain milling has less idle time and that idle time decreases faster with firm size in this sector, which is again consistent with higher labor specialization (and so better time coordination).

Table 2: Heterogeneity in Labor Specialization by Sector

	Carpentry	Welding	Grain milling
	(1)	(2)	(3)
Panel A. Across Firm Specialization in Prod. Steps			
Avg. share of firms doing a step	0.870	0.932	0.816
Slope of share of firms doing a step with firm size	-0.001	-0.003	0.053
Panel B. Specialization in Production vs. Non-prod.			
Avg. diff. in entr. and empl. share of time in non-prod.	0.319	0.352	0.618
Slope of share of empl. time in non-prod. with firm size	0.004	-0.005	-0.001
Slope of share of entr. time in non-prod. with firm size	0.025	0.016	0.063
Panel C. Specialization within Production Steps			
Avg. share of employees performing a step	0.831	0.880	0.705
Slope of share of empl. performing a step with firm size	-0.022	-0.019	-0.064
Avg. share of entrepreneurs performing a step	0.686	0.647	0.306
Slope of share of entr. performing a step with firm size	-0.012	-0.024	-0.055

Notes: Sample: all surveyed firms and employees. The rows starting with "slope" report the OLS coefficient of a regression of the corresponding variable on firm size. See Figures 3, 5, and 6 for variable definition.

Implications. The sectoral heterogeneity is consistent with notion that the prevalence of product customization reduces the ability of entrepreneurs to specialize labor. The reason is that grain milling, which we have shown to have more standardized products, has substantially more specialization.

The cross-sectoral analysis is also helpful as it shows that specialization is, at least in principle, possible in this setting. It thus suggests that limited managerial skills or institutional features such as contract enforcement and lack of trust are not the only drivers of lack of specialization.[35] This heterogeneity also reassures us that the limited specialization in carpentry and welding is not simply due to measurement error, as the measurement of time use is the same across the three sectors.

[35]Since entrepreneurs in grain-milling are exposed to the same institutional setting and have similar managerial ability as those in carpentry and welding, the comparison across sectors is more consistent with the notion that product customization leads to limited specialization, than vice versa. In fact, for the causal link to go from specialization to customization, we would need to find some other reasons that limit labor specialization in grain milling in the first place.

5 Model: Organizing Artisanal Firms

We develop a model of within-firm specialization, optimal firm size, and occupational choice. We use the model to formalize the two-way relationship between labor specialization and firm size and to show how barriers to specialization affect firm-level and aggregate productivity.

5.1 Environment

We consider a static, closed economy with one sector, manufacturing. Each agent chooses whether to be a worker or start a firm. Firm owners choose employment as well as the allocation of workers to tasks, thereby determining firm productivity.

Agents and Demographics. The economy is populated by a measure 1 of agents who differ in their ability $z \in [0, z_{\max}]$. The distribution of ability in the population is given by $G(z)$. Each agent supplies one unit of labor and has linear utility over consumption.

Individuals can start a firm and become entrepreneurs (owners) o, or join the labor market as employees (workers) w. The endogenous distribution of ability in the two occupations is given by F_o and F_w with $G(z) = F_o(z) + F_w(z)$. Ability is not affected by an individual's occupational choice, but we refer to a generic firm owner's ability as $\hat{z}$. The equilibrium share of workers is $F_w(z_{\max})$, and we use $\mathcal{W}$ to denote the corresponding set of workers.

Firms produce a homogeneous good, so aggregate output in the economy is equal to the integral over individual firms' production:

$$Y = \left[\int Y(z) dF_o(z) \right] \tag{5.1}$$

We normalize the aggregate price index to 1.

Labor Market. There is a spot market for labor. The ability z of both entrepreneurs and workers is private information at the time of hiring. Therefore, there is a single labor market that randomly matches owners and employees. When a firm

owner chooses employment, she chooses only the mass of workers; the composition is determined by the equilibrium distribution $F_w(z)$.

Production Lines. The total output of a firm is equal to the sum of the output produced by all individuals in the firm: the owner $\hat{z}$ as well as the mass of $(n-1)$ workers she hired. By default, each individual is assigned one *production line*. In order for the production line to yield any output, a fixed set of tasks must be completed. Completing all tasks takes an individual their full unit of time. A fraction D of the tasks are complex; the remainder $1-D$ are simple. In the context of carpentry, simple tasks could be interpreted as basic production steps suck as drying the wood or sanding, whereas complex tasks would include negotiating with customers and suppliers or working on complex production steps such as designing or thicknessing.

Blueprints and Task Productivity. The value of the output from each production line is a function of two components. First, it depends on the ability of the firm owner, $\hat{z}$. This captures the quality of the "idea" (the blueprint) or also the reputation of the shop. The contribution of entrepreneurial talent is non-rival: the value of everybody's output benefits from the ability of the entrepreneur, irrespective of who does what within the firm.

Second, the value of the output from each production line depends on net task productivity, which is a measure of the average ability level with which the D complex tasks are performed. In the absence of labor specialization, this is simply equal to the individual z or $\hat{z}$ corresponding to the production line. With labor specialization, tasks within each production line can be traded. For instance, the entrepreneur could take over negotiating with customers while her employee cuts the wood for the entrepreneur's door. The employee's net task productivity would then be a function of the entrepreneur's ability as well.

Unbundling Cost and Assignment of Tasks. Trading tasks is costly. In order to assign parts of a production line to a different individual, tasks must be *unbundled*. For example, if the entrepreneur is the one negotiating with customers on all orders, she must then communicate exactly what the customer wants to the employee producing the order. The total cost of unbundling incurred by a firm depends on how many tasks in each production line are delegated to others.

Task assignment in the firm is summarized by μ, which consists of two functions. For all pairs of employees $\{z, z'\}$ in the firm, $\mu_c(z, z')$ and $\mu_s(z, z') \in \{\mathcal{W} \times \mathcal{W}\}$ 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 $\hat{z}$. For example, $\mu_c(z, \hat{z})$ measures 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 that an employee with ability z' performs.

Firm-Level Output. Formally, the output of a firm of size n, owned by an individual with ability $\hat{z}$ who chooses task assignment μ, is given by Equation (5.2). This general setup is heavy on notation. In the model characterization section, we describe a simpler, and empirically relevant, case that provides all the economic intuition.

$$
\begin{aligned}
Y(\hat{z}, n, \mu) &= y(\hat{z}, \hat{z}, \mu) + (n-1) \int y(z, \hat{z}, \mu) \frac{dF_w(z)}{F_w(z_{\max})} \qquad (5.2)\\
&\text{where,} \quad \forall z \in \{\mathcal{W} \cup \hat{z}\}\\
y(z, \hat{z}, \mu) &= \hat{z}^{\lambda} \quad \tilde{z}(z, \hat{z}, \mu)^{1-\lambda} \quad \mathbb{I}_{y(z,\hat{z},\mu)}\\
\tilde{z}(z, \hat{z}, \mu) &= \exp\left\{ \mu_C(z, \hat{z}) \log(\hat{z}) + (n-1) \int \mu_C(z, z') \log(z') \frac{dF_w(z')}{F_w(z_{\max})} \right\} (1 - \kappa(\mu_C(z, z)))\\
\mathbb{I}_{y(z,\hat{z},\mu)} &= \mathbb{I}_{[\mu_C(z,\hat{z}) + (n-1)\int \mu_C(z,z') \frac{dF_w(z')}{Fw(z_{\max})} \geq 1]} \quad \mathbb{I}_{[\mu_S(z,\hat{z}) + (n-1)\int \mu_S(z,z') \frac{dF_w(z')}{Fw(z_{\max})} \geq 1]}
\end{aligned}
$$

where λ governs the relative importance of the non-rival entrepreneurial idea in firm productivity and $\kappa(.)$ is the unbundling cost. Note that the overall weight of the ability $\hat{z}$ of the entrepreneur in the firm productivity increases in λ, and if $\lambda = 0$ there is perfect symmetry between the owner and her workers in the production function. The indicator function $\mathbb{I}_{y(z,\hat{z},\mu)}$ enforces that all tasks, single and complex, are completed on each production line.

In order for an assignment be *feasible*, no individual in the firm can spend more than their one unit of time across all tasks. Formally, an assignment μ is feasible if

and only if

$$\forall z \in \{\mathcal{W} \cup \hat{z}\}: \quad D\mu_C(\hat{z}, z) + (1-D)\mu_S(\hat{z}, z) + (n-1)\int \left(D\mu_C(z', z) + (1-D)\mu_S(z', z)\right) \frac{dF_w(z')}{F_w(z_{\max})} \leq 1 \tag{5.3}$$

5.2 Choices

We next describe the choices of economic agents in this model. All individuals choose whether to be workers or start their own firms. Entrepreneurs also choose 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 μ to maximize profits:

$$\pi(\hat{z}) = \max_{\{\mu, n \geq 1\}} Y(\hat{z}, n, \mu) - (n-1)\int w(z, \hat{z}, \mu)\frac{dF_w(z)}{F_w(z_{\max})} - \chi(n) \tag{5.4}$$
$$\text{s.t.}(5.2), (5.3)$$

The cost $\chi(n)$ captures all other costs incurred by the firm. These include hiring costs, capital expenditures, credit frictions, and any other auxiliary costs or frictions that scale with firm size. For brevity, we refer to $\chi(n)$ as the hiring cost. Since firm size is at least equal to 1—the owner herself, $\chi(1)$ corresponds to the fixed cost of setting up a firm.

Wages. As soon as workers and entrepreneurs match, their ability is publicly observed. The wage is determined by a standard Nash bargaining protocol.

A worker's outside option is equal to $\overline{w}$. The level of $\overline{w}$ is endogenous and adjusts to clear the labor market. The owner's outside option is her profits when producing with one fewer worker. The surplus of the match is therefore a function of worker ability z as well as the entrepreneur's ability $\hat{z}$ and the assignment of tasks in the firm μ: $S(z, \hat{z}, \mu) = y(z, \hat{z}, \mu) - \overline{w}$. The hiring cost $\chi(n)$ is sunk at the time of bargaining and is therefore not directly included in the surplus.

The worker has bargaining power ω. His wage is given by

$$w(z, \hat{z}, \mu) = (1 - \omega)\overline{w} + \omega y(z, \hat{z}, \mu) \tag{5.5}$$

Occupational Choice. Each agent observes their ability z and chooses whether to be a worker or an entrepreneur. Workers make no further choices since there is random matching in the labor market and tasks within the firm are assigned by entrepreneurs.[36] Profits conditional on entering are known, since firm owners hire a representative sample of workers. Wage earnings, on the other hand, depend on who the worker happens to match with. An individual with ability z therefore starts a firm if and only if her profits are higher than her expected wage in the labor market :

$$\mathbb{I}_o(z) = 1 \quad \Longleftrightarrow \quad \pi(z) \geq \int w(z, \hat{z}, \mu)\,(n(\hat{z}) - 1)\frac{dF_o(\hat{z})}{\int (n(\hat{z}) - 1)\,dF_o(\hat{z})} \tag{5.6}$$

5.3 Equilibrium

Finally, we define an equilibrium in our setting, which simply requires that all agents maximize and that the wage level clears the labor market; that is, the total labor demand of entrepreneurs is equal to the mass of individuals choosing not to start a firm.

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

1. *firm owners choose size and task assignment to maximize profits as in (5.4);*
2. *individuals choose their occupation according to (5.6);*
3. *the labor market clears: $\int (n(z) - 1) dF_o(z) = \int dF_w(z)$;*
4. *$\frac{F_o(z)}{F_o(z_{max})}$, $\frac{F_w(z)}{F_w(z_{max})}$ are consistent with the occupational choice—that is, $F_w(z) = \int (1 - \mathbb{I}_o(z)) dG(z)$ and $F_o(z) = \int \mathbb{I}_o(z) dG(z)$*

[36] Task assignment maximizes surplus, so there is no disagreement between workers and the entrepreneur.

5.4 Characterization

In this section, we analyze how the costs and benefits of specialization (i.e., $\kappa(.)$ and λ) affect the allocation of talent within firms and how it shapes firm size, productivity, and ultimately the allocation of talent between occupations. 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 properties of the economy's equilibrium. All proofs are in Appendix B.

Assumptions. Throughout this section, we work under three assumptions. The assumptions restrict model parameters to the empirically relevant case and furthermore simplify the assignment problem within the firm in order to clarify the key economics.

ASSUMPTION **1.** The unbundling cost is given by $\kappa(x) = 1 - \exp\{-\hat{\kappa}(x)\}$, where $\hat{\kappa}(x) = {\kappa_0}^{1/\kappa_1} \frac{(1-x)^{1+1/\kappa_1}}{(1+1/\kappa_1)}$, and where x is the share of complex tasks that are not delegated: $x = \mu_C(z, z)$. The hiring cost is given by $\chi(n) = \chi_0^{1/\chi_1} n^{1+1/\chi_1} (1+1/\chi_1)^{-1}$.

ASSUMPTION **2.** Each entrepreneur spends, in equilibrium, at least some time on simple tasks: $\mu_S(\hat{z}, \hat{z}) + (n-1) \int \mu_S(z', \hat{z}) dF_w(z') > 0 \ \forall \hat{z}$.

ASSUMPTION **3.** Workers' bargaining weight ω satisfies $\omega \leq \left(\frac{\partial(\max_\mu y(\hat{z}, z, \mu))}{\partial z} \right)^{-1} \forall \{z, \hat{z}\}$.

Assumption 1 allows for closed-form solutions and to parameterize unbundling and hiring costs by two key parameters, κ_0 and χ_0. Assumption 2 is motivated by the empirical evidence showing that even in the largest firms, the entrepreneur spends some of her time on simple tasks. It is a joint assumption on the parameters of the model, which simplifies the exposition in this section, but will be relaxed in the quantitative analysis.[37] Assumption 3 is also motivated by the data, where we show that there is relatively little variation in worker compensation.[38] Further, we find that entrepreneurs are positively selected, which is guaranteed in the model by Assumption 3.

Occupational Choice. The model yields a familiar sorting of talent into occupations as a function of their skill sensitivity.

[37] In the estimated model in the next section, this assumption will hold at our estimates, but not necessarily in the counterfactuals.

[38] A low ω translates into little *variation* in wages as a function of skill. Since the level of wages $\overline{w}$ is an equilibrium object, assumptions on ω do not directly translate into the *level* of wages.

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 $z \geq z_0$.*

The higher an individual's ability, the larger the returns from starting a firm relative to remaining a worker. While earnings in both occupations are increasing in individual ability z, the relationship is steeper for entrepreneurs. For them, ability affects not only their own output but also that of their workers (as long as $\lambda > 0$ and $\kappa_0 < \infty$). For workers instead, wages do not increase one for one with ability, since they are a function of a fixed equilibrium component, their own ability, and the one of the entrepreneur they match with.

Lemma 1 guarantees that the entrepreneur is the most skilled individual in each firm. This result simplifies the within-firm allocation of talent, which we analyze next.

Production-line output. When the owner is more skilled than her employees and is not time constrained (Assumption 2), she is the only one who takes on any complex tasks of others. To see this, fix a share of complex tasks that a worker z delegates, $1 - \mu_c(z, z)$. Given that share, it is always best to assign complex tasks to the most skilled individual, since net task productivity $\tilde{z}(z, \hat{z}, \mu)$ is increasing in the ability of the person performing complex tasks, and the unbundling cost does not depend on to whom tasks are delegated.[39] The assignment problem therefore reduces to choosing $\mu(z) \equiv \mu_c(z, z)$, the share of each worker's complex tasks she performs herself. The remainder is done by the owner, who also performs all of her own complex tasks.

Suppressing the indicator function for simplicity, we can rewrite the output of each production line as

$$y(z, \hat{z}, \mu) = \underbrace{\hat{z}^{\lambda}}_{\text{non-rival}} \left(\underbrace{z^{\mu(z)} \hat{z}^{1-\mu(z)}}_{\text{task productivity}} \quad \underbrace{[1 - \kappa(\mu(z))]}_{\text{unbundling cost}} \right)^{1-\lambda} \tag{5.7}$$

The output produced by a worker z is a geometric average, weighted by the parameter λ, of the ability of the firm owner $\hat{z}$ and the net task productivity. Net task productivity is itself a geometric average of the worker's and the owner's ability. The weight on the owner is equal to the share of complex tasks she takes over from this worker, and thus it depends on the allocation of tasks within the firm.

[39] We made this assumption to gain analytical tractability.

Artisanality. Equation 5.7 showcases the two parameters that modulate the extent of *artisanality* in our model: the non-rival component of talent λ and the unbundling cost κ. Both parameters shape the ability of skilled entrepreneurs to leverage their talent by passing it through to their workers. If λ is low, the productivity of the production line is mostly determined by the ability of the person performing the complex tasks. As a result, assigning the complex tasks to the entrepreneur is necessary to increase firm productivity. If, in addition, the unbundling cost κ is high, such labor specialization is costly and the entrepreneur would choose $\mu(z)$ close to 1. Overall, both λ and κ determine the firm productivity and the return to managerial ability, but κ additionally shapes the internal organization of the firm.

We interpret artisanal manufacturing as a setting in which λ is low, making vertical labor specialization important, and at the same time κ is high, making such specialization costly. In contrast, modern manufacturing, and scalability of talent, could be achieved either through effective specialization (low κ) or through other arrangements that effectively commodify labor, making the quality of the output less dependent on the ability of each individual worker (high λ).

Labor Specialization. For each worker, the optimal share of complex tasks delegated to the entrepreneur equates the marginal benefit—the difference in abilities between worker and entrepreneur—to the marginal unbundling cost.

$$\underbrace{\log \hat{z} - \log z}_{\text{marginal benefit of assigning complex task to } \hat{z}} = \underbrace{\kappa_0^{1/\kappa_1} \mu(z)^{1/\kappa_1}}_{\text{marginal cost of delegating complex tasks}} \tag{5.8}$$

The complex pairwise assignment problem therefore has a simple solution, which highlights the specific nature of labor specialization in our model. Consistent with the empirical evidence presented in Section 4, specialization happens along the vertical dimension: the most skilled individual in the firm is the one who specializes in the complex tasks.[40] The level of specialization in the firm is governed by κ_0, while the curvature κ_1 modulates the extent to which delegation is a function of the ability gap between worker and owner.

[40] Our model does not include a notion of horizontal specialization. We decided to focus on the vertical specialization for two reasons: (i) we found stronger evidence for it, and (ii) as we describe in the next section, we can (relying on the structure of the model) measure the gains from vertical specialization. It would, instead, be very hard to measure the gains from horizontal specialization.

In order to directly map model and data, we now characterize the equilibrium time use of entrepreneurs and workers.

DEFINITION **1** (Average Labor Specialization)**.** Let the total time spent on complex tasks by the entrepreneur $\hat{z}$ be $\hat{\theta}(\hat{z}) \equiv D\left(\mu_C(\hat{z},\hat{z}) + \frac{(n-1)}{F_w(z_{max})}\int \mu_C(\hat{z},z')dF_w(z')\right)$ and that of one of her workers z be $\theta(z,\hat{z}) \equiv D\left(\mu_C(z,\hat{z}) + \frac{(n-1)}{F_w(z_{max})}\int \mu_C(z,z')dF_w(z')\right)$. We define the *average labor specialization* in the firm, $\overline{\theta}(\hat{z})$, to be the difference between the time spent on complex tasks by the entrepreneur and her average employee:

$$\overline{\theta}(\hat{z}) \equiv \hat{\theta}(\hat{z}) - \int \theta(z',\hat{z})\frac{dF_w(z')}{F_w(z_{max})}.$$

Equipped with this definition, Lemma 2 formalizes the degree of labor specialization and its relationship with firm size.

LEMMA **2** (Labor Specialization)**.** *Consider a firm of size n.*

1. *The time spent on complex tasks by a worker with ability z and by the entrepreneur $\hat{z}$ is equal to*

$$\theta(z,\hat{z}) = D\left(1 - \frac{1}{\kappa_0}\left(\log\hat{z} - \log z\right)^{\kappa_1}\right) \tag{5.9}$$

$$\hat{\theta}(\hat{z}) = D\left(1 + \frac{n-1}{\kappa_0}\int \left(\log\hat{z} - \log z\right)^{\kappa_1}\frac{dF(z)}{F_w(z_{max})}\right). \tag{5.10}$$

2. *Average labor specialization in the firm is equal to*

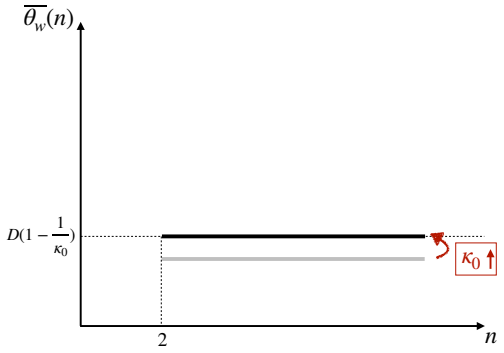
$$\overline{\theta}(\hat{z}) = D\left(\frac{n}{\kappa_0}\int \left(\log\hat{z} - \log z\right)^{\kappa_1}\frac{dF(z)}{F_w(z_{max})}\right) \tag{5.11}$$

It is declining in the unbundling cost κ_0 and increasing in firm size n at a rate that decreases in κ_0.

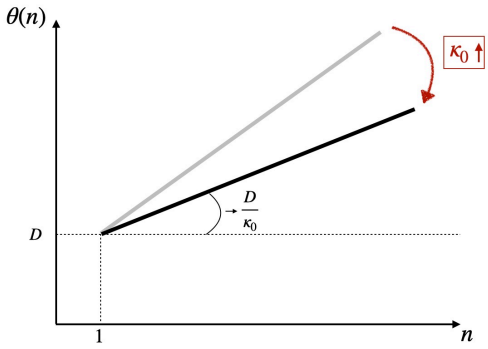
Figure 7 illustrates the relationship between labor specialization, firm size, and the unbundling cost κ_0. For ease of exposition, we set $\kappa_1 \to 0$. Under this parameterization, the share of complex tasks each worker delegates to the entrepreneur is independent of the ability gap. Both panels of Figure 7 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).

Figure 7: Labor Specialization, Firm Size, and Unbundling Cost κ_0

(a) Workers Time in Complex Tasks



(b) Entrepreneurs Time in Complex Tasks



The share of time each worker spends on complex tasks is independent of firm size n. As long as the entrepreneur has capacity left to take on complex tasks, which is guaranteed by Assumption 2, 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 there are for the entrepreneur to take on.

An increase in the unbundling cost ($\kappa_0 \uparrow$) leads each worker to delegate fewer tasks to the entrepreneur and spend more time on complex tasks. For entrepreneurs, a higher unbundling cost affects the *slope* of complex time share with size, as Figure 7 illustrates. The intuition for this is straightforward. In a firm of size one, the unbundling cost of course has no impact on the time the entrepreneur spends on complex task, since she is the sole worker. But each employee she hires delegates fewer tasks to the entrepreneur, and hence her share of time spent on complex tasks grows more slowly in 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 $Y(\hat{z}, n) = \mathbb{Z}(\hat{z}, n, \mu)\, n$ where*

$$\mathbb{Z}(\hat{z}, n, \mu) = \underbrace{\hat{z}^{\lambda}}_{\text{non-rival}} \quad \underbrace{\tilde{Z}(\hat{z}, n, \mu)^{1-\lambda}}_{\text{aggregate task-productivity}}$$

$$\tilde{Z}(\hat{z}, n, \mu) = \left(\frac{1}{n} \hat{z}^{1-\lambda} + \underbrace{\frac{n-1}{n} \int \tilde{z}(z, \hat{z}, \mu)^{1-\lambda} \frac{dF_w(z)}{F_w(z_{max})}}_{\text{dilution from firm size}} \right)^{\frac{1}{1-\lambda}},$$

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

and $\mathbb{Z}(\hat{z}, n, \mu)$ *is strictly declining in* n *as long as both* $\lambda < 1$ *and* $\kappa_0 > 0$.

Lemma 3 isolates the two components of firm productivity. The first term, $\hat{z}^{\lambda}$, captures the unique role of the entrepreneur. As long as $\lambda < 1$, the entrepreneur's ability passes through to firm productivity in a non-rival way, that is, independently of time use within the firm.

The second term, firm-level task-productivity $\tilde{Z}(\hat{z}, n, \mu)^{1-\lambda}$, is a weighted average of the ability of all individuals completing the complex tasks. Since the entrepreneur completes all her own tasks, her individual task productivity is $\hat{z}$. The task productivity of each employee, however, is less than the entrepreneur's ability, as long as there is less than full specialization ($\kappa_0 > 0$). The entrepreneur's and her workers' task productivity is then aggregated to the firm level, with a weight of $1/n$ on the entrepreneur and $(n-1)/n$ on workers. Increasing the size of the firm would therefore decrease its productivity: more weight is given to the lower task productivity of workers. This mechanism generates decreasing returns to scale originating from limited specialization within the firm.

Optimal Firm Size. Labor specialization and firm size are closely intertwined. Lemma 2 showed one side of this two-way causal relationship: small average firm size limits labor specialization. Lemma 4 shows that the opposite is also true: barriers to labor specialization reduce the optimal firm size.

Lemma **4** (Firm Size)**.** *The optimal firm size* n *of each entrepreneur* $\hat{z}$ *solves*

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

It is declining in the marginal hiring cost $\chi'(n)$ and, as long as $\lambda < 1$, in the unbundling cost κ_0.

Profit maximization implies that at the optimal size, the marginal cost of hiring an additional worker is equal to marginal revenues. The marginal cost is equal to 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, firm-level productivity is decreasing in employment, since each additional worker is less skilled than the entrepreneur (as long as $\lambda < 1$ and $\kappa_0 > 0$). In choosing firm size, the entrepreneur takes into account the decreasing returns to scale arising from productivity dilution.

Overall, two kinds of frictions can keep firms small. The hiring cost $\chi(n)$ directly constrains optimal firm size by making expansion costly. The unbundling cost κ_0 reduces firm size both through production dilution and because it reduces average productivity for any firm size, as Lemma 3 shows.

There is also an apparent complementarity between the frictions. The benefits from relaxing the external wedge $\chi'(n)$ are limited if internal barriers to labor specialization hinder firm productivity and generate strong decreasing returns to scale.

Why Do Firms Exist? Two Polar Cases. To complete the intuition behind our model of a firm, Lemma 5 outlines two polar cases that span the space from the lowest to the highest degree of artisanality in production.

LEMMA **5** (Self-Employment within the Firm)**.** *Depending on the weight of non-rival entrepreneurial talent in production (λ) and the size of the unbundling cost (κ_0) the model spans two polar types of firms:*

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

2. ***Self-Employment within the Firm.*** *If $\lambda = 0$ and $\kappa_0 \to \infty$, then $Y(\hat{z}, n, \mu) = \overline{z}(\hat{z})\, n$, 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 ($\kappa_0 = 0$) or entrepreneurial talent is entirely non-rival ($\lambda = 1$), firm productivity is equal to the entrepreneur's 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, there are high returns to entrepreneurial ability and hence talented firm owners optimally run large firms.

In the opposite extreme, delegation is prohibitively costly ($\kappa_0 \longrightarrow \infty$). All individuals in the firm behave as if they were self-employed, completing all the tasks required for their production line. If in addition talent is purely rival—that is, productivity is purely a function of the ability of the individual performing the complex tasks and $\lambda = 0$–then firm productivity firm is simply the average ability of all its workers. In this benchmark, the only reason why firms exist is to potentially share fixed costs. All firms are identical in size, and entrepreneurs have no means to leverage their ability.

Overall, Lemma 5 shows that, in our model, the notion of what a firm is crucially depends on the parameters κ_0 and λ, hence on the extent of artisanality and the resulting internal organization of the firm.

Equilibrium and Aggregate Implications. So far, we have considered the solution to the problem of one entrepreneur. Next, we turn to the overall economy. We prove the main proposition for the case of $\kappa_1 = 0$. This allows us to highlight the main economics without further assumptions on the population distribution of talent, $G(z)$. The estimated model indeed confirms that κ_1 is small in our environment.

PROPOSITION **1** (Aggregate Effects of the Unbundling Cost κ_0)**.** *Suppose that $\lambda < 1$ and $\kappa_1 = 0$. 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. *average labor specialization $\overline{\theta}(n)$ in all firm sizes;*
2. *the slope of the relationship between average labor specialization and firm size;*
3. *the average ability of firm owners: z_0 increases;*
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 $\overline{\mathbb{Z}} \equiv \int \mathbb{Z}(z, n(z), \mu(z)) n(z) \frac{dF_o(z)}{F_o(z_{max})}$;*
6. *the wage $w(z, \hat{z}, \mu)$ of all workers z in all firms $\hat{z}$.*

Proposition 1 summarizes the effect of a decline in the unbundling cost κ_0 on the economy. When the cost is high, the economy is made of many small firms, owned by low-productivity managers and internally organized with limited specialization. The result is low aggregate productivity, low demand for workers, and consequently low wages. The returns to managerial ability are limited as firm owners are not able to leverage their talent.

Reducing the delegation 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.

6 Estimation: Bringing the Model to the Data

We now bring our model to the data. As a first step, we use heterogeneity across sectors and regions to validate the theoretical predictions from Section 5. Then, we parameterize the model to make it amenable to a quantitative analysis. Finally, we discuss identification and the results from the estimation.

6.1 Empirical Validation of the Theoretical Predictions

We provide two qualitative tests to support the model predictions of Section 5.4.

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 to provide supporting evidence. As discussed in Section 3, the degree of standardization is remarkably similar in carpentry and welding, but it is larger in grain milling, suggesting a lower κ_0 in that sector.

In Table 3, we show that the key predictions of Proposition 1 hold across sectors. Carpentry and welding are almost identical in terms of specialization within firms, average size, returns to managerial ability, and selection into entrepreneurship. In

grain milling, on the other hand, there is more labor specialization, firms are larger, and the returns from managerial ability as well as the skill gap between entrepreneurs and their employees are larger.[41]

Table 3: Cross-Sectoral Heterogeneity

	Carpentry	Welding	Grain milling
	(1)	(2)	(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: the gap in the average share of time in non-production tasks between entrepreneurs and employees. Panel B: Sample: all firms. We report coefficients from regression of three dependent variables on the (standardized) index of managerial ability. Panel C: Sample: all interviewed entrepreneurs and employees. We report 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.

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 operating at the regional level. 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.[42] We calculate

[41] 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.

[42] We restrict our focus to carpentry and welding since we have shown in Table 3 that grain

the average firm size in each sub-county, rank them based on this statistic, and then divide the sub-counties in two groups with roughly equal numbers of firms.[43] Finally, within each group of sub-counties, we estimate the return to managerial ability by regressing log revenues on sector dummies and either the previously used managerial ability index or the years of education of the entrepreneurs.

The results are shown in Table 4. Consistent with Lemma 3, we find higher returns to managerial ability within the set of sub-counties with the smallest firms.

Table 4: 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 with carpentry and welding firms (grain milling excluded). Robust standard errors are in parentheses. Regressions are weighted by sampling weights.

6.2 Model Parameterization

We extend the model laid out in Section 5 along three dimensions. First, in order to account for the heterogeneity in firm sizes observed in the data, we allow entrepreneurs to differ not only in their managerial ability z but also in the hiring cost χ_0. We assume that z is drawn from a generalized Pareto distribution with scale and location normalized to one and shape given by σ_z. Further, χ_0 follows a normal distribution with mean $\overline{\chi_0}$ and standard deviation σ_χ.[44]

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.

[43]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.

[44]Given the normalization of scale and location, the variance of productivity is monotonically increasing in the shape parameter σ_z.

Second, to precisely match the time use within firms, we allow for an overhead amount d of non-production tasks at the firm level. This overhead time must be supplied by the entrepreneur and does not affect productivity.[45]

Third, we make a small modification to the functional form of the hiring cost. We assume that the hiring cost has to be paid only on hired labor $(n-1)$ but also assume that the entrepreneur has to pay a fixed cost to operate a firm χ_f. The overall "hiring" cost for a firm of size n is therefore $\chi(n) = \chi_f + \chi_0^{1/\chi_1} (n-1)^{1+1/\chi_1} (1+1/\chi_1)^{-1}$. We choose this functional form to separate the cost of starting a firm from the true hiring cost, providing additional flexibility. For simplicity, we will still refer to the composite cost $\chi(n)$ as the hiring cost.

We also assume that our empirical measure of the managerial score discussed in Section 4 is a noisy proxy of the underlying managerial ability, denoted $s(z)$. Specifically, we assume that the (normalized) managerial score is equal to the (normalized) log of managerial ability plus an additive, normally distributed term.[46]

Table 5 summarizes the full economic environment that we take to the data and links each economic block to the main parameters modulating it.

6.3 Targeted Moments and Identification

Given the parameterization of the model, we are left with 12 parameters to be pinned down. Our survey was uniquely designed to measure firms' start-up and fixed operating costs. We can thus use it to directly calibrate χ_f.[47] The remaining 11 parameters do not have straightforward empirical counterparts and are jointly estimated.

Targeted moments. We target 150 moments, based on pooled data for carpentry and welding. Table 6 lists 21 summary moments that capture the main economic concepts we are targeting. We leave the full list of moments as well as details on how each one is calculated to Appendix C.1.[48]

[45] For this reason, we do not include d when calculating the measures of firm-level specialization in the counterfactuals, which we now define as $\overline{\theta}(\hat{z}) - d$, where $\overline{\theta}(\hat{z})$ is as previously defined.

[46] This last assumption is particularly important since it allows us to accommodate enough heterogeneity in managerial ability to match the empirical distribution of log revenues while matching the observed empirical relationship between firm revenue, workers' earnings, and managerial ability index.

[47] See Appendix C.1.5 for details.

[48] For example, while we target the deciles of the distributions of firm sizes and revenue, we include in the Table 6 only their means and standard deviations.

Table 5: Summary of the Economic Environment and Parameters

	Equation	Parameters
Final Output	$Y = \int Y(z) dF_o(z)$	
Firm Output	$Y(\hat{z}, n) = \mathbb{Z}(\hat{z}, n, \mu)\, n$	
Firm Productivity	$\mathbb{Z}(\hat{z}, n, \mu) = \hat{z}^{\lambda} \left(\frac{1}{n} \hat{z}^{1-\lambda} + \frac{n-1}{n} \int \tilde{z}(z, \hat{z}, \mu)^{1-\lambda} dF_w(z) \right)^{\frac{1}{1-\lambda}}$	λ
Net Task Productivity	$\tilde{z}(z, \hat{z}, \mu) = z^{\mu(z)} \hat{z}^{1-\mu(z)} \left[1 - \kappa \left(\mu(z) \right) \right]$, $\mu(z) \equiv \frac{\theta(z, \hat{z})}{D}$	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	$\kappa(x) = 1 - \exp\left\{ -{\kappa_0}^{1/\kappa_1} \frac{x^{1+1/\kappa_1}}{(1+1/\kappa_1)} \right\}$	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	$w(z, \hat{z}, \mu) = (1 - \omega)\overline{w} + \omega \hat{z}^{\lambda} \tilde{z}(z, \hat{z}, \mu)^{1-\lambda}$	ω
Measurement Error	Observe $s(z) = \log z + \epsilon$, $\epsilon \sim N(0, \sigma_\epsilon)$	σ_ϵ
Complex Share (Workers)	$\theta(z, \hat{z}) = D \left(1 - \frac{1}{\kappa_0} \left(\log \hat{z} - \log z \right)^{\kappa_1} \right)$	D, κ_0, κ_1
Complex Share (Entrepr.)	$\hat{\theta}(\hat{z}) = d + D \left(1 + \frac{n-1}{\kappa_0 F_w(z_{max})} \int \left(\log \hat{z} - \log z \right)^{\kappa_1} dF_w(z) \right)$	d, D, κ_0, κ_1

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 3 and 4. 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.[49] 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 managerial ability. Workers' earnings, in fact, convey useful information on the productivity of each individual worker, which

[49] 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.

unfortunately we do not observe in the data.

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 simple routine that we developed in our previous work (Bassi et al., 2022b). Details are included in Appendix C.2, where we also show that the parameters are well-identified: the likelihood function is single-peaked around the estimated parameters, and we verify that our estimation procedure recovers the true parameters when we run the estimation on a synthetic set of moments generated by the model itself.[50]

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 Appendix C.2, but in the last column of Table 6, we include the *key parameters* that are linked to each moment.[51]

As Lemma 2 highlights, the within-firm allocation of time is tightly linked to the share of complex tasks in production (D), the overhead time (d), and the parameters of the unbundling cost (κ_0, κ_1). Importantly, our unique data on the slope of 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 biggest identification challenge is to pin down the degree of the non-rivalry of entrepreneurial talent λ. To identify this key parameter, we use the fact that, conditional on labor specialization, λ modulates the pass-through of entrepreneurial to worker productivity. When λ is large, workers inherit the ability of their entrepreneurs, and there is a lot of heterogeneity in worker productivity across firms

[50] To be precise, we draw random vectors of parameters around the estimates. We then compute moments from those vectors and show that if we target the model-generated moments, our routine recovers the true set of parameters.

[51] In Appendix C.2, we explain in detail how we compute the Jacobian and select the *key parameters*. A parameter is deemed key for a moment if its effect on that specific moment is at least twice as large as the average effect of all other parameters.

but little within-firm heterogeneity.[52] In our data, we do not directly observe worker productivity. We can, nonetheless, rely on a key feature of our setting: workers' earnings are an increasing function of their productivity, modulated by workers' bargaining weight ω (see equation 5.5).

The second block of moments therefore includes several statistics on workers' earnings, which allow us to separately pin down 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. In practice, since earnings are likely 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, normalized by their standard deviation, across more versus less productive firms.

The third block of moments includes the distribution of firm revenue and its relationship with the measured managerial score. These moments discipline the variance of managerial talent (σ_z) and the noise term (σ_ϵ) in our empirical proxy. As shown by Lemma 3, a large σ_z increases the variance of firm productivity and hence the variance of revenues. Given σ_z, a large σ_ϵ would then flatten the relationship between revenue per worker and managerial score through a standard attenuation bias.

Finally, the last block of moments help us to pin down the parameters of the hiring cost: χ_0, χ_1, and σ_χ. In fact, Lemma 4 shows that, given the firm productivity and the average wage, these parameters directly map to the firm size distribution and its relationship with managerial ability.

Importance of Time-Use Data. The identification argument highlights why it is crucial to collect time use data within the firm. Without this data, it would be impossible to identify the barriers to labor specialization within the firm and separate them from any other constraint that keeps firm small.

Even observing aggregate measures of labor specialization would not be enough. For example, assume that we observe firm size, but we have information only on average labor specialization across all firms. It would be impossible to distinguish whether firms are small because they are not specialized or whether they are not specialized because their small size does not make it worthwhile to do so. Instead,

[52] It is important to emphasize that a low κ_0 has the same effect of a high λ. Therefore, our identification strategy unfortunately relies on the assumed relationship between time spent on complex tasks and productivity (i.e., on $\mu(z)$).

our unique data, by showing that specialization increases only weakly with firm size, allows us to pin down directly the barriers to specialization (κ_0). Then, given κ_0 and the other parameters, the structure of the model predicts a firm size distribution, and comparing the predicted distribution with the observed one pins down all other potential barriers to firm growth (χ_0).

Table 6: 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	0.229	D
(ii) Average of Entrepreneurs	0.457	0.447	d, D
(iii) Average for Self-Employed	0.341	0.341	d, D
(iv) Average of Low-Skilled Workers	0.137	0.173	D
(v) Average of High-Skilled Workers	0.217	0.177	D
(vi) Slope w/ Size (Entr)	0.021	0.021	κ_0, κ_1
(vii) Slope w/ Size (Low-Skilled Work)	0.002	-0.001	κ_1, ω, χ_1
(viii) Slope w/ Size (High-Skilled Work)	0	-0.001	κ_1, ω, χ_1
(ix) Slope w/ Log(Earn)	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	0.196	$\sigma_\epsilon, \lambda, \omega$
(ii) Log(Earn) on Log(Rev p.w.)	0.191	0.196	$\lambda, \omega, \sigma_\chi$
(iii) Norm. Earn Gap by Rev p.w.	0.389	0.718	$\lambda, \chi_1, \overline{\chi_0}$
(iv) Norm. Earn Gap by Man. Ability	0.137	0.327	$\omega, \chi_1, \sigma_\epsilon$
C. Distribution of Firm Revenues			
(i) Std of Log(Rev)	0.726	0.637	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	0.385	$\sigma_\epsilon, \omega, \overline{\chi_0}, \sigma_z$
D. Firm Size Distribution			
(i) Average Size	5.701	5.894	$\omega, \chi_1, \overline{\chi_0}$
(ii) Std of Log(Size)	0.489	0.586	$\sigma_\chi, \overline{\chi_0}$
(iii) Std of Size	2.263	2.537	σ_χ, χ_1
(iv) Log(Size) on Man. Ability	0.1	0.089	$\chi_1, \omega, \sigma_\epsilon$
(v) Size Gap by Man. Ability	0.275	0.339	$\omega, \chi_1, \sigma_\epsilon$

Notes: The table shows the empirical moments used in estimation and the 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.1.

6.4 Estimation Results and Model Fit

Despite being relatively parsimonious, the model matches the data well, as Table 6 shows. Figure 8 illustrates the fit for some of the key moments: the model matches

Table 7: List of Parameters and their Estimated Values

Param.	Value	Param.	Value	Param.	Value
χ_f	$0.1\overline{\pi}(z)$	κ_0	0.078^{-1}	κ_1	0.684
λ	0.223	ω	0.373	$\overline{\chi_0}$	14.608
χ_1	0.47	σ_χ	4.371	σ_z	0.966
D	0.189	d	0.181	σ_ϵ	2.28

the heterogeneity between firms in terms of size and revenue, as well as the time allocation within firms. Crucially, the model accounts for the key relationships between specialization and firm size shown in Figure 5.[53]

Table 7 includes the estimated values of all parameters. A few are easy to interpret and worthwhile to discuss. First, the value of λ implies that the entrepreneur is able to pass through approximately 22% of her ability to her workers.[54] 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 more than 90% of her complex tasks ($\mu(z) \approx 0.92$). Given the value of λ, this means that the productivity pass-through due to vertical specialization is $\sim 6\%$—that is, about 30% of the direct pass-through due to λ.

Second, the value of ω shows a prominent role for piece-rate, consistent with the evidence described in Section 3: more productive workers are compensated for around 40% of their higher output.[55]

Third, we estimate a very large heterogeneity in managerial ability at the top of the distribution. The estimated value of the shape parameter σ_z implies that the managerial ability of the 98^{th} percentiles is approximately 9 times that of the 80^{th} percentiles.[56] It is important to notice this estimate since the dispersion of talent across individuals in the economy is a key driver of the aggregate losses from artisanality. If all potential entrepreneurs were of similar skills, the inability of relatively

[53] We describe the model fit for all 150 moments in Appendix C.3.

[54] Recall that the overhead time d is not included in the labor specialization. For this reason, if the overhead time allows entrepreneurs to pass-through some of their talent, this channel would be captured by a higher value of λ.

[55] Importantly, we verify that the estimated value of ω is small enough to satisfy Assumption 3, making the single-crossing hold in our estimated model.

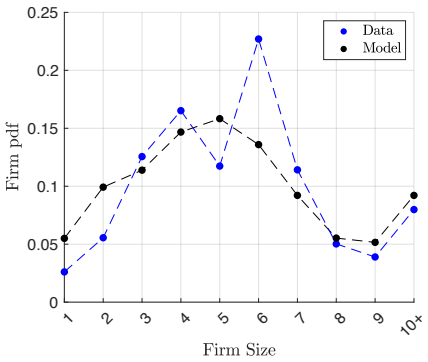
[56] The 80^{th} percentiles of the ability distribution correspond roughly to the marginal entrepreneur, given an average firm size ~ 6.

high-skilled entrepreneurs to leverage their talent would not be very consequential.

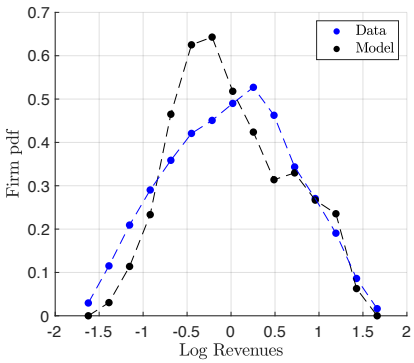
Finally, we find that κ_1 is small, which shows that Proposition 1, which characterized the equilibrium for $\kappa_1 = 0$, considers an empirically relevant case.[57]

Figure 8: Model Fit for Firm Heterogeneity and Time Allocation

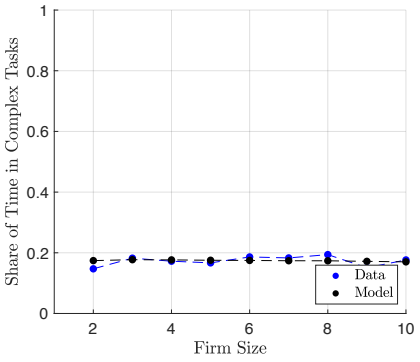
(a) Firm Sizes



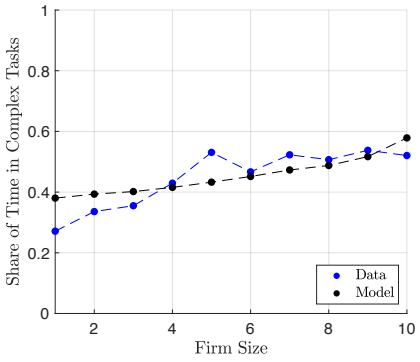
(b) Firm Revenues



(c) Employees' Complex Tasks



(d) Entrepreneurs' Complex Tasks



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

7 Quantification: the Cost of Artisanality

We use the estimated model for three purposes. First, we quantitatively assess how firms are organized. That is, we vary the degree of "artisanality," as captured by the unbundling cost κ_0 and the non-rival component of firm productivity λ, and measure

[57] One way to assess the magnitude of κ_1 is to calculate the average implied gap in 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 her complex tasks, whereas a workers at the 90^{th} percentile completes $\sim 94\%$ of hers.

the extent to which firms in our context are close to the case of self-employment within the firm. Second, we study the economic mechanisms through which the artisanal business model affects aggregate productivity. Third, we show that the artisanal business model is inherently difficult to scale and, as such, limits the returns to other interventions aimed at spurring firm growth.

7.1 Quantifying the Internal Organization of Firms

The degree of *artisanality* is modulated by two parameters: the unbundling cost κ_0 and the importance of the non-rival component λ. They both affect the ability of entrepreneurs to pass on their talent to their workers. The parameter λ determines the extent to which labor specialization is necessary to leverage entrepreneurial talent, and κ_0 determines the extent to which such specialization is costly.

Given the importance of both parameters, we perform four counterfactual exercises with different combinations of κ_0 and λ, keeping all other parameters at their estimated values. The results are in Table 8, where we compute several aggregate statistics for each counterfactual and compare them to the benchmark economy (column 1). In this section, we focus on the first four rows: labor specialization, average firm size, total output (which is the same as labor productivity), and total consumption.[58] We discuss the other rows, as well as columns 6 and 7, in the next section.

In column 2, we shut down labor specialization entirely ($\kappa_0 \to \infty$). In column 3, we consider a lower value of κ_0 calibrated to match the relationship between firm size and specialization we observe in grain milling.[59] We choose this value of κ_0 as an amount of labor specialization that is, at least in principle, attainable in our setting. In columns 4 and 5, we vary both κ_0 and λ to span the two polar cases discussed in Lemma 5. In column 4, we set $\lambda = 0$ and $\kappa_0 \to \infty$ so that the productivity of workers is independent from that of the entrepreneur. This is the case we refer to as *self-employment within the firm*. In column 5, we set $\lambda = 1$: irrespective of the extent of labor specialization, each worker inherits the ability of the entrepreneur. This is the case with *scalable entrepreneurial talent*.

This exercise offers two key takeaways. First, shutting down any specialization,

[58] Consumption is equal to output minus resources spent on hiring costs.

[59] Recall from the previous section that the parameter κ_0 is closely tied to the slope of the regression of labor specialization on firm size. For this reason, we can interpret this counterfactual as computing a hypothetical economy in which the unbundling cost is as small in carpentry and welding as it is in grain milling.

and even going as far as the extreme case of column 4, has only a relatively modest impact on firm size, output, and consumption. This result provides a quantitative answer to the question, Why do firms exist? Our baseline economy is not far from the polar case of self-employment within the firm.[60] Therefore, we conclude that, in our setting, the role of firms as vehicles for leveraging talent is relatively small. Firms must thus exist mainly as tools for sharing fixed production costs, such as the cost of the premises or the cost of machines that can be used by many individuals.

Second, the degree of artisanality in carpentry and welding in Uganda has sizable costs in terms of aggregate productivity and firm size. For example, column 3 shows that even reducing κ_0 just enough to fit the level of specialization observed in grain milling would increase aggregate productivity by 30%. The average firm size would also increase, by almost one employee, showing that firms are small, at least in part, because of their lack of specialization. Column 5 shows that in a benchmark economy with fully scalable talent, which could resemble modern manufacturing, the average firm size would be more than 20 employees, and both productivity and consumption would increase dramatically. We should, of course, take this case with a grain of salt since we are entirely changing the production technology. Nonetheless, it is a useful benchmark to show that the artisanality of manufacturing could play a major role in understanding cross-country differences.

Table 8: Model Counterfactuals for Artisanality (κ_0, λ) and Hiring Cost $(\overline{\chi_0})$

Moment	Bench	$\kappa_0 \to \infty$	Low κ_0	$\lambda = 0, \kappa_0 \to \infty$	$\lambda = 1$	High $\overline{\chi_0}$	Low $\overline{\chi_0}$
Specialization	0.09	0	0.29	0	0	0.08	0.11
Firm Size	5.89	5.38	6.84	3.08	21.46	5.11	7.03
Output	1	0.86	1.3	0.41	23.52	0.85	1.23
Consumption	1	0.91	1.2	0.56	11.96	0.9	1.13
Pass-through of Man. Ability	0.28	0.22	0.39	0	1	0.28	0.28
Log(Size) on Man. Ability	0.36	0.31	0.42	0	0.83	0.3	0.44
Average Man. Ability	1	0.94	1.1	0.64	2.1	0.91	1.12
Workers' Earnings	1	0.91	1.17	0.56	2.16	1.47	1.12
Entrepreneurs Profits	1	0.9	1.2	0.56	12.48	0.9	1.13

Notes: Statistics are created from model counterfactuals. Average managerial ability, output, consumption, workers' earnings, and entrepreneurs' profits are normalized to 1 for the benchmark values.

[60] Even in this case, firms are not of size one because of the presence of a fixed operating cost.

7.2 Understanding the Cost of Artisanal Manufacturing

We now analyze the mechanisms through which the aggregate effects of artisanality manifest and contrast them with the impact of changing the hiring cost.

Productivity Pass-through and Allocation of Labor. The degree of artisanality affects average firm size and aggregate productivity through two main channels.

First, a decline in artisanality (i.e., either a decrease in κ_0 or an increase in λ) amplifies the pass-through of entrepreneurial ability to workers' productivity. In the case of λ, this happens mechanically, whereas a decrease in κ_0 leads to a reorganization of production and hence more labor specialization. This result can be seen in row 5 of Table 8 and in the left panel of Figure 9, where we plot average worker productivity as a function of entrepreneurial ability.[61] In the model estimated for carpentry and welding (blue circles), average worker productivity increases in managerial ability, but the pass-through is smaller than one. In the extreme case of self-employment within the firm (gray squares), the pass-through is zero. In the case with scalable entrepreneurial talent (pink diamonds), the pass-through is exactly one. Reducing κ_0 (red triangles) also amplifies the pass-through.

Second, a decline in artisanality improves the allocation of labor across firms. Since entrepreneurs are more skilled than workers, the increase in pass-through just discussed boosts the average productivity for all firms, leading to an upward shift in labor demand. As a result, wages rise to clear the market and marginal entrepreneurs shift into wage work, improving the average managerial ability through a standard selection effect. Furthermore, the larger productivity pass-through steepens the relationship between firm size and managerial ability: if more skilled entrepreneurs have higher worker productivity, they choose to run larger firms. These results can be seen in rows 6 and 7 of Table 8 and in the right panel of Figure 9, which plots average firm size as a function of managerial ability.

In summary, the level of artisanality modulates the relationship between the (exogenous) distribution of talent in the economy and the endogenous distribution of firm productivity. If λ is low and κ_0 is high, the distribution of firm productivity simply mirrors the one of talent. While if λ is high or κ_0 is low, it reflects mainly the right tail of the talent distribution, as the skilled entrepreneurs hire most workers and

[61] Average worker productivity is computed as $\hat{z}^{\lambda} \int \tilde{z}(z, \hat{z}, \mu)^{1-\lambda} \frac{dF_w(z)}{F_w(z_{\max})}$, where $\hat{z}$ is managerial ability.

pass to them their ability. Overall, a decline in artisanality leads to a reallocation of labor toward talented entrepreneurs through an extensive as well as an intensive margin, and everyone benefits as a result (rows 8 and 9 of Table 8).

The Two-way Relationship between Firm Size and Specialization. Next, we compare the effects of a reduction in artisanality with a change in the hiring cost. We consider either an increase (green triangles) or a decrease (orange arrows) in the average hiring cost $\overline{\chi_0}$, which is calibrated to generate changes in firm sizes similar to those of the κ_0 counterfactuals (see columns 2, 3, 6, and 7 of Table 8).

The right panel of Figure 9 confirms that the reduction in $\overline{\chi_0}$ has a similar effect on firm size as the change in κ_0. However, the left panel of the same figure, as well as rows 1 and 5 of Table 8, show very different effects on worker productivity and labor specialization. A change in $\overline{\chi_0}$ has only a minimal impact on worker productivity and pass-through (purely operating through the selection of workers) and a small compositional effect on labor specialization, driven by the fact that larger firms are more specialized.

This last result shows that while our model encompasses both a causal relationship from labor specialization to firm size and one from firm size to specialization, the former is quantitatively stronger. For this reason, our estimated model is more consistent with the notion that firms are small because they are not specialized than the notion that they are not specialized because they are small.

7.3 Returns to Interventions in the Presence of Artisanality

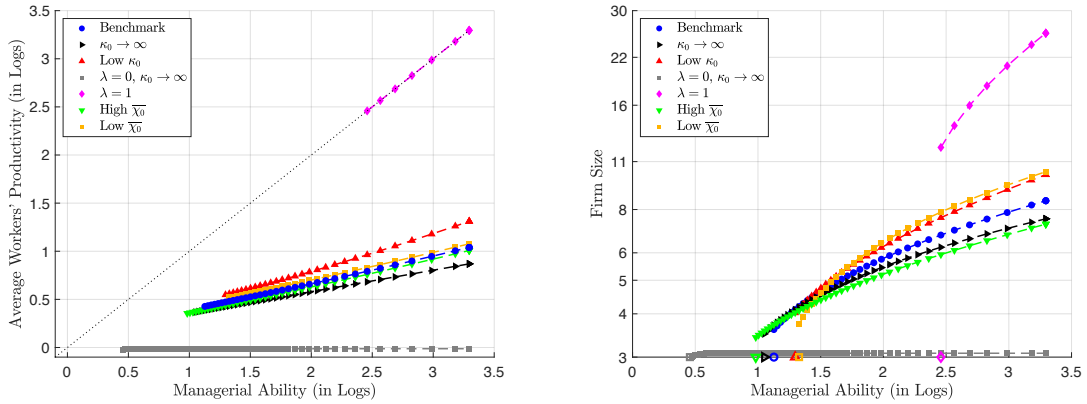
To conclude the section, we perform two exercises to show that the returns from development interventions depend on how labor is organized in the economy.

Relaxing Constraints to Firm Growth. We first study the effect of reducing the hiring cost $\overline{\chi_0}$, which could be interpreted as a policy aimed at relaxing credit or hiring constraints, on stimulating firm growth.

We start from the three alternative values of κ_0: the one estimated for carpentry and welding (in blue), a benchmark with no specialization—that is, $\kappa_0 \to \infty$ (in black)—and the value $\kappa_0 = \overline{\kappa_0}$ picked to match the level of specialization observed in grain milling (in red). We then vary the average hiring cost $\overline{\chi_0}$. The results are

Figure 9: Firm-Level Effects of Changing λ, κ_0, and $\overline{\chi_0}$

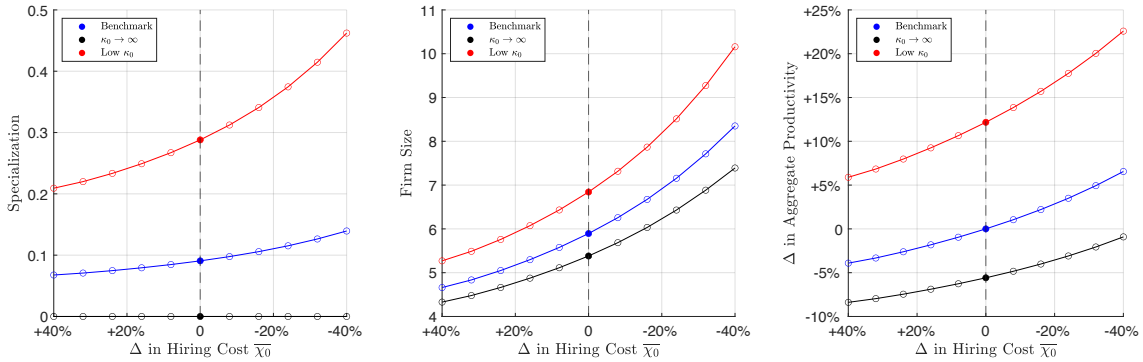
(a) Workers' Productivity (b) Firm Size



Notes: The figure shows the average productivity of workers (left panel) and firm size (right panel) as a function of managerial ability. The right panel also marks on the x-axis (with empty markers) the lowest skilled among the set of managers. The blue circles are for the benchmark estimation of the model. The other lines represent different model counterfactuals, as explained in the text.

Figure 10: Aggregate Effects of Changing the Hiring Cost $\overline{\chi_0}$

(a) Specialization (b) Average Firm Size (c) Aggregate Productivity



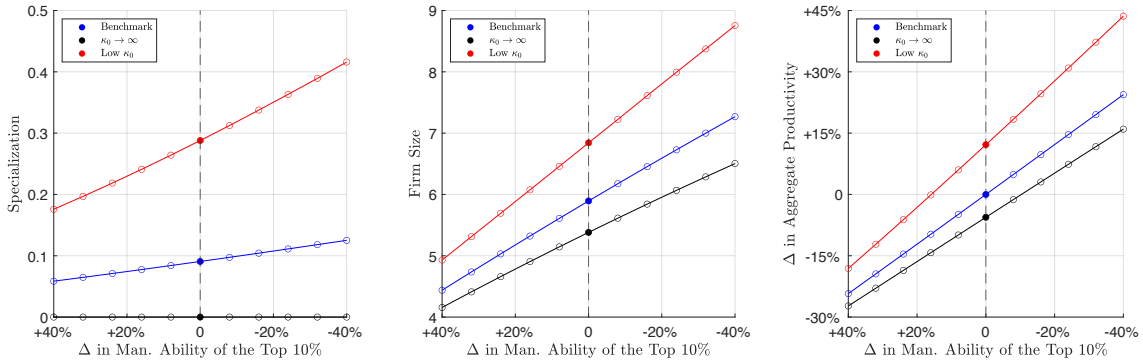
Notes: The figure shows labor specialization (left panel), average firm size (middle), and aggregate productivity (right) as a function of changes in the hiring cost ($\overline{\chi_0}$). Each lines corresponds to different values of the unbundling cost (κ_0).

shown in Figure 11, where we include three key statistics: specialization, firm size, and changes in aggregate output (hence, labor productivity).

The unbundling cost has a large impact on the return to other development policies. Relative to our benchmark, calibrating κ_0 to match the larger specialization observed in grain milling would increase the effect of a reduction in the hiring cost

Figure 11: Aggregate Effects of Changing the Managerial Ability Distribution

(a) Specialization (b) Average Firm Size (c) Aggregate Productivity



Notes: The figure shows labor specialization (left panel), average firm size (middle), and aggregate productivity (right) as a function of changes in the managerial ability of the top 10% of the distribution. Each lines corresponds to different values of the unbundling cost (κ_0).

on productivity by 60% and on firm size by 35%.

Increasing Managerial Talent. We then consider the effect of changing the distribution of talent at the top of the distribution. Specifically, we consider the 10% most skilled *individuals* and change their managerial ability.[62] This exercise is meant to capture the effect of training programs aimed at increasing the managerial ability of the most promising entrepreneurs. As for the previous case, we consider three different values of the unbundling cost.

The unbundling cost again modulates the returns from the intervention, which is approximately 30% higher if we calibrate κ_0 to the level of specialization of grain-milling. Training the most skilled entrepreneurs is more valuable if they are able to specialize on complex tasks and pass-through their newly acquired knowledge to their workers. In a world with limited specialization, training one entrepreneur only increases her productivity. In a world with specialized firms, it could increase the productivity of a whole village.

Taken together, these exercises highlight a key takeaway of our work. Barriers to within-firm labor specialization make artisanal manufacturing a business model that is difficult to scale. As a result, the returns from policy interventions aimed at

[62] Notice that we pick the 10% most skilled individuals, without conditioning on their occupation. In the baseline estimation, they would all be entrepreneurs in equilibrium.

spurring firm growth may be limited because entrepreneurs face strong decreasing returns to scale coming from obstacles to leveraging their talent.

8 Conclusion

This paper makes two contributions. First, we collect unique data on time-use within a representative sample of manufacturing firms in a developing country. Second, combining this survey data with a model, we study barriers to labor specialization inside the firm and quantify their link to low returns to scale and low productivity.

Our key finding is that even though most firms in our data have enough employees to potentially specialize, they do so only to a limited extent: the internal organization of labor resembles that of self-employed individuals sharing a production space. We show that this limited specialization reduces the extent to which entrepreneurs can leverage their talent by hiring more workers, thus leading to smaller and less-productive firms. Using our survey data, we argue that the limited specialization is a consequence, at least in part, of the prevalence of product customization, which creates communication and coordination costs within the firm and makes it difficult to "unbundle" the production process into separate tasks that can be performed by specialized individuals.

Our work has important implications for industrial policy. First, we show that the return to development interventions is likely to vary across sectors (or countries) depending on the internal organization of firms and their implied business scalability. In our context, supply-side interventions to help firm growth (such as credit or hiring subsidies) would have a larger aggregate impact if targeted toward grain milling rather than carpentry or welding. This underscores the importance of designing context-specific industrial policies (Juhász et al., 2023).

Second, our results highlight the importance of demand-side policies in fostering development (Goldberg and Reed, 2022). The lack of product standards and small market size prevalent in low income countries could contribute to the prevalence of the artisanal business model, hindering productivity, and even reducing the effectiveness of supply-side interventions. A policy priority is therefore to promote product standards and connect firms with larger markets. This would create an opportunity for entrepreneurs to scale up their business and leverage the capabilities they already possess, thus increasing productivity and firm size, and, ultimately, reducing poverty.

References

Adhvaryu, A., V. Bassi, A. Nyshadham, J. Tamayo, and N. Torres (2023): "Organizational Responses to Product Cycles," *Available at SSRN 4403515*. 6

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. 1

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. 1

Atack, J. (1987): "Economies of Scale and Efficiency Gains in the Rise of the Factory in America, 1820-1900," *Quantity and Quiddity: Essays in US Economic History*, 286–335. 3

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. 7

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. 7

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. 18

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. 7

Banerjee, A. V. and E. Duflo (2014): "Do firms want to borrow more? Testing credit constraints using a directed lending program," *Review of Economic Studies*, 81, 572–607. 1

Bassi, V., M. E. Kahn, N. Lozano Gracia, T. Porzio, and J. Sorin (2022a): "Searching for Customers, Finding Pollution," Tech. rep., mimeo. 16

Bassi, V., R. Muoio, T. Porzio, R. Sen, and E. Tugume (2022b): "Achieving scale collectively," *Econometrica*, 90, 2937–2978. 9, 13, 16, 6.3, A.2.1, A.2, C.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. 72

Becker, G. S. and K. M. Murphy (1992): "The division of labor, coordination costs, and knowledge," *The Quarterly journal of economics*, 107, 1137–1160. 1

Blattman, C. and L. Ralston (2015): "Generating employment in poor and fragile states: Evidence from labor market and entrepreneurship programs," *Available at SSRN 2622220*. 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. 1, 1

Bloom, N., L. Iacovone, M. Pereira-López, and J. Van Reenen (2022): "Management and misallocation in Mexico," Tech. rep., National Bureau of Economic Research. 5

Bloom, N., A. Mahajan, D. McKenzie, and J. Roberts (2010): "Why do firms in developing countries have low productivity?" *American Economic Review*, 100, 619–23. 1

Boehm, J. and E. Oberfield (2023): "Growth and the Fragmentation of Production," Tech. rep., mimeo. 7

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. 1

Bruhn, M., D. Karlan, and A. Schoar (2018): "The impact of consulting services on small and medium enterprises: Evidence from a randomized trial in Mexico," *Journal of Political Economy*, 126, 635–687. 1

Caliendo, L., F. Monte, and E. Rossi-Hansberg (2015): "The anatomy of French production hierarchies," *Journal of Political Economy*, 123, 809–852. 7

Caliendo, L. and E. Rossi-Hansberg (2012): "The impact of trade on organization and productivity," *The quarterly journal of economics*, 127, 1393–1467. 6

Chandler, A. D. (1990): *Scale and scope: The dynamics of industrial capitalism*, Harvard University Press. 1

Dessein, W. and T. Santos (2006): "Adaptive organizations," *Journal of political Economy*, 114, 956–995. 4, 6

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. 73

Freund, L. (2022): "Superstar Teams: The Micro Origins and Macro Implications of Coworker Complementarities," *Available at SSRN 4312245*. 7

Garicano, L. (2000): "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108, 874–904. 1

Garicano, L. and E. Rossi-Hansberg (2006): "Organization and inequality in a knowledge economy," *The Quarterly journal of economics*, 121, 1383–1435. 1, 4.2.3

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, 8

Goldin, C. and L. F. Katz (1998): "The origins of technology-skill complementarity," *The Quarterly journal of economics*, 113, 693–732. 3

Hardy, M. and J. McCasland (2022): "Are small firms labor constrained? experimental evidence from ghana," *American Economic Journal: Applied Economics*. 1

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

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, 4, 3.2, 17

Hsieh, C.-T. and B. A. Olken (2014): "The Missing "Missing Middle"," *The Journal of Economic Perspectives*, 28, 89–108. 1, 1, 3.1

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, 8

Juhász, R., N. J. Lane, and D. Rodrik (2023): "The New Economics of Industrial Policy," Tech. rep., National Bureau of Economic Research. 8

Juhász, R. and C. Steinwender (2018): "Spinning the web: Codifiability, information frictions and trade," *NBER working paper*. 18

Katz, L. F. and R. A. Margo (2014): "Technical change and the relative demand for skilled labor: The united states in historical perspective," in *Human capital in history: The American record*, University of Chicago Press, 15–57. 3

Kohlhepp, J. (2023): "The Inner Beauty of Firms," . 7

Lucas, R. (1978): "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 9, 508–523. 5.4

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. 73

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. 1, 73

——— (2014): "What are we learning from business training and entrepreneurship evaluations around the developing world?" *The World Bank Research Observer*, 29, 48–82. 73

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

——— (2021): "Training Entrepreneurs," Tech. rep., VoxDev Lit. 1, 2

Piore, M. J. and C. F. Sabel (1984): "The second industrial divide: possibilities for prosperity," . 4, 3.2

Sokoloff, K. L. (1984): "Was the transition from the artisanal shop to the nonmechanized factory associated with gains in efficiency?: Evidence from the US Manufacturing censuses of 1820 and 1850," *Explorations in Economic History*, 21, 351–382. 3

Startz, M. (2019): "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade," *mimeo*. 1, 16

Vickery, S., C. Dröge, and R. Germain (1999): "The relationship between product customization and organizational structure," *Journal of Operations Management*, 17, 377–391. 4

Vitali, A. (2022): "Consumer Search and Firm Location: Theory and Evidence from the Garment Sector in Uganda," *mimeo*. 1, 16

Yang, X. and J. Borland (1991): "A microeconomic mechanism for economic growth," *Journal of political economy*, 99, 460–482. 1

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 of the tasks by the average firm in parenthesis. Panel B shows the production steps for the core product in the three sectors with the share of production time accounted for by each production 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) Carpentry		**(ii) Welding**		**(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 pre-specified step. The steps are listed in the typical order of implementation. This information is available only for firms producing the core product, and the statistics in Panel B are conditional on doing a given step. The data from Panel B comes from the survey module asking the entrepreneur and each employee about whether they usually work on a given step. This information is available for entrepreneurs and employees but only in firms with at least one employee.

A.2 Firm Characteristics and Representativeness: Details

A.2.1 Firm Characteristics and Output Market

Figure A.1 reports the distribution of revenues per worker, profits and wages across sectors. Three findings emerge. First, the firms in our sample are not only relatively large, but also productive: the average firm generates about $233-278 per month in revenues per worker across our sectors. To put these numbers into perspective, per capita GDP in Uganda was $60 per month in 2018. Second, there is large dispersion in labor productivity and profits. In Bassi et al. (2022b), we show that these differences across firms are strongly correlated with managerial ability, thus suggesting that there would be gains from reallocating resources to higher ability managers. Third, these firms offer remunerative employment with wages of about $52-74 per month, but profits are substantially higher, so that entrepreneurs earn significantly more than employees.

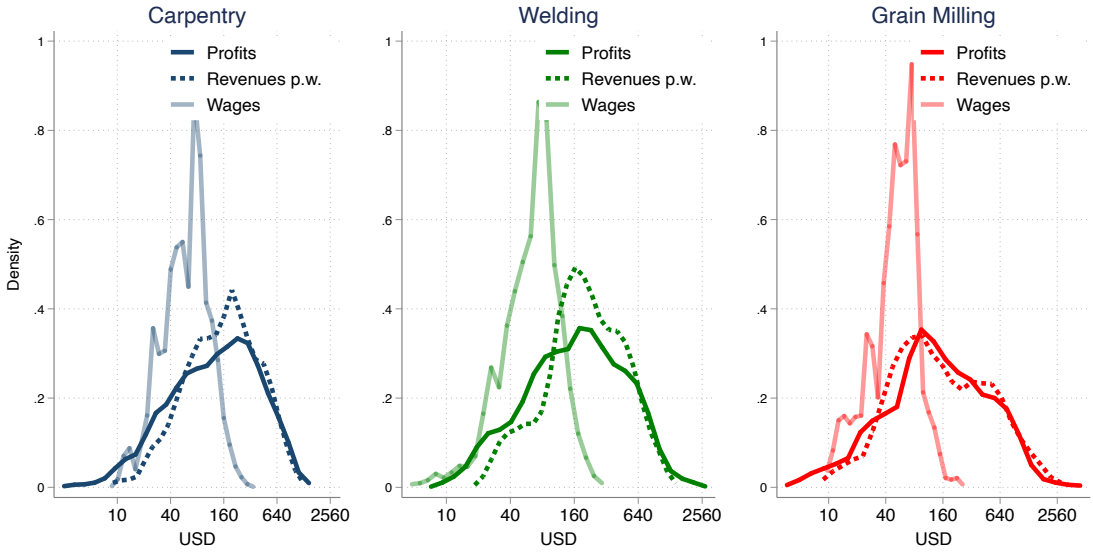
Panel A of Table A.2 shows that our sample comprises of a mix of relatively large and small, yet well-established firms. Across the firm size distribution, we consistently observe that firms have been in business for at least 8 years and are typically registered with the local authority (over 80% on average). The average firm uses about 20% of the pre-specified modern machines listed in our survey, indicating that the production process is at least partly mechanized for firms of all sizes. The table also shows that employees are slightly less educated than entrepreneurs on average. These manufacturing firms offer stable employment: on average, employees have been at the firm for about 3.5 years. About 90% of employees are paid at least a piece-rate component (the share paid exclusively piece-rate is around 80%).

Large and small firms differ on important dimensions, such as managerial ability, monthly revenues and profits. For example, the average large carpentry firm (with firm size exceeding five workers) earns $1,491 in monthly revenues relative to $975 earned by the average small carpentry firm (with firm size five or below). However, Panel B shows that by and large *all* firms primarily operate in local markets, relying on similar strategies to promote sales. Few firms report marketing expenses and instead, personal interactions with customers are key: for close to 53% of firms, directly talking to customers is their main strategy to signal product quality. In line with this, most sales are made through walk-ins. These features of how firms find customers and sell to them are consistent with significant difficulties in accessing

wider markets, as documented by a growing literature on output market frictions in developing countries. Consistent with this, lack of demand is indicated as a main constraint to growth by the majority of firms. These results confirm that in all three sectors, firms not only produce output, but also generate demand through personal interactions with final customers.

These are sectors with very limited imports and exports. Less than 1% of the firms in our data export. Using the 2019 VAT and customs data for Uganda, we show that imports and exports in the carpentry sector amount to just about 3% and 1% of domestic sales.[63]

Figure A.1: Distribution of Revenues per Worker, Profits, and Wages



Notes: Sample: all surveyed firms. The x-axis is in log scale. Revenues per worker (profits): average revenues per worker (profits) in the three months preceding the survey. Wages: total employee earnings in the last month. Revenues per worker, profits and wages are trimmed at the top 1%. 1 USD = 3,800 UGX for monetary amounts.

A.2.2 Representativeness of Our Sample

We assess the representativeness of our sample by comparing it to two sources of administrative data on firms available for Uganda: (i) the 2010 Census of Business

[63] In the supplemental appendix, we show that the price of carpentry goods in Uganda is low relative to that of similar IKEA products across the world, which can explain why penetration of large multinationals producing standardized products through imports is limited.

Table A.2: Descriptives on Basic Firm Characteristics by Size

	Carpentry		Welding		Grain Milling	
	<=5	>5	<=6	>6	<=6	>6
Number of firms	211	311	254	179	86	74
Panel A. Firm Characteristics						
Monthly revenues (USD)	975	1491	1174	2386	1046	3181
Monthly revenues per worker (USD)	259	204	275	284	186	316
Firm age (yrs.)	10.0	10.8	8.2	10.7	13.5	10.0
Formal license	69.7	81.0	83.1	89.4	84.9	98.6
Mechanization rate (%)	22.5	24.2	16.4	19.4	16.9	22.9
Monthly profits (USD)	177.5	265.2	209.3	367.5	130.6	403.0
Managerial score	10.6	11.9	10.3	12.3	10.1	14.0
Entrepreneur education (yrs.)	9.5	10.3	9.8	10.5	10.8	11.2
Employee education (yrs.)	8.8	9.1	10.2	10.3	7.9	7.9
Employee tenure (yrs.)	3.4	3.7	3.2	3.7	4.0	3.6
Monthly wage (USD)	69.4	77.8	67.5	79.9	41.8	65.7
Employees paid piece-rate (%)	95.5	91.3	94.0	90.7	78.1	83.7
Panel B. Sales and Marketing						
Share of sales to final consumers	97.2	92.3	99.2	99.0	60.6	57.0
Sold outside Uganda	0	1.2	1.4	0.1	0.0	0.3
Sales are done through walk-ins	79.6	79.6	76.3	71.1	95.3	91.1
Any marketing expenditure	4.7	10.8	7.9	20.9	6.8	9.2
Talks to consumers to communicate quality	54.4	56.8	54.0	46.3	42.7	52.7
Lack of demand as a main constraint	52.1	56.8	57.7	44.8	60.9	50.6

Notes: Means are reported by median firm size in each sector. Sample: surveyed firms. Panel A presents summary statistics at the firm level. Monthly revenues: average revenues in the three months preceding the survey (trimmed at top 1%). 1 USD = 3,800 UGX for monetary amounts. Mechanization rate: types of machines used relative to all potential pre-specified types of machines for producing the core product. This is reported for firms that produce doors in carpentry, windows in welding, and maize flour in grain milling. Managerial score: index going from -1 to 27 and based on multiple survey questions. See Bassi et al. (2022b) for details. Panel B presents descriptive statistics on common sales and marketing channels, e.g. whether the firm talks to consumers to communicate product quality (0/1). Sold outside Uganda: dummy variable set to 1 if the firm sold to final consumers, other firms or wholesalers outside Uganda in the last three months.

Establishments, conducted by the Uganda Bureau of Statistics (UBOS), and (ii) the Corporate Income Tax (CIT) data collected by the Uganda Revenue Authority (URA), which is available from 2013 to 2020. The UBOS census is meant to cover the entire population of firms, including both formal and informal establishments.

The CIT data should include all firms with more than $40,000 in yearly revenues, which are required to pay the corporate income tax (CIT). Throughout the analysis, we keep all the firms in our 52 sampled sub-counties, and focus on carpentry and welding, as these are our two key sectors of interest (as discussed in Section 3).

In Panel (a) of Figure A.2, we first compare the number of firms in our initial survey (conducted in 2018/19) with the number of firms in the UBOS census and in the URA data. The figure shows two key results. First, we find that our data includes twice as many firms as the UBOS census. This reassures us that our survey thoroughly covered our sample sub-counties.[64] Second, even when restricting our sample to firms above the CIT revenue threshold, our data still includes about 8 times more firms than the URA data. This is important because it shows that our data also includes many "large" firms above the URA revenue threshold (it also suggests that compliance with CIT could be relatively weak in this setting).

In Panel (b) we narrow in on the comparison with the URA data, focusing on 2018/19, the same year as our initial survey. The figure shows two key results. First, it shows that there is significant overlap in the distribution of revenues in our data and the URA data.[65] This is reassuring, because it implies that our results apply also to the typical "large" and formal firm in these sectors. Second, Panel (b) shows that the URA data includes a handful of firms with very large yearly revenues above $1 million, which are not covered in our survey data. However, we note that these firms at the top of the sales distribution in the URA data have a wage bill-to-sales ratio of less than 5% (and in some cases close to 1%), thus suggesting that they may be large importers, plausibly foreign owned, rather than manufacturers.

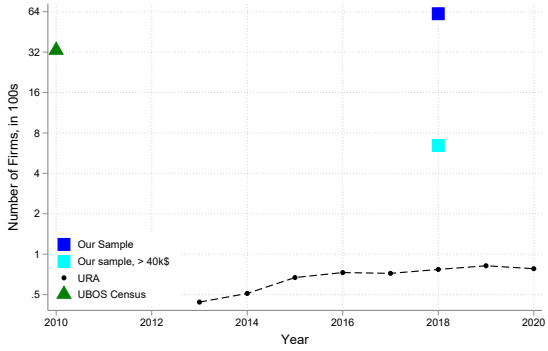
Overall, the analysis in this section shows that our sample is representative of both small and large firms. While we have not been able to reach some of the very large formally registered firms, we notice that those are so few that in the aggregate they still constitute just a relatively small share of total sales, and an even smaller share of the total employee earnings.

[64] Since the UBOS census is from 2010, some of this difference can reflect net firm entry until 2018/19. Of course, it is also possible that there was some non-compliance in the UBOS census.

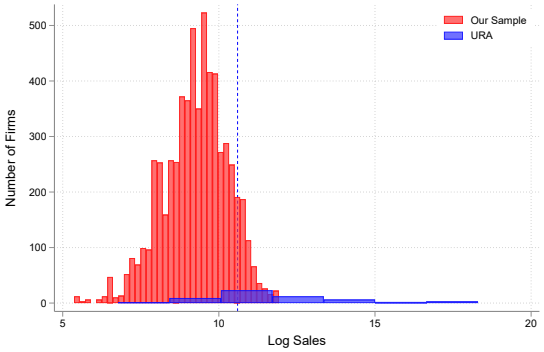
[65] As Panel (b) of Figure A.2 shows that many of our firms are close to the URA threshold, this can help explain why the number of firms in our data above the threshold is so much larger than in the URA data (as firms right around the threshold may be less likely to register for CIT).

Figure A.2: Comparison with Firm Census and Corporate Income Tax Datasets

(a) Number of Firms



(b) Distribution of Revenues



Notes: Sample: The left panel reports the number of firms in carpentry and welding: (i) in the 2010 UBOS census, (ii) in our sample, both overall and restricting to firms with yearly sales above the threshold to register for corporate income tax ($40,000), and (iii) in the Corporate Income Tax dataset of the URA. The right panel reports the distribution of log yearly revenues in our sample and in the firms in the 2018/19 CIT data of URA. The vertical line represents the threshold to be included in the URA data ($40,000).

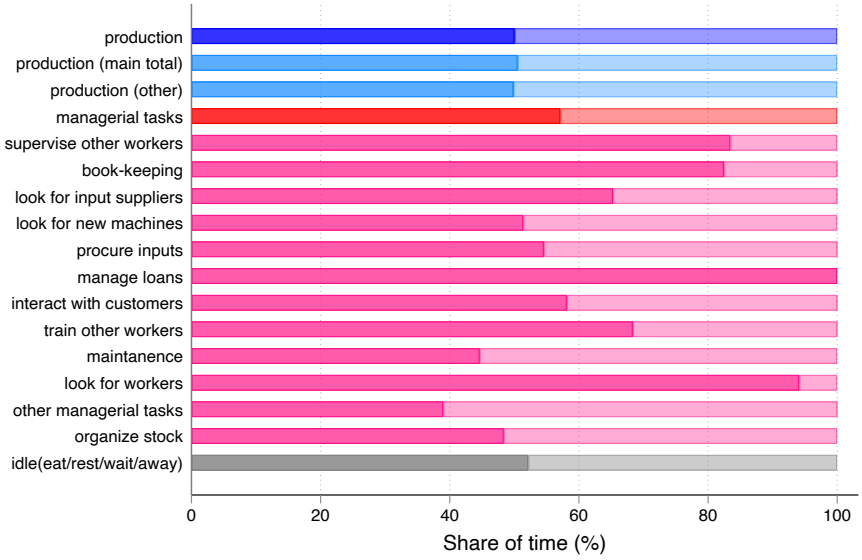
A.3 Additional Evidence on Labor Specialization

In this section, we report several additional results and robustness checks on labor specialization that are mentioned in Section 4 of the main text.

Limited Specialization in Non-production Tasks Between Employees. Figure A.3 replicates Figure 4 but comparing employees above and below median skills (splitting employees by earnings within each firm). The Figure shows much higher overlap and substantially less evidence of specialization in non-production tasks for more skilled employees. For instance, when focusing on the headline summary categories of "production" and "managerial tasks" (depicted in the first and fourth bars starting from the top of the figure) we clearly see that the time allocation of higher and lower skilled employees is very similar.

Non-production Tasks Are More Complex. In Appendix Table A.3 we provide evidence to support the claim that non-production tasks are more complex. We do so by focusing on employees, and 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

Figure A.3: Task Allocation Between Production and Non-production Tasks by Employee Skills



Notes: Replication of Figure 4 in the main text but comparing skilled and unskilled employees. Dark bars: skilled employees. Light bars: unskilled employees. Sample: all surveyed firms in the Carpentry and Welding sectors. Time use reported by interviewed employees. 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.

non-production tasks are indeed more complex tasks, 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 their time on non-production tasks earn substantially more: going from no involvement in non-production tasks 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 earnings returns from involvement in non-production tasks, thus suggesting that they are more complex—i.e., it is more challenging to perform them well. 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 various non-production categories. We find that supervision/training, interaction with customers and input procurement are the categories driving the gains in earnings. Taken together, columns 1 and 2 suggest that non-production tasks are more complex, as those employees assigned to them earn more. 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.

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.

Variation in Task Difficulty Within Production. In column 3 of Table A.3 we show that even within production there is evidence of vertical differentiation in terms of task difficulty. To do so, 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. Consistent with this result, in Appendix Figure A.7 we show that higher-skilled employees (as measured by earnings) spend a larger share of their time on difficult production steps.

Entrepreneurs Are More Skilled than Employees. In Appendix Table A.4 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 sector 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

Table A.3: Heterogeneity in Skill Intensity of Tasks

	(Log) Employee Earnings		
	(1)	(2)	(3)
Time Share Non-prod. Tasks	0.302		0.259
	(0.089)		(0.097)
Supervise/Train (0/1)		0.213	
		(0.071)	
Customer Int. (0/1)		0.088	
		(0.044)	
Input Procurement (0/1)		0.102	
		(0.041)	
Org. Stock (0/1)		0.005	
		(0.042)	
Other Managerial Tasks (0/1)		-0.047	
		(0.075)	
Avg. Difficulty of Prod. Steps Performed			0.270
			(0.095)
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, robust standard errors in parentheses. Sample: all interviewed employees in the Carpentry and Welding sectors with non-missing time use responses. The dependent variable is the log of total employee monthly earnings. Column 1: for each employee we compute their share of time on non-production tasks and use this as independent variable. Column 2: we include dummy variables taking value one if the employee performs each task (the reference groups 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, looking for new machines, looking for 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, where the weights are the time spent on each step. Each employee was asked to state on a 1 to 5 scale their self-reported ability to perform each step. We compute the difficulty level of each step by taking the average of employee responses to this question. Demographic controls: age, years of education, tenure at the firm, and a dummy for whether the worker received vocational training.

tenure; however differences between employees are overall less pronounced, compared to differences between the entrepreneurs and the employees, apart from education, where the gaps are similar in columns 1 and 2.[66] This confirms that entrepreneurs

[66]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 further the gap between entrepreneurs and employees in column 1.

Table A.4: 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	
	(0.145)		(0.397)		(0.275)	
Skilled (0/1)		0.739		3.892		1.562
		(0.211)		(0.601)		(0.236)
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, robust standard errors in parentheses. Sample: all interviewed entrepreneurs and employees (odd columns) in the Carpentry and Welding sectors; all interviewed employees (even columns) in the Carpentry and Welding sectors. 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.

on average are more skilled than employees, and the main margin of heterogeneity in skills is between entrepreneurs and employees.

Relationship Between Specialization and Firm Size: Robustness. Table A.3 shows that not all non-production activities are equally complex, as proxied by the correlation between the share of time spent on a given task and employee earnings: supervision/training is associated with the highest earnings returns (and is therefore the most complex task according to this definition), and customer interaction and input procurement also have positive returns; on the other hand, the return is not significant for stock organization and other managerial tasks, indicating that these tasks may be less complex.

Table A.5: Task Allocation and Firm Size: Robustness

	Carpentry and Welding		Grain Milling	
	Owner	Worker	Owner	Worker
	(1)	(2)	(3)	(4)
All Non-Production Tasks	0.021	0.000	0.063	-0.001
	(0.005)	(0.003)	(0.012)	(0.007)
Cust. Int. + Superv. + Input Proc.	0.021	0.004	0.052	0.001
	(0.005)	(0.002)	(0.012)	(0.004)
Customer Interactions	0.004	0.001	0.018	-0.001
	(0.002)	(0.001)	(0.006)	(0.003)
Supervision/Training	0.014	0.003	0.036	0.002
	(0.004)	(0.001)	(0.011)	(0.003)
Input Procurement	0.003	-0.000	-0.002	0.000
	(0.002)	(0.001)	(0.004)	(0.000)

Notes: Sample: all surveyed firms. Non-Prod. Tasks refer to non-production tasks (same definition as in Figure 5). All regressions include region dummies and sector dummies when necessary. Robust standard errors clustered at the firm level are reported 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).

In Tables A.5 and A.6 we test the robustness of the findings in panels (a) and (b) of Figure 5 by focusing only on the share of time spent in the three most complex non-production activities. Specifically, Table A.5 reports the slopes (and standard errors) of the relationship between share of time spent by the entrepreneur (columns 1 and 3) or employees (columns 2 and 4) with firm size. We calculate these slopes using: (i) the share of time in all non-production tasks (as in Figure 5); (ii) the share of time in any of customer interactions, supervision/training and input procurement; (iii) the share of time separately in each of customer interactions, supervision/training and input procurement. Comparing the first and second row of the table, we see that when restricting attention to the three most complex non-production tasks, the slopes are very similar to those reported in the main text: the relationship between time spent in non-production tasks remains positive for entrepreneurs, and close to zero for workers, and the slopes for entrepreneurs remain larger in grain milling. The last three rows of the table show that the relationship between firm size and share of time in non-production tasks for entrepreneurs is driven primarily by customer interactions and supervision/training.

Table A.6: Average Time Allocation of Tasks: Robustness

	Carpentry and Welding		Grain Milling	
	Owner	Worker	Owner	Worker
	(1)	(2)	(3)	(4)
All Non-Production Tasks	0.457	0.178	0.652	0.108
Cust. Int. + Superv. + Input Proc.	0.322	0.088	0.481	0.063
Customer Interactions	0.091	0.036	0.132	0.040
Supervision/Training	0.151	0.028	0.333	0.023
Input Procurement	0.081	0.024	0.016	0.000
All Non-Prod., >1 yr. Tenure Firms	0.463	0.177	0.685	0.114
All Non-Prod., >2 yr. Tenure Firms	0.450	0.163	0.662	0.093

Notes: Sample: all surveyed firms. Non-Prod. Tasks refer to non-production tasks. 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). >1 year tenure firms (row 6) refers to firms that only comprised of employees with above one year of experience at the firm. One year denotes the 90th percentile of on-the-job training duration provided to employees. >2 year tenure firms refers to firms that only comprised of employees with above two years of experience at the firm. The robustness of our results to the sample restrictions in rows 6 and 7, which represent 55.4% and 33.2% of the sample respectively, suggests that our findings are not driven purely by training/apprenticeship considerations.

In Table A.6 we test robustness of the average involvement of entrepreneurs and employees in non-production tasks to the same alternative definitions of non-production tasks. Regardless of the specific definition of non-production tasks considered, we find that: (i) entrepreneurs spend a higher share of time in non-production tasks compared to employees; (ii) entrepreneurs in grain milling specialize in non-production tasks more than in the other two sectors; (iii) the gap between entrepreneurs and employees is larger in grain milling. Taken together, the evidence from tables A.5 and A.6 confirms that the results from Figure 5 are not sensitive to the specific definition of non-production tasks considered.

Entrepreneur's Involvement in Production Not Driven by Apprenticeship Motives. Figure 5 shows that entrepreneurs spend a large share of their time in production activities. One possible concern is that this may reflect training motives, whereby entrepreneurs spend time in production to train employees as part of apprenticeship schemes. To address this, in the last two rows of Table A.6 we calculate again the share of time spent by entrepreneurs and employees in non-production tasks but restricting the sample to firms (i) employing no workers with less than one year

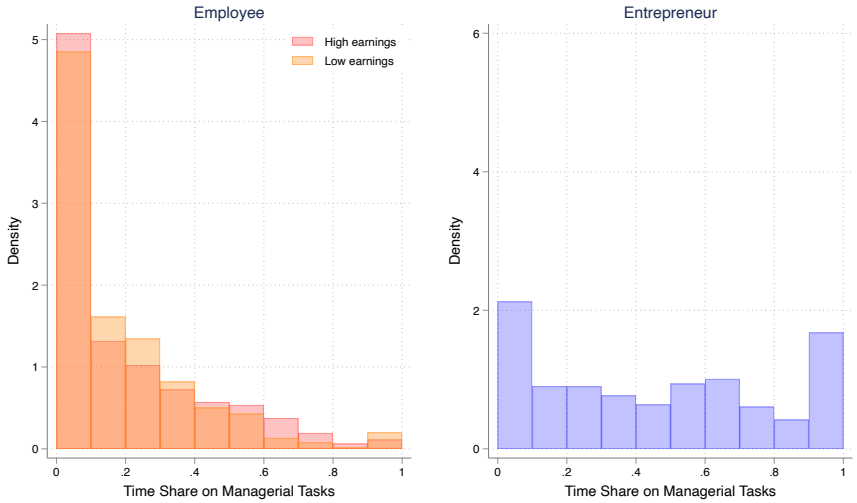
of tenure, and (ii) employing no workers with less than two years of tenure. We choose these tenure cutoffs because the 90^{th} percentile of the distribution of length of on-the-job training in our sample is one year. Therefore, by applying these sample restrictions, we remove firms where at least some employees are likely to be in training. Comparing the last two rows of the table with the first row (which includes all firms) we notice that these sample restrictions do not affect the results. This confirms that training motives cannot be a main explanation for the substantial involvement of entrepreneurs in production.

Heterogeneity Across Workers Is Limited and Does Not Vary with Firm Size. We explore whether the findings in panel (a) of Figure 5 that there is little heterogeneity in specialization among employees and this does not vary with firm size can mask the fact that some workers may be spending a lot of time in non-production tasks, while others very little. Figure A.4 reports the distribution of time share in non-production tasks among workers (left panel) and entrepreneurs (right panel). Since our measurement of time use refers to the last day worked (rather than being an average measure over a longer time horizon), 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.[67] We farther validate this in Figure A.5, 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 remarkably 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.

Consistency of Measurement of Non-production Tasks by Firm Size. One potential concern with the results in Figure 5 is that our measurement of non-production tasks may not be consistent across the size distribution. For instance,

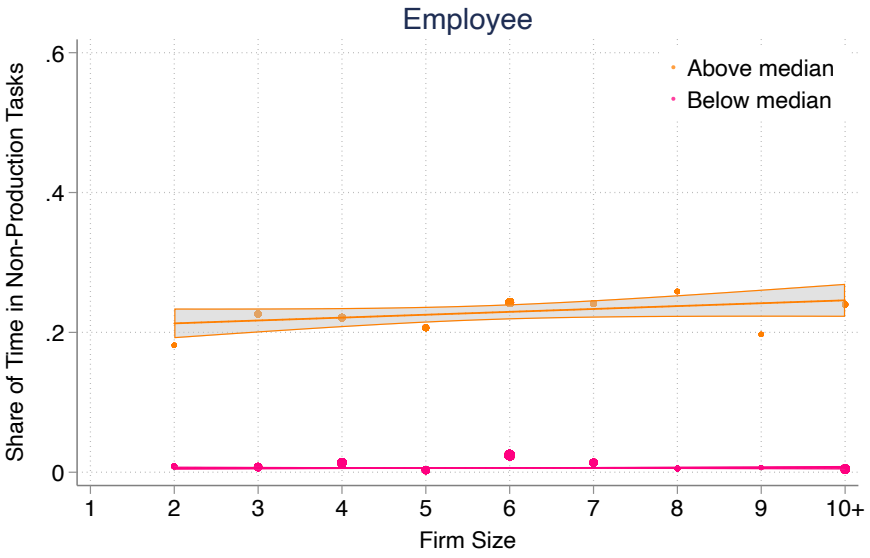
[67] The right panel shows that there is more variation among entrepreneurs, and this is consistent with panel (b) of Figure 5, where we see that entrepreneurs in larger firms spend more time in non-production tasks, so that there is heterogeneity among entrepreneurs.

Figure A.4: Distribution of Time Allocation of Employees and Entrepreneurs



Notes: Distribution of share of time spent on non-production (or "managerial") tasks. Left panel: Employees, classified as high and low earnings by their earnings within each firm (above and below the median). Right panel: Entrepreneurs. Sample: all surveyed firms in the Carpentry and Welding sectors.

Figure A.5: Limited Increase in Specialization Across Employees with Firm Size

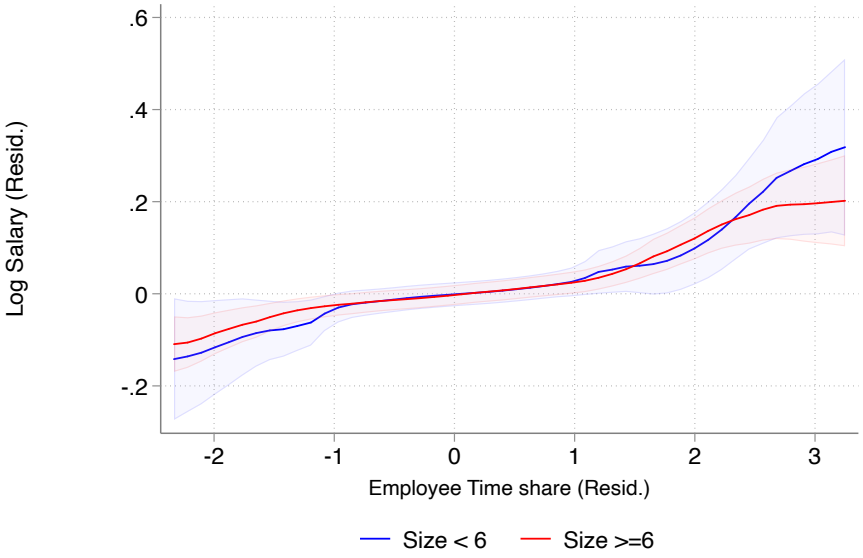


Notes: Share of time employees spend on non-production tasks by firm size. Orange lines correspond to employees that spend an above-median share of time on non-production tasks within each firm-size group; pink lines correspond to those that spend a below-median share. Sample: all surveyed firms in the Carpentry and Welding sectors.

what we label in the survey as "customer interactions" might represent a slightly different activity in smaller and larger firms, and this activity might be differentially

complex in smaller and larger firms. One way to address this concern is to show that the return (in terms of employee earnings) of being involved in non-production tasks are similar across the size distribution. We show that this is the case in Figure A.6, which reports the non-parametric regression of log employee earnings on the share of time spent on non-production tasks, splitting employees by whether they work in firms below or above the median size (and controlling for other observable worker characteristics). The figure confirms that the return to working on non-production tasks is very similar in smaller and larger firms, thus suggesting that the complexity of non-production tasks is similar across the firm size distribution.

Figure A.6: Consistency in Measurement of Non-production Tasks Across the Size Distribution



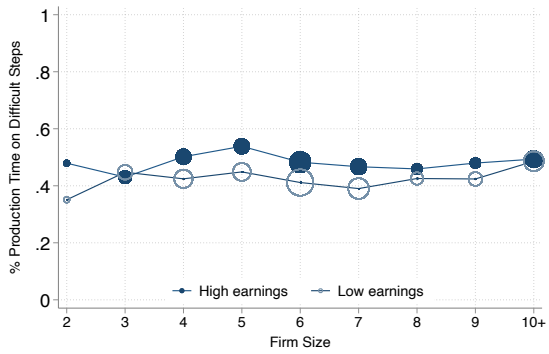
Notes: Sample: all surveyed firms in the Carpentry and Welding sectors. The figure shows the non-parametric regression (with 95% confidence intervals) between residualized log employee total earnings in the last month and the residualized share of time spent in non-production tasks, splitting the sample by whether the employee is working in firms of size below or above the sample median. To residualize both variables, we regress them on employee characteristics, as well as sector and region dummies.

Labor Specialization Between Difficult and Simple Steps. To study specialization across production steps of different difficulty, we exploit a survey question where each employee ranked (on a scale 1 to 5) their ability to perform each production step conducted by the firm (regardless of whether the particular employee performs that step). We use this to rank production steps and then we split them

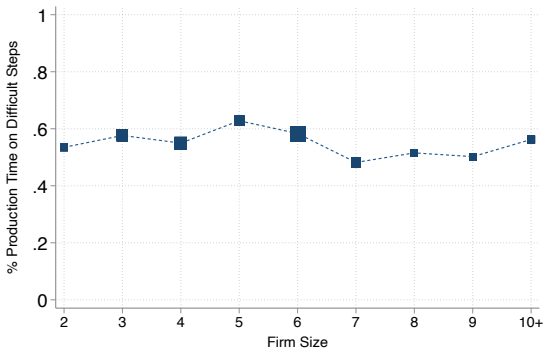
by above/below median difficulty. In Figure A.7, we study how employees and entrepreneurs allocate their production time to simple and difficult steps. If an individual only works on difficult steps, the share of time in difficult steps would be 100%. The Figure shows that: (i) high skilled employees 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 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.7: Task Allocation Across the Size Distribution, by Task Difficulty

(a) High & Low Earning Employees



(b) Entrepreneurs

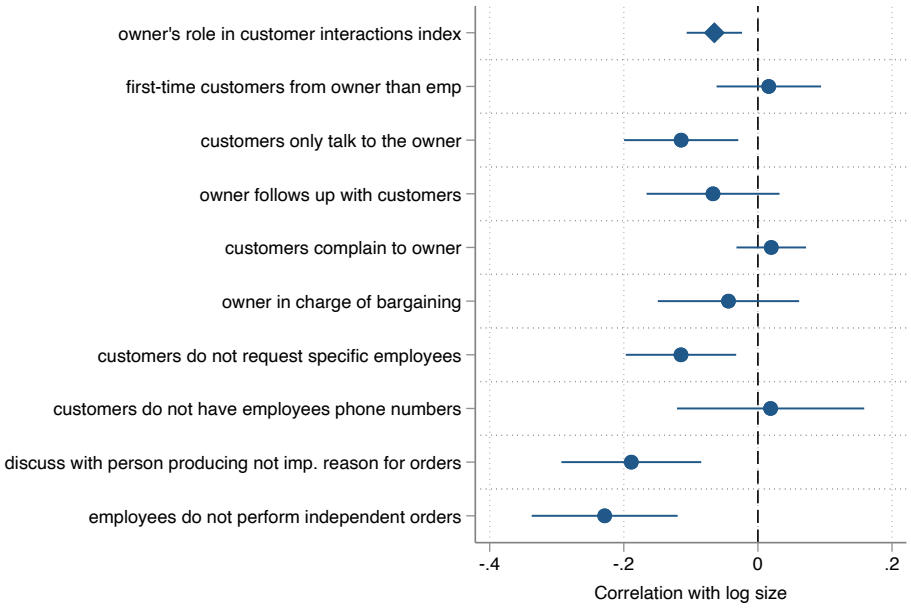


Notes: Sample: all surveyed firms in the Carpentry and Welding sectors. Panels (a) and (b) represent the share of worker and entrepreneur production time spent on difficult production steps. Production steps are categorized as difficult (or simple) using employee assessments of their ability to perform each production step conducted by the firm, following the procedure described in the text. Steps that are ranked as having above median levels of difficulty are regarded as 'difficult'.

In Larger Firms, Employees Play a Larger Role in Customer Generation. We use the follow-up phone survey to provide evidence that in larger firms employees play a larger role in customer generation and interactions. To do so, we regress several indicators for the involvement of employees and entrepreneurs in these activities on firm size, controlling for region and sector fixed effects. The indicators are all dummy variables, and we recode each so that a value of 1 is consistent with entrepreneurs playing a prominent role in customer interactions, while 0 means that employees play a larger role. We first regress each of these indicators on firm size separately, and then

combine them in an index by taking their simple average. The results are presented in Figure A.8. The coefficients on most indicators are negative, and the coefficient on the overall index is negative and significant. This confirms that entrepreneurs tend to play a less important role in customer interactions in larger firms, and so is in line with the results in panel (d) of Figure 5.

Figure A.8: Entrepreneur's Involvement in Customer Interactions and Firm Size



Notes: Correlation between entrepreneurs' role in customer interactions and firm size, where the markers represent point estimates of the coefficients, and the lines represent the 90% CI. The "owner's role in customer interaction index" is the mean of the other nine measures (all dummies). Variable definition, starting from the top: (i) overall index; (ii) whether first-time customers originate primarily from entrepreneurs' contacts rather employees'; (iii) whether customers only talk to the owner (vs. talking also to employees); (iv) whether it is the owner who follows up with customers after the order is placed (vs. whether the employees do it); (v) whether customers would complain to the owner if they are unhappy about the product (vs. whether they would complain to the employees); (vi) whether the owner is in charge of bargaining (vs. whether the employees can also do it); (vii) whether customers do not request specific employees when placing an order; (viii) whether the customer does not have the phone number of the person producing; (ix) whether an important reason why the customers buy on order is to discuss the details of the product with the person producing; (x) whether the worker performs independent orders. All regressions include sector fixed effects and region fixed effects. Sample: Carpentry and Welding sector.

Correlation Between Specialization and Productivity. We study the correlation between specialization and firm productivity. We follow two approaches. First, we use data from the initial survey, and run an OLS regression of log revenues on the share of time spent by the entrepreneur on non-production tasks (computed as in Fig-

ure 5), and controlling for firm size, sector and region fixed effects. The results are in column 1 of Table A.7. The coefficient on the share of the entrepreneurs' time on non-production tasks shows that going from spending no time at all on non-production activities to spending all the entrepreneur's time on non-production activities is associated with a 13% increase in revenues. Since we are controlling for firm size, we interpret this as an increase in (revenue) productivity. However, the coefficient is not significant at conventional levels. We note though that our available measure of revenues is at the monthly level, while the share of time in non-production tasks comes from the time-use survey and so refers to the last day worked. Variation across firms in this variable (conditional on firm size) could therefore likely capture sizable measurement error or random noise. In light of this, it is perhaps not surprising that the estimate in column 1 is not significant.

Our second approach overcomes this measurement challenge by creating a measure of specialization that is less subject to daily variation and therefore noise. To do so, we exploit 14 questions from the follow-up survey, where we ask for several practices related to specialization of the entrepreneur on non-production activities and to the division of labor among employees. We average the resulting 14 (dummy) variables into one index of specialization, which goes from 0 to 1 and takes larger values if the firm implements more of these practices. In column 2, we regress log monthly revenues (at the time of the follow-up survey) on this measure of specialization, again controlling for firm size. The estimates show that going from no specialization to full specialization (based on the index) is associated with an increase in revenues of 49%, and this is significant at the 10% level. Figure A.9 reports coefficients from separate regressions of revenues on each of the 14 components of the index, again controlling for firm size. Reassuringly, we see that most of the coefficients are positive.

Taken together, this evidence provides suggestive evidence that more specialized firms have higher revenue productivity.

Heterogeneity in Specialization by Sector. Figures A.11, A.11 and A.12 replicate Panel (a) of Figure 3, panels (a) and (b) of Figure 5 and panels (a) and (b) of Figure 6 by sector. Throughout, we see that the similarity between carpentry and welding is remarkable, while there is more specialization in grain milling firms, along all dimensions.

Table A.7: Correlation Between Revenues and Specialization

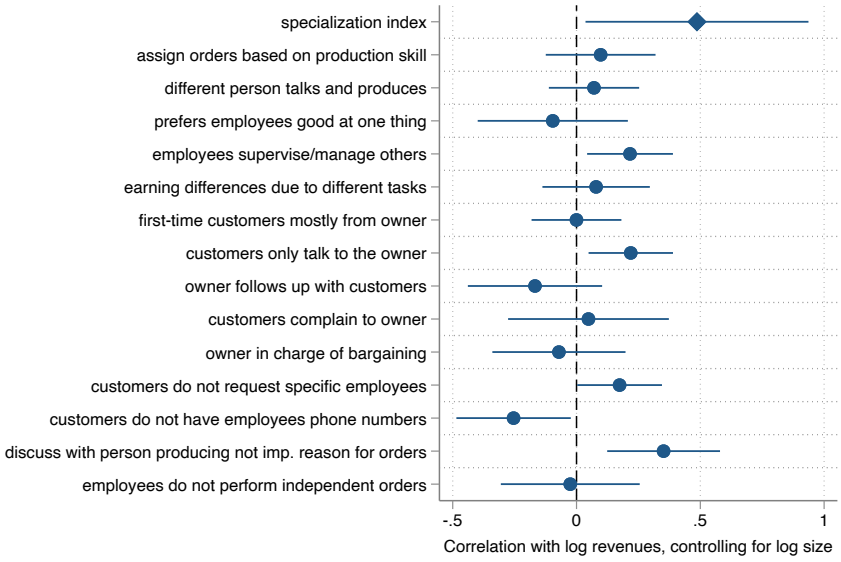
	$\log(Rev_0)$ (1)	$\log(Rev_t)$ (2)
% time of non-prod. tasks by entrepreneur$_0$	0.131	
	(0.108)	
log(size$_0$)	0.688	
	(0.076)	
Specialization index$_t$		0.487
		(0.273)
log(size$_t$)		0.665
		(0.109)
Adjusted R^2	0.312	0.096
Observations	890	568

Notes: Sample: all surveyed firms in the Carpentry and Welding sectors. Non-prod. tasks refer to non-production tasks. The specialization index is constructed in the same way as in Figure A.9. All regressions include region fixed effects and sector fixed effects. Robust standard errors are reported in parentheses. Subscript 0 represents baseline data, and subscript t indicates follow-up data.

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 carpentry and welding. Second, while in grain milling employees in larger firms are significantly less idle (apart from around lunch time), this is not the case in carpentry and welding, 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, idle time could reflect the presence of coordination and communication costs since workers, due for example to poor coordination, might have to be idle while waiting for others to finish their tasks. The larger idle time in carpentry and welding is thus consistent with our claim that the prevalence of customization creates sizable communication and coordination costs. Second, 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. This validates

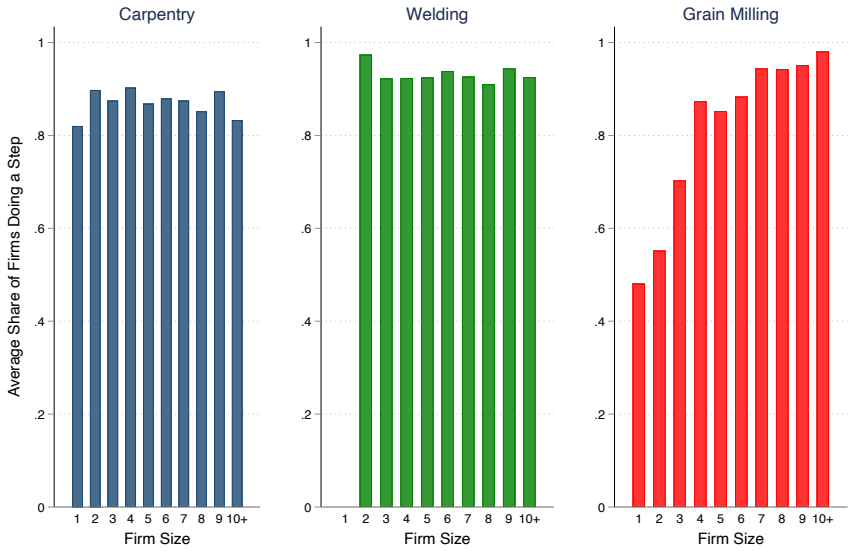
Figure A.9: Correlation Between Proxies for Specialization and Revenues



Notes: This figure depicts the correlation between within-firm specialization and firm size, where the markers represent point estimates of the coefficients, and the lines represent 90% CIs. The specialization index is the mean of the other 14 measures. The 14 measures include all the variables from Figure A.8, plus the following variables from the top (all dummies): (i) whether the owner assigns orders to employees based on production skills (vs other reasons such as who is available or who knows the customer well); (ii) whether the person talking to customers is usually or always different than the person producing; (iii) whether the owner prefers having employees outstanding in one task and poor at another (vs having employees that are equally fair in both); (iv) whether half of the employees or more regularly supervise other workers; (v) whether the main reason for earnings differences among employees is that they work on different tasks. All regressions include sector fixed effects and region fixed effects. Sample: all firms in Carpentry and Welding sectors.

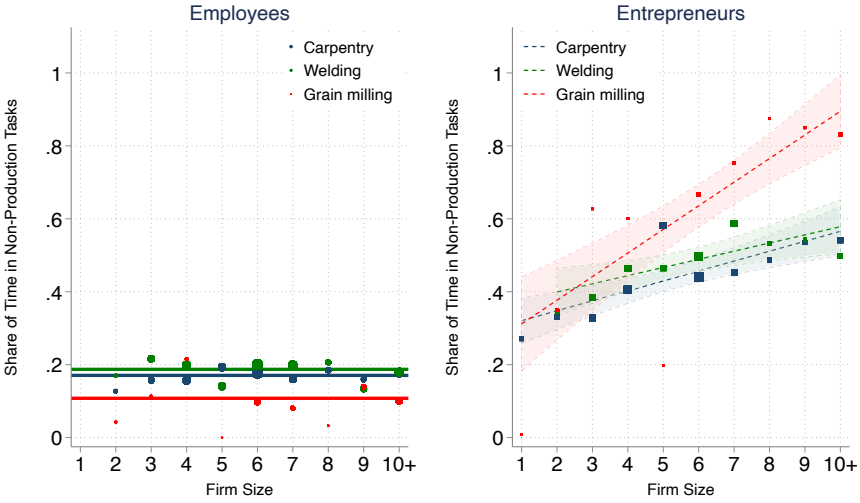
the heterogeneity analysis in the previous figures using a different measure of labor specialization.

Figure A.10: Heterogeneity in Share of Firms Performing a Production Step by Sector



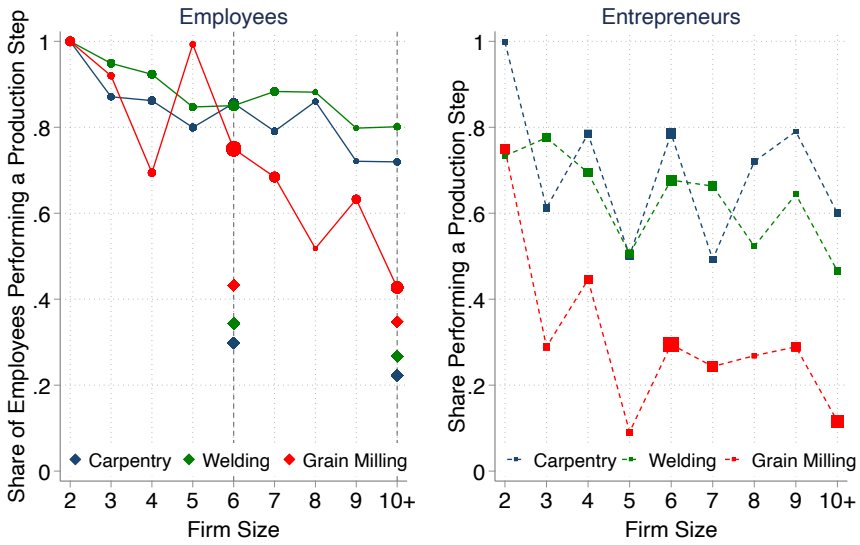
Notes: Replication of Figure 3, Panel (a), in the main text, separately for each sector.

Figure A.11: Heterogeneity in Time Allocation to Non-production by Sector



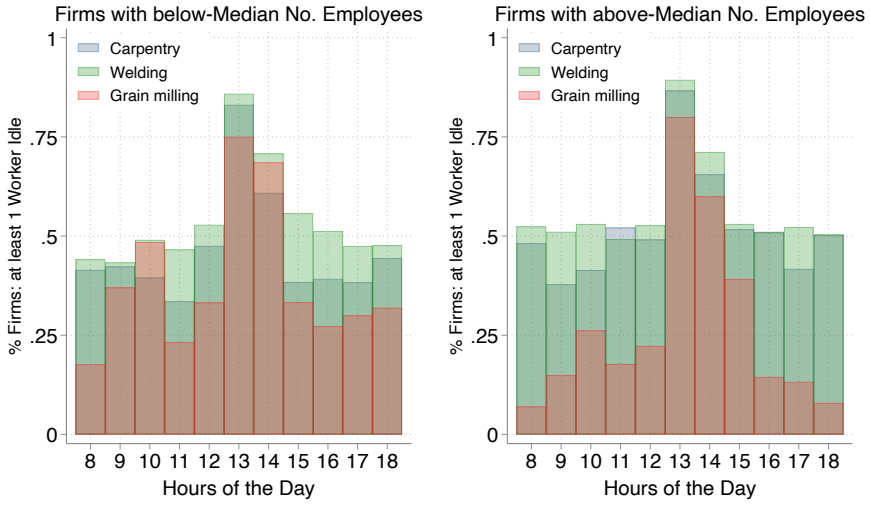
Notes: We replicate panels (a) and (b) of Figure 5, but separately for each sector. All surveyed firms. Shaded areas: 95% confidence intervals. The size of dots and squares represent the number of firms in each size group. Time use reported by interviewed entrepreneurs and employees.

Figure A.12: Heterogeneity by Sector in Specialization within Production Across Steps



Notes: We replicate panels (a) and (b) of Figure 6, but separately by sector. These figures depict the share of employees (left panel) and entrepreneurs (right panel) working on a production step by sector. The navy, green, and red lines correspond to the carpentry, welding and grain milling sectors, respectively. The size of the dots represents the within-sector weight of firms. The navy, red and green diamond markers in the employee panel represent the share of employees that would work on a production step under full specialization (reported for firms of size 6 and 10 only).

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



Notes: These figures depict the pattern of workers' idle time in a day by sector. 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 lines correspond to the carpentry, welding, and grain milling sectors, respectively. The left panel depicts the share of firms among those with the size below the median of its sector, the right panel is for firms of above-median size.

B Online Appendix - Model

B.1 Proofs

PROOF OF LEMMA 1. We want to show that $\exists z_0 \geq 0$ such that

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

The proof proceeds in three steps

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

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

1. From Equation (5.4) combined with the solution to the wage bargaining in Equation (5.5), we can write the derivative of profits wrt to owner ability as

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

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

   Turning to expected wages,

$$\frac{\partial \mathbb{E}_{\hat{z}}(w(z, \hat{z}, \mu))}{\partial z} = \omega \int \frac{\partial y(z, \hat{z}, \mu)}{\partial z} \frac{dF_o(\hat{z})}{F_o(z_{\max})} \leq 1 \tag{B.2}$$

   by Assumption 3.

2. Suppose instead that $\pi(0) > \int w(0, \hat{z}, \mu) \frac{dF_o(\hat{z})}{F_o(z_{\max})}$ and, 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 $\pi(z_{\max}) < \int w(z_{\max}, \hat{z}, \mu) \frac{dF_o(\hat{z})}{F_o(z_{\max})}$ and, 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 D – the total amount of complex tasks in his production line – minus the share of tasks delegated to the entrepreneur. Using Equation (5.8), this can easily be rewritten as $\theta(z, \hat{z}) = D\left(1 - \frac{1}{\kappa_0}(\log \hat{z} - \log z)^{\kappa_1}\right)$. 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. $\hat{\theta}(\hat{z}) = D\left(1 + \frac{n-1}{\kappa_0}\int (\log \hat{z} - \log z)^{\kappa_1} \frac{dF(z)}{F_w(z_{max})}\right)$

2. The expression for $\overline{\theta}(\hat{z})$ follows directly from using the expressions above in the definition of average labor specialization.

$$\begin{aligned}
\frac{\partial \overline{\theta}(\hat{z})}{\partial \kappa_0} &= -D\frac{n}{\kappa_0^2}\int (\log \hat{z} - \log z)^{\kappa_1} \frac{dF(z)}{F_w(z_{max})} && \leq 0 \\
\frac{\partial \overline{\theta}(\hat{z})}{\partial n} &= D\frac{1}{\kappa_0}\int (\log \hat{z} - \log z)^{\kappa_1} \frac{dF(z)}{F_w(z_{max})} && \geq 0 \\
\frac{\partial^2 \overline{\theta}(\hat{z})}{\partial \kappa_0 \, \partial n} &= -D\frac{1}{\kappa_0^2}\int (\log \hat{z} - \log z)^{\kappa_1} \frac{dF(z)}{F_w(z_{max})} && \leq 0
\end{aligned}$$

■

Proof of Lemma 3. Rearranging Equation (5.2) and using Equation (5.7) gives the result.

■

Proof of Lemma 4.

Let $\overline{w}(\hat{z},\mu) \equiv \int w(z,\hat{z},\mu)\frac{dF_w(z)}{F_w(z_{\max})}$. The equation in Lemma (4) follows from taking the first-order condition of (5.4) with respect to n. Further,

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

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

Solving for n,

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

which is declining in χ_0.

Using the envelope theorem,

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

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

■

Proof of Lemma 5.

1. When $\lambda = 1$, $Y(\hat{z},n,\mu)$ directly collapses to $\hat{z}n$. When $\kappa_0 = 0$, then $\mu(z) = 1 \forall z$. Note that Assumption (2) guarantees that the entrepreneur has capacity to take on all complex tasks of her workers. With $\mu(z) = 1$, we have again that $Y(\hat{z},n,\mu) = \hat{z}\,n$. Optimal firm size is then simply given by

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

   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 $\mu(z) = 0\ \forall z$. Hence $\tilde{z}(z,\hat{z},\mu) = z$. Moreover, if $\lambda = 0$, we get that $Y(\hat{z},n,\mu) = \hat{z} + (n-1)\int z\frac{dF_w(z)}{F_w(w_{max})}$ and optimal firm size solves

   $$\chi'(n) = \int z\frac{dF_w(z)}{F_w(w_{max})}$$

   and is thus 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

$$Y(\hat{z}, n) = \hat{z} + (n-1)\,\hat{z}^{\lambda+\frac{1-\lambda}{\kappa_0}} \int_0^{z_0} z^{(1-\lambda)(1-\frac{1}{\kappa_0})} \frac{dF(z)}{F(z_0)} \tag{B.3}$$

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 $\mathbb{Z}(\hat{z}, z_0)$. That is,

$$\mathbb{Z}(\hat{z}, z_0) \equiv \hat{z}^{\lambda+\frac{1-\lambda}{\kappa_0}} \int_0^{z_0} z^{(1-\lambda)(1-\frac{1}{\kappa_0})} \frac{dF(z)}{F(z_0)} \tag{B.4}$$

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

$$\pi(\hat{z}; z_0) = \hat{z} + (n-1)(1-\omega)\left[\mathbb{Z}(\hat{z}, z_0) - \overline{w}\right] - \chi(n) \tag{B.5}$$

where n is equal to

$$n = \frac{1}{\chi_0}\left[(1-\omega)\left(\mathbb{Z}(\hat{z}, z_0) - \overline{w}\right)\right]^{\chi_1} \tag{B.6}$$

The expected wage of a worker z is equal to

$$\mathbb{E}\left(w(z;\mu)\right) = (1-\omega)\overline{w} + \omega\mathbb{Z}_w(z; z_0) \tag{B.7}$$

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

$$\mathbb{Z}_w(z_0) \equiv z^{(1-\lambda)(1-\frac{1}{\kappa_0})} \int_{z_0}^{z_{\max}} \hat{z}^{\lambda+\frac{1-\lambda}{\kappa_0}} \frac{n(z)dF(z)}{\int_{z_0}^{z_{\max}} n(z)dF(z)} \tag{B.8}$$

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

$$z_0 + (n-1)(1-\omega)\left[\mathbb{Z}(\hat{z}, z_0) - \overline{w}\right] - \chi(n) = (1-\omega)\bar{w} + \omega\mathbb{Z}_w(z_0), \tag{B.9}$$

$$\int_{z_0}^{\overline{z}} n(z)f(z)dz = 1. \tag{B.10}$$

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 $\overline{w}^*$.

$$z_0 + (n-1)(1-\omega)\left[\mathbb{Z}(z_0, z_0, \kappa_0^*) - \overline{w}^*\right] - \chi(n) = (1-\omega)\bar{w}^* + \omega\mathbb{Z}_w(z_0, \kappa_0^*) \tag{B.11}$$

$$n^*(z_0) = \frac{1}{\chi_0}\left[(1-\omega)\left(\mathbb{Z}(z_0, z_0, \kappa_0^*) - \overline{w}^*\right)\right]^{\chi_1} \tag{B.12}$$

Combining the two equations:

$$n^* = \frac{1}{\chi_0}\left[\frac{(1-\omega)\bar{w}^* + \omega\mathbb{Z}_w(z_0, \kappa_0^*) - z_0 + \chi(n^*)}{n^* - 1}\right]^{\chi_1} \tag{B.13}$$

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

$$\chi_0 dn = \chi_1 \left[\frac{(1-\omega)\bar{w}^* + \omega \mathbb{Z}_w(z_0, \kappa_0^*) - z_0 + \chi(n^*)}{n^* - 1}\right]^{(\chi_1 - 1)} \times \left[\frac{(1-\omega)d\bar{w}^* + \omega \frac{\partial \mathbb{Z}_w(z_0,\kappa_0^*)}{\partial 1/\kappa_0} d(1/\kappa_0) + \chi'(n)dn}{n-1} - dn\frac{(1-\omega)\bar{w}^* + \omega \mathbb{Z}_w(z_0, \kappa_0^*) - z_0 + \chi(n^*)}{(n-1)^2}\right] \tag{B.14}$$

$$\chi_0 dn = \chi_1 \left(\chi_0 n\right)^{\frac{\chi_1 - 1}{\chi_1}} \frac{(1-\omega)d\bar{w}^* + \omega \frac{\partial \mathbb{Z}_w(z_0,\kappa_0^*)}{\partial 1/\kappa_0} d(1/\kappa_0)}{n-1} \tag{B.15}$$

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 $\kappa_1 = 0$, average specialization simplifies to $\bar{\theta} = D\frac{n}{\kappa_0}$, which is increasing in $1/\kappa_0$.

2. $\frac{\partial \bar{\theta}}{\partial n} = \frac{D}{\kappa_0}$ 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

$$y(z, \hat{z}, \mu) = \hat{z}^{\lambda}\left(z^{\mu(z)}\hat{z}^{1-\mu(z)}[1-\mu(z)]\right) \tag{B.16}$$

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 ($\mu(z)$ 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 $w(z,\hat{z},\mu) = (1-\omega)\overline{w}+\omega z^{(1-\lambda)(1-\frac{1}{\kappa_0})}\hat{z}^{\lambda+\frac{1-\lambda}{\kappa_0}}$ 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.

■

C Online Appendix - Estimation

In this section, we provide details to complement Section 6 on the model's estimation. First, in Subsection C.1 we explain how all the moments are calculated. Then, in Subsection C.2 we provide details on our estimation procedures and its effectiveness. Finally, in Subsection C.3 we include the full description of the model's fit since in the main text we only offered a summary of it.

C.1 Details on Empirical Moments and Calibration

We describe the computation of the 150 moments targeted in the model estimation, and of the calibrated fixed operating costs. The analysis pools the carpentry and welding sectors together, and all moments are computed from the baseline survey. The calibrated fixed cost, instead, uses the follow up survey which was designed on purpose. We start by describing the 150 moments, and organize the discussion by dividing moments into four groups, following the four panels of Table 6.

C.1.1 Allocation of Time to Complex Tasks (Table 6, Panel A)

We start from Panel A of Table 6. The Average Time on Complex Tasks (row (i)) is computed as the average firm-level share of time in non-production tasks, including the entrepreneur and all employees in the firm. 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.6. Rows (iv) and (v) report the average share of time in non-production tasks for employees, splitting them by below and above median salary (we use salary as a proxy of skill).[68] The slope for entrepreneurs in row

[68] In order 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 within

(vi) is taken from column 1 of Table A.5, that is, we regress the share of time of the entrepreneur in non-production tasks on firm size, top-coding firm size at 10 workers (as in the analysis in Section 4) and controlling for sector and region fixed effects. Rows (vii) and (viii) of Panel A report the coefficients from a similar regression for high- and low-skilled workers separately, where again we split workers by below and above median salary for this, as described above. The results are reported in columns 2 and 3 of Table C.1. Finally, row (ix) reports the coefficient from a regression of share of time in non-production tasks on log employee earnings, controlling for region and sector fixed effects. The results are in column 1 of Table C.1.

We also target several moments related to the distribution of specialization in complex tasks across the size distribution (shown in Figure 8 and in Table C.7). Specifically, we calculate the share of time of entrepreneurs in non-production tasks in each firm size group, from firms with no employees to firms with 10 workers or above (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, as described above). This yields other 9 moments, also shown in Table C.7.

C.1.2 Distribution of Earnings (Table 6, Panel B)

The coefficient in row (i) of Panel B, Table 6, is from a regression of employee log monthly earnings on the index of managerial ability, controlling for region and sector fixed effects. The results are in column 4 of Table C.1. We then standardize this coefficient by dividing it by the standard deviation of employee log earnings.[69] The coefficient in row (ii) is from a regression of employee log monthly salary on log revenues per worker, controlling for sector and region fixed effects. The results are in column 5 of Table C.1. Row (iii) of Panel B 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 de-

each firm size group, and are split above and below median salary within each firm size group.

[69] 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.

viation. 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 Table C.6 and also visualized in Figure C.4). 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.[70] This produces other 20 targeted moments (again shown in Table C.6 and in Figure C.4).

C.1.3 Distribution of Firm Revenues (Table 6, 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, controlling for region and sector fixed effects. The results are in column 1 of Table C.2. 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. 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 Table C.5 and visualized in Figure C.4). In addition, we also target the pdf of residualized log firm revenues (visualized in Figure 8). 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 (in Table C.5).

C.1.4 Firm Size Distribution (Table 6, Panel D)

Average size in the first row of Panel D is uncensored. The standard deviation of log size and of size in the second and third row 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, and controlling for sector and region fixed effects. The results are in column 2 of

[70] 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.

Table C.2. Finally, row (v) shows the average gap in firm size between entrepreneurs above and below median managerial ability. To compute this, we simply create the distribution of firm size (censored at 10 workers), separately for above and below median managerial ability firms. 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 Table C.4 and in Figure C.4). 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 8).

Table C.1: Moments, Employee Level Regressions

Dep. Var.:	Worker Share of Time in Non. Prod.			log(Salary)	
Sample:	All	High Skilled	Low Skilled	All	All
	(1)	(2)	(3)	(4)	(5)
log(Salary)	0.033				
	(0.012)				
Firm Size		0.002	0.000		
		(0.003)	(0.004)		
Managerial Score (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: all firms in Carpentry and Welding. In Col (2) and (3), employees are classified as high- and low-skilled within sector by size groups. All firms with more than 10 employees are classified as having 10 employees. log(Salary) is in US Dollar. For model estimation, we standardize the coefficient in Col (4) by dividing it by the standard error of residualized log salary, trimmed at 5% and 95% level, controlling for region and sector FEs. In Col (1), robust standard error is reported; in the other columns, robust standard errors clustered at the firm level are reported. The Managerial Score variable is standardized. For its definition, see Table A.2.

C.1.5 Calibration of Start-up Cost

In Appendix Table C.3 we present estimates of the start-up capital in carpentry and welding (column 1), and compare this with monthly profits in the first year of

Table C.2: Moments, Firm Level Regressions

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

Notes: Sample: all firms in Carpentry and Welding. Robust standard errors clustered at firm level are reported. The Managerial Score variable is standardized. For its definition, see Table A.2.

Table C.3: 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	281.802
Median	657.895	65.789	189.781
Obs.	308	303	2627
Sample	Follow-up		Initial survey

Notes: Sample: all surveyed firms in carpentry and welding sectors. All numbers are in USD. Column (1) and (2) summarize data from the follow-up survey, Column (3) summarizes data from the initial survey. Start-up Capital is defined as reported savings and money received from others (including but not limited to loans and gifts from family/bank/money-lender) in order to start the business. We trimmed the sample at 1% level and excluded all 0 values. Monthly profit (first year) is defined as the reported average monthly profits in the first year as owner. We trimmed the sample at 1% level. Column (3) is the average monthly profit in the three months before the initial survey, the sample is trimmed at 1% level.

operation (column 2), and monthly profits just before the initial survey (column 3).[71] To calculate the start-up capital, we exploit a unique survey module included in the follow-up survey, where each entrepreneur was asked to report: (i) all personal savings and (ii) all external sources of funds (e.g., loans, gifts) received to start the business, i.e., to cover all expenses needed in the first month of operation. We sum (i) and (ii) to create a measure of start-up capital. The value of the average start-up capital in the sample of carpentry and welding firms is $903. In the same survey module about

[71] Note that the number of observations is around 300 firms 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 overall survey length.

start-up costs, entrepreneurs were also asked to recall their profit in an average month in the first year as owners of the business. The average monthly profit in the first year of operation is \$106.7, or 11.8% of the mean start-up capital. The median value of the start-up capital is \$657.9, and the median monthly profit in the first year of operation is about \$65.79, or 10% of the median start-up capital. Given these results, we calibrate the start-up cost χ_f as 10% of average profits (as shown in Table 7).

C.2 Estimation Procedure and Identification

The objective of the estimation is to find a parameter vector ϕ^* that solves

$$\phi^* = \arg\min_{\phi \in \mathbb{F}} \mathbf{L}, (\phi) \tag{C.1}$$

$$\mathbf{L}(\phi) \equiv \sum_x \left[\omega_x \left(T_x \left(m_x \left(\phi \right), \hat{m}_x \right) \right)^2 \right] \tag{C.2}$$

where $\mathbb{F}$ is the set of admissible parameter vectors. The choice of the function T_x minimizes the sum of the percentage deviations between model-generated and empirical moments. We introduce a weighting factor ω_x to give equal weight to each one of the 21 groups of moments that we target, shown in Table 6. In practice, we then use a fairly standard simulation-based minimization routine to solve the Problem C.1, following our previous Bassi et al. (2022b).

Figure C.1 shows key outcomes of the estimation routine plotting the likelihood function $\mathbf{L}(\phi)$ for the best 5,000 parameter draws together with the estimated values of each parameter (in red dash line). The figure shows that the likelihood function is smooth and single-peaked around our estimated values.

Jacobian Matrix. To formally explore the connection between parameters and moments, we compute the elasticity of each (model generated) moment to each estimated parameter. We start from the estimated vector of parameters ϕ^*, and we create 22 alternative vectors, two for each parameter j, one that keeps all parameters except for j constant and decreases j by 5% ($\underline{\phi}(j)$), while one that does the same, but increasing j by 5% ($\overline{\phi}(j)$). We then use the model to compute the vectors of moments corresponding to each vector of parameters and use them to compute $\Delta_{jr} = m_r(\overline{\phi}(j)) - m_r(\underline{\phi}(j))$, which measures how much moment r would change if we

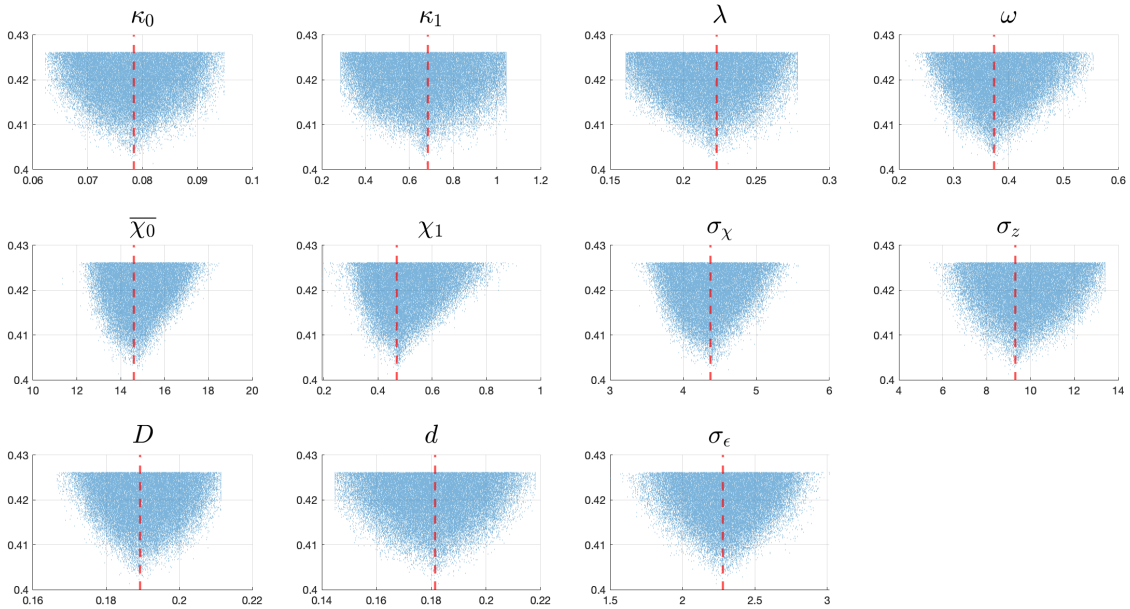
changed only parameter j by 10% around the estimated vector.

To ease comparison, we normalize Δ_{jr} for each moment r so that, when rounded, it sums to 22 across all parameters. The result of this procedure is the Jacobian matrix shown in Figure C.2, which illustrates which parameter is most important for each moment. Magnitudes in the matrix are easy to interpret: if all parameters have an identical impact on a specific moment, we should see a value of 2 for each parameter in the corresponding row; if only two parameters matter for a moment, with equal relevance, then we should see a value 11 for those parameters and 0 otherwise, and so on. The results are intuitive and they connect different blocks of parameters to the moments that we would expect. In Table 6 in the main text, we report for each moment all the parameters that have value bigger or equal than 4 in the Jacobian matrix – i.e. that are at least twice as relevant as the average parameter.

Testing the Estimation Routine. To test the effectiveness of our estimation routine, we show that it recovers the true parameters if we target a synthetic set of moments generated by our own model. Specifically, we proceed as follows. First, we extract 11 random vectors of parameters around the estimated one. Second, we use the model to generate 11 sets of moments; one for each of the 11 vectors of parameters. Third, we run 11 separate estimation procedures targeting each set of moments. We keep all the inputs identical across the different estimations, and we follow step by step our main estimation routine. The only difference is that while our main routine consists of 200 separate simulation strings, for this exercise we run only 20 simulation string for each set of parameters (to keep the estimation time relatively brief).

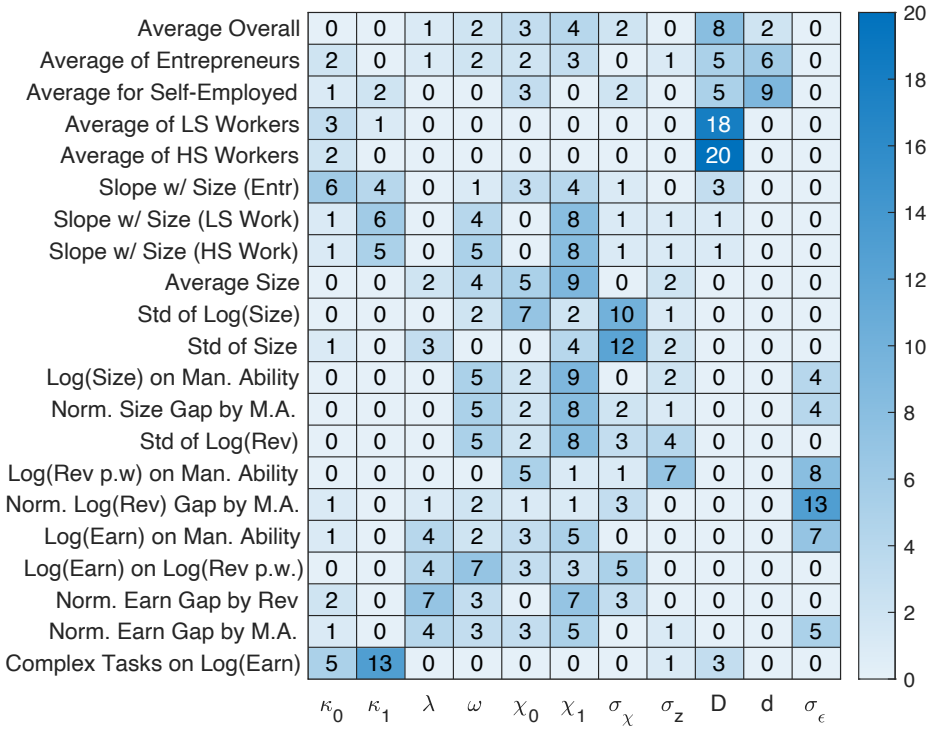
The results from this exercise are shown in Figure C.3, which plots our estimates against the true parameters, together with a 45 degrees line. While the fit is not perfect, overall our estimates are always close to the true values, suggesting that all parameters are very well identified.

Figure C.1: Likelihood Functions around Estimated Parameters



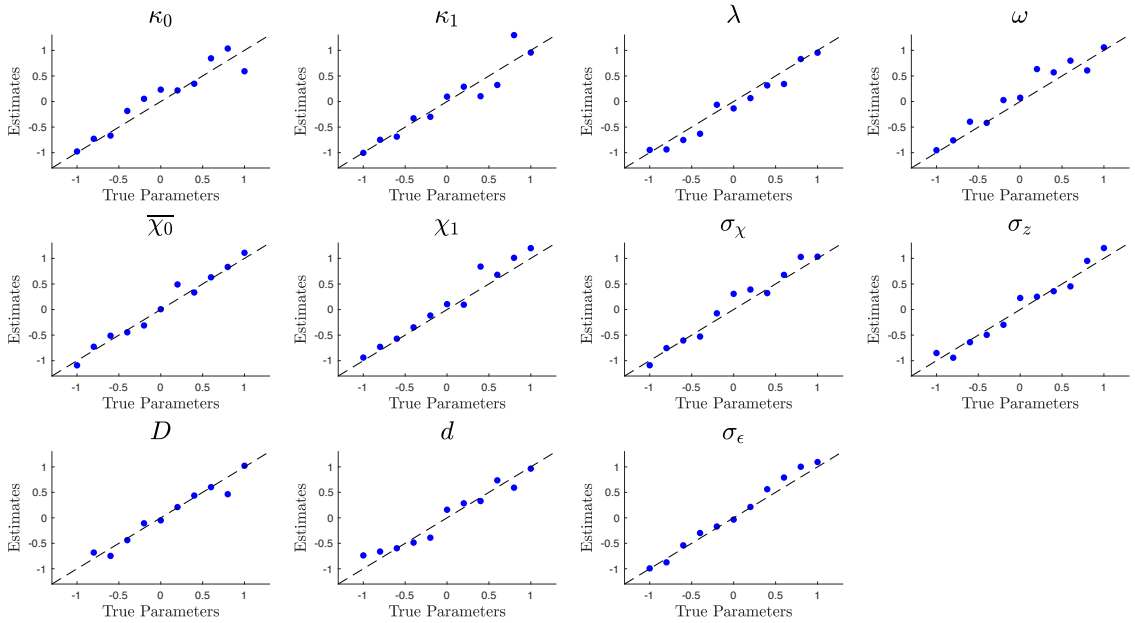
Notes: The figure shows the values of likelihood function for the best 5,000 draws along our estimation chain plotted against each parameter. Each panel is one of the 11 estimated parameters, and the red dotted line shows their estimated values. We notice that the likelihood are single peaked, and achieve their minimum values (i.e. minimized distance between the model and the data), around the estimates.

Figure C.2: Normalized partial derivatives of moments to parameters



Notes: The matrix includes the normalized values of Δ_{jr} computed as described in the text. Each row is a moment (or a summary of moments) and each column is a parameter.

Figure C.3: Model's Fit for Synthetic Estimations



Notes: the figure show the results from 11 separate estimation of the model. We extract 11 random vectors of parameters around the estimated values and compute model-generated moments for each vector. We then run a (shortened) version of our estimation routine targeting the 11 different sets of moments. The figure shows, separately for each parameter, that our estimates track well the true underlying parameters.

C.3 Further Details on Model's Fit

In the main text, Table 6 includes only a summary of all the 150 parameters, and Figure 8 illustrates the model fit for a few key aspects. Tables C.4, C.5, C.6, and C.7 include all the 150 moments. To help visualizing the model's fit, Figure C.4 includes the 80 moments related to the joint distribution of firms' revenues, workers earnings, and managerial ability. While the fit is not perfect, the figure shows that our model is able to broadly account for these key empirical patterns.

Table C.4: Full Moments, Size Distribution

Moments	Data	Model	Moments	Data	Model
A. PDF of Firm Sizes			B. CDF of Firm Size, < Med. Man. Ability		
Size = 1	0.026	0.055	Size =1	0.029	0.07
Size = 2	0.056	0.099	Size =2	0.083	0.164
Size = 3	0.126	0.114	Size =3	0.292	0.303
Size = 4	0.165	0.147	Size =4	0.459	0.475
Size = 5	0.117	0.158	Size =5	0.567	0.644
Size = 6	0.227	0.136	Size =6	0.828	0.78
Size = 7	0.114	0.092	Size =7	0.898	0.865
Size = 8	0.05	0.055	Size =8	0.944	0.913
Size = 9	0.039	0.052	Size =9	0.951	0.95
Size = 10+	0.08	0.092	Size =10+	1	1
			C. CDF of Firm Size, > Med. Man. Ability		
			Size =1	0.026	0.04
			Size =2	0.081	0.145
			Size =3	0.196	0.233
			Size =4	0.361	0.355
			Size =5	0.48	0.502
			Size =6	0.702	0.638
			Size =7	0.822	0.738
			Size =8	0.873	0.8
			Size =9	0.916	0.866
			Size = 10+	1	1

Table C.5: Full Moments, Log Revenue Distribution

Moments	Data	Model	Moments	Data	Model
A. PDF of Log Revenue			B. CDF of Log(Rev), < Med. Man. Ability		
= -1.623	0.03	0	5 pctile	-1.522	-1.008
= -1.389	0.115	0.03	15 pctile	-1.074	-0.745
= -1.154	0.209	0.114	25 pctile	-0.791	-0.576
= -0.919	0.29	0.233	35 pctile	-0.513	-0.454
= -0.685	0.359	0.465	45 pctile	-0.254	-0.325
= -0.450	0.421	0.625	55 pctile	0.012	-0.196
= -0.215	0.451	0.643	65 pctile	0.24	-0.052
= 0.019	0.49	0.518	75 pctile	0.429	0.147
= 0.254	0.527	0.424	85 pctile	0.747	0.417
= 0.489	0.463	0.314	95 pctile	1.320	0.937
= 0.723	0.343	0.33	C. CDF of Log(Rev), > Med. Man. Ability		
= 0.958	0.27	0.267	5 pctile	-1.298	-0.851
= 1.192	0.191	0.235	15 pctile	-0.673	-0.534
= 1.427	0.086	0.063	25 pctile	-0.391	-0.325
= 1.662	0.016	0	35 pctile	-0.184	-0.129
			45 pctile	0.075	0.064
			55 pctile	0.261	0.276
			65 pctile	0.457	0.516
			75 pctile	0.709	0.757
			85 pctile	1.016	0.987
			95 pctile	1.673	1.229

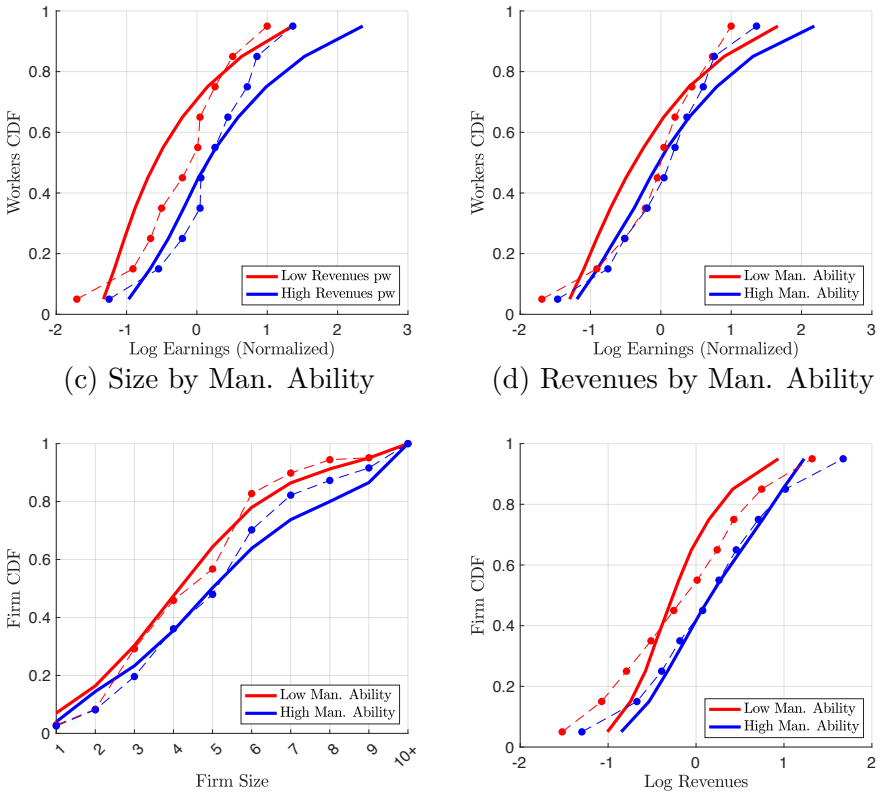
Table C.6: Full Moments, Earning Distribution

Moments	Data	Model	Moments	Data	Model
A. CDF of Earnings, < Med. Rev.			C. CDF of Earnings, < Med. Man. Ability		
5 pctile	-1.707	-1.329	5 pctile	-1.692	-1.297
15 pctile	-0.91	-1.173	15 pctile	-0.91	-1.095
25 pctile	-0.659	-1.031	25 pctile	-0.514	-0.911
35 pctile	-0.502	-0.88	35 pctile	-0.219	-0.716
45 pctile	-0.206	-0.699	45 pctile	-0.05	-0.498
55 pctile	0.012	-0.481	55 pctile	0.045	-0.244
65 pctile	0.045	-0.206	65 pctile	0.201	0.04
75 pctile	0.257	0.149	75 pctile	0.441	0.399
85 pctile	0.508	0.636	85 pctile	0.735	0.9
95 pctile	0.999	1.356	95 pctile	0.999	1.663
B. CDF of Earnings, > Med. Rev.			D. CDF of Earnings, > Med. Man. Ability		
5 pctile	-1.249	-0.971	5 pctile	-1.468	-1.194
15 pctile	-0.546	-0.664	15 pctile	-0.753	-0.9
25 pctile	-0.206	-0.406	25 pctile	-0.514	-0.643
35 pctile	0.045	-0.188	35 pctile	-0.195	-0.38
45 pctile	0.056	0.021	45 pctile	0.045	-0.154
55 pctile	0.257	0.271	55 pctile	0.201	0.095
65 pctile	0.441	0.587	65 pctile	0.37	0.402
75 pctile	0.716	0.985	75 pctile	0.603	0.793
85 pctile	0.854	1.528	85 pctile	0.759	1.308
95 pctile	1.363	2.357	95 pctile	1.36	2.182

Table C.7: Full Moments, Time on Complex Tasks

Moments	Data	Model	Moments	Data	Model
A. Entrepreneur, by Size			B. Low-Skilled Worker, by Size		
Size = 1	0.271	0.38	Size = 2	0.1	0.172
Size = 2	0.336	0.393	Size = 3	0.116	0.175
Size = 3	0.355	0.402	Size = 4	0.154	0.174
Size = 4	0.429	0.416	Size = 5	0.096	0.173
Size = 5	0.531	0.433	Size = 6	0.14	0.173
Size = 6	0.467	0.451	Size = 7	0.161	0.172
Size = 7	0.523	0.473	Size = 8	0.161	0.172
Size = 8	0.507	0.488	Size = 9	0.114	0.17
Size = 9	0.538	0.517	Size = 10+	0.142	0.168
Size = 10+	0.521	0.579	C.High-Skilled Worker, by Size		
			Size = 2	0.194	0.176
			Size = 3	0.248	0.18
			Size = 4	0.191	0.179
			Size = 5	0.237	0.178
			Size = 6	0.233	0.177
			Size = 7	0.205	0.176
			Size = 8	0.227	0.176
			Size = 9	0.188	0.174
			Size = 10+	0.21	0.172

Figure C.4: Model Fit for Distribution of Firm Size, Revenues, and Workers' Earnings



Notes: The figure compares empirical moments (dash lines) with their model-generated counterparts (solid lines). The top two panels show the distribution of worker's earnings separately by firm that have above/below median revenues (left) or above/below median managerial ability index (right). The bottom panels shows, separately for managers that are above/below the median ability index, the distribution of firm sizes (left) and of revenues (right).