The Organization of Work: Time Use and the Returns to Specialization*

Julieta Caunedo[†] Claudia Macaluso[‡] Pamela Medina[§]

April 2026

Preliminary Draft

Abstract

This paper examines how gains from specialization are shaped by the arrival of new technologies and how workers respond to the reorganization of production. We develop a theory in which firms choose how to partition job activities given technology, and workers allocate time across those partitions. Partitions are costly but generate productivity gains through workers' endogenous time allocation. Technological change alters both the returns to partitions and workers' incentives to allocate time within them. More productive firms create more partitions, assign a larger share of activities to specialists, and exhibit greater task concentration within activities. We discipline the model using a novel firm-level survey from Lima, Peru, collected through a digital calendar-style platform. Returns to task specialization depend on firm size. In large firms, wages are approximately 70 log points higher at the 75th percentile of worker specialization, than for the average. In small firms, wages are statistically indistinguishable along worker specialization. Recovering firms' partition structures from workers' joint time use, we find that specialist partitions command substantial within-firm wage premia (beyond worker heterogeneity): the head partition paying roughly 90% more than non-specialist roles. A calibrated version of the model quantifies the response of organizational shifts and worker time reallocation to the arrival of task-displacing technologies, and their effect on aggregate productivity and wage inequality.

Keywords: Time use, Task allocation, Labor productivity, Firm organization, Technology adoption

JEL Classification: J24, M54, O33, D24

*We thank Joaquin Rivadeneyra and Brooke Hansbrough for excellent research assistance. We benefited from comments from various audiences, including Duke University and Georgetown University. The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Richmond or any part of the Federal Reserve System.

†Cornell University. Email: julieta.caunedo@cornell.edu

‡Federal Reserve Bank of Richmond. Email: claudia.macaluso@rich.frb.org

§University of Toronto. Email: pamela.medinaquispe@rotman.utoronto.ca

1 Introduction

The hypothesis that specialization induces productivity gains, and that the essence of economic growth is producing ever more with similar inputs, dates back to Adam Smith (1776). A growing body of work emphasizes that technological progress transforms the content of work, but typically focuses on the extensive margin of tasks performed within the job rather than the intensive margin of time reallocation (Garicano et al., 2026; Freund and Mann, 2025). However, as new technologies arrive, not only firms reorganize production, but also workers of different skills reallocate their time within jobs. In this paper, we study how economic and technological advancement reshapes the returns to different activities within existing jobs and, in turn, how firms and workers reorganize activities and time within jobs to take advantage of technology.

Technological progress often changes not whether a job exists, but how time is allocated across the activities that comprise it. Software developers today spend less time on unspecialized coding and debugging and more time on design, communication, and system-level problem solving than their predecessors. Retail workers similarly devote less time to processing transactions and more time to customer-facing and advisory activities. These examples illustrate a broader phenomenon: technological change reshapes the composition of work within occupations, and this reshaping is "reduced-form" evidence of skill-biased technological change.

Our understanding of within-job time reallocation, however, remains limited. Existing work typically defines jobs by occupational titles and studies whether the tasks they entail are substitutes or complements to technology.[1] We take a complementary approach and focus on how workers allocate and specialize their time across activities, allowing the content of jobs to evolve even when occupational titles remain unchanged. More broadly, we argue that technological change affects productivity not only through the adoption of new technologies, but also through the endogenous reorganization of production within firms.

To study this process, we combine a new theory of the endogenous organization of production with novel data on workers' time allocation. In the model, firms choose how to partition production activities across workers, while workers allocate time across activities within those partitions. The interaction between firm specialization and worker focus determines the organization of production, the distribution of wages, aggregate productivity, and the response of firms to technological change. Our framework micro-founds how workers allocate their time across on-the-job activities, moving beyond the assumption that working

[1] A notable exception is Bandiera et al. (2020), who provide compelling evidence on the time allocation of managers and executives.

hours are fully fungible across tasks and workers. Technologies are chosen by heterogeneous firms operating under decreasing returns to scale. A technology is defined as a partition of production activities that allows workers to focus on a narrower set of activities and thereby increases labor productivity. Partitions are costly: each entails a fixed setup cost and a variable cost that increases with the measure of activities included. Firms therefore optimally choose both the number and breadth of partitions. Workers, in turn, allocate their time across the activities assigned to their partition. Specialization and focus are jointly endogenous outcomes of firms' organizational choices and workers' time-allocation decisions.

To discipline the model, we design, field, and analyze a novel survey that measures on-the-job time use for a broad range of white-collar workers. The survey covers 66 workers across 18 firms in Lima, Peru, between September 2024 and March 2025. It combines a manager survey on firm characteristics and production activities, with a worker survey implemented through a digital calendar-style platform that records detailed task allocation throughout the workday. This approach yields fine-grained information on how workers allocate time across activities, including core activities, non-core activities, and unscheduled time.

We find that focus, defined as concentrating work time in a narrower set of activities, is largely driven by firms' organizational choices. More productive firms adopt more specialized organizational structures, partitioning activities across workers and creating greater scope for specialization. As a result, these firms become larger and generate higher returns to worker focus. Consistent with this mechanism, workers at the 75th percentile of the focus distribution earn roughly 70 log points more than otherwise comparable workers in large firms, while no distinguishable premium exists in small firms.

We document two additional findings on work organization. First, firms organize production into a small number of specialist partitions. Recovering each firm's partition structure from workers' joint time use, we find that firms operate on average about two partitions, including a non-specialist pool. These partitions do not align closely with occupational titles: within a given occupation, time is spread across most business functions.

Second, specialist partitions command substantial wage premia within the same firm. Workers in a firm's head specialist partition earn about 90 percent more than its non-specialists, and each lower-ranked specialist partition pays roughly 0.58 times the wage of the partition above it.

We use the theory and the survey moments to discipline a calibrated quantitative version of the model. The calibration matches moments on task concentration, activity prevalence, wage differences across partitions, and the relationship between firm size and organizational structure. It further allows us to recover key technological parameters, including the productivity gains from partitioning and the fixed and variable costs of specialization.

Armed with the calibrated model, we conduct two counterfactual exercises to understand how technological shocks affect the organization of work. First, we reduce the fixed cost of creating partitions, capturing a development path along which economies become richer and coordination frictions decline. Lower partitioning costs induce firms to climb a "specialization staircase": more firms adopt additional partitions, specialized employment expands, and aggregate output rises. Importantly, these gains are entirely organizational. Growth arises from the endogenous reorganization of activities within firms — a mechanism we refer to as *organizational deepening*.

Second, we consider technological shocks that displace activities at different points of the activity space. Displacement of activities performed by non-specialists induces firms to expand specialization by creating additional partitions. In contrast, displacement of activities performed by specialists leads firms to collapse its partitions back into the non-specialized pool. The aggregate consequences of activity-displacing innovations, therefore, depends not only on how many activities are displaced, but also on where those activities lie within the production process.

Related Literature Our work sits at the intersection of the literatures on economic growth, technology adoption, and firm organization, with an original focus on how firms' choice of structure and workers' time allocation interact and vary with development and technology.

A large body of research studies how work is organized across workers — both total hours, which vary systematically with development and technology (Bick et al., 2018; Aguiar et al., 2021), and the allocation of tasks within firms. On the latter, Bassi, Lee, Peter, Porzio, Sen and Tugume (2023) find that Ugandan manufacturing workers who specialize in fewer tasks are more productive; Caunedo and Kala (2022) show that technology adoption reallocates time across tasks among Indian agricultural workers; and Walker, Shah, Miguel, Egger, Soliman and Graff (2024) document pervasive slack and unscheduled time in developing-country workplaces.

Much less is known about how workers allocate time across activities *within jobs*. The managerial literature is the main exception: Bandiera, Prat, Hansen and Sadun (2018) use diary data on Italian managers to show that much managerial time goes to delegable administrative tasks, and Bandiera, Prat, Hansen and Sadun (2020) link practices that promote specialization and reduce fragmentation to higher productivity. We extend this evidence beyond upper management and connect time allocation to a model of specialization, technology, and firm productivity.

On the tasks-and-technology side, Autor, Chin, Salomons and Seegmiller (2024) show that new job specialties have shifted toward high-paid professional and low-paid service work, and Autor and Thompson (2025) argue that automation's wage effects hinge on whether displaced tasks raise or lower the expertise of the remaining tasks — both consistent with our view that technology acts through the within-job mix of activities. Relatedly, Atencio-De-Leon, Lee and Macaluso (2025) study how turnover shapes task specialization in Peru, and Acemoglu, Akcigit and Johnson (2026) model the interaction of technology adoption and development. Closest to our measurement approach, Bonney, Breaux, Dinlersoz, Foster, Haltiwanger and Pande (2026) document AI diffusion within firms at the level of business functions and worker tasks.

Our organizational mechanism builds on the theory of knowledge hierarchies and the division of labor — the specialization-versus-coordination trade-off of Becker and Murphy (1992), the endogenous knowledge hierarchies of Garicano (2000), and their extension to trade and firm organization by Caliendo and Rossi-Hansberg (2012) — but emphasizes time allocation within *horizontal* activity partitions rather than knowledge assignment across *vertical* layers.[2] We also complement work linking firm organization to technology, such as Guadalupe and Wulf (2010) on trade and corporate hierarchies and Bresnahan et al. (2002) on information technology and workplace reorganization, by observing within-job time allocation directly and using it to discipline the model.

Relative to this body of work, our contribution is to measure time allocation at a finer resolution and to recover, firm by firm, the *partition* structure of activities — distinguishing firm-level specialization (how the activity space is divided) from worker focus (how time concentrates within a partition). This lets us document heterogeneity in partitions across firms. We are then able to show how organizational scale shapes the returns to focus, which provides, to our knowledge, the first direct evidence on this margin.

The remainder of the paper is organized as follows. Section 2 introduces our theoretical framework and discusses its implications for understanding the effects of technology adoption on firm structure and the organization of labor. Section 3 describes our data and methodology. It includes the clustering procedure used to identify empirical partitions, and provides descriptive statistics of our sample. Section 4 analyzes patterns of time allocation and specialization across different activities and presents our main empirical results on the relationship between task concentration and wages. Section 5 presents the calibration of the model and the two counterfactual exercises. Section 6 concludes.

[2] Our contribution is close in spirit to Garicano, Li and Wu (2026) on how AI redraws job boundaries and Freund and Mann (2025) on the labor-market effects of job transformation.

2 A Model of Time Specialization on the Job

2.1 Environment

A continuum of firms, each indexed by productivity z drawn from a distribution $G(z)$ with support $[\underline{z}, \infty)$ and finite mean, produce a homogeneous good — the numeraire — using a decreasing-returns-to-scale technology with returns parameter $\nu < 1$. The economy is populated by $\bar{N}$ workers who inelastically supply one unit of time each.

A *job* consists of a continuum of activities $i \in [0, 1]$, ordered from most to least productive: core tasks lie near $i = 0$ and ancillary tasks near $i = 1$. The function $f(i) \geq 0$ is a technology-specific weight function that determines the relative importance of each activity in the production aggregator and in the time constraint; it need not integrate to one. Different jobs may share the same activity space but differ in f: a unspecialized-intensive job concentrates density at higher i, while a complex job places more weight at lower i.

The function $\tilde{g}(x_i, i)$ maps the time x_i a worker spends on activity i into that activity's productivity contribution. More time on any activity raises its contribution ($\tilde{g}_x > 0$), but with diminishing returns ($\tilde{g}_{xx} < 0$). Activities are ordered so that their ranking is independent of the time devoted to them: for any $i_1 < i_2$, $\tilde{g}(x, i_1) > \tilde{g}(x, i_2)$ for all $x > 0$, making lower-indexed activities strictly more productive at every time budget. Furthermore, the marginal return to time is higher for lower-indexed activities ($\tilde{g}_{xi} < 0$): an additional unit of time yields a larger productivity gain on a core task than on an ancillary one, so that workers optimally concentrate their time budget on the most productive activities within their partition. We also impose the boundary condition $\tilde{g}(0, i) \geq 0$ for all i: zero time on an activity yields a non-negative productivity contribution.

Each worker faces a time budget

$$\int_{\mathcal{A}} x_i \, f(i) \, di \;\leq\; 1, \tag{1}$$

where $\mathcal{A} \subseteq [0, 1]$ is the set of activities to which the worker is assigned — her *partition*.[3]

[3]Our framework contrasts with the knowledge-based hierarchy approach of Garicano (2000) and Caliendo and Rossi-Hansberg (2012), where workers are assigned to *problems* of different difficulty, and with the task-based framework of Acemoglu and Restrepo (2019), where the key margin is which tasks are performed by labor versus capital. Here the endogenous object is the time allocation *within* a worker's assigned partition, and technology operates by reshaping the primitives $\tilde{g}$, f, and the cost of partitioning.

2.2 Production and Partitioning Technology

Let p denote the output price (normalized to one in equilibrium). A firm with productivity z that hires n workers and assigns all of them to the full activity space $[0,1]$ produces

$$y_z \;=\; z^{1-\nu}\left(n\int_0^1 \tilde{g}(x_i,i)\,f(i)\,di\right)^{\nu}. \tag{2}$$

The term inside the parentheses is the worker's efficiency units of labor. We refer to these workers as *non-specialists*, because their assignment spans the full activity space.

Firms can reorganize work by carving the activity space into contiguous partitions $\mathcal{A}_\ell = [a_{\ell-1}, a_\ell]$, assigning a dedicated group of workers to each. A worker in partition ℓ enjoys a productivity premium $\phi_\ell > 1$ relative to a worker in partition 0; we normalize $\phi_0 = 1$ for partition 0 and order $\phi_1 > \phi_2 > \cdots > \phi_L > \phi_0 = 1$, so that the head partition (which contains the most productive activities) yields the largest premium. The premium ϕ_ℓ is a reduced-form stand-in for productivity gains from task concentration — encompassing learning-by-doing, switching costs, and cognitive load reduction that arise when workers specialise on a narrower activity set. The ordering $\phi_1 > \phi_2 > \cdots > \phi_L > \phi_0 = 1$ is an assumption: core activities (low index) benefit most from dedicated attention. The partition structure is summarized by boundaries $0 = a_0 < a_1 < \cdots < a_L = a^{\max}$, with the domain $[a^{\max}, 1]$ constituting partition 0.

With L higher-productivity partitions plus partition 0, firm output is

$$y_z \;=\; z^{1-\nu}\left(\sum_{\ell=0}^{L} n_\ell\,\phi_\ell \int_{\mathcal{A}_\ell} \tilde{g}(x_i,i)\,f(i)\,di\right)^{\nu} \;=\; z^{1-\nu}\left(\sum_{\ell=0}^{L} n_\ell\,\bar{T}_\ell\right)^{\nu}, \tag{3}$$

where $\bar{T}_\ell \equiv \phi_\ell \int_{\mathcal{A}_\ell} \tilde{g}(x_i,i)\,f(i)\,di$ is the efficiency units per worker delivered by partition ℓ.

Partitioning is costly. The firm pays a fixed cost I for each specialized partition it creates, which we interpret as the cost of setting up a new role, team, or dedicated workflow. It also pays a variable cost $c_{\text{prop}} \cdot a^{\max}$ that is proportional to the total range of activities that have been specialized (the size of the specialist domain), which we interpret as the cost of reorganizing workflow across specialist groups.

The firm's full problem is to choose the number of partitions L, the interior boundaries $\{a_\ell\}_{\ell=1}^{L-1}$, the outer boundary $a^{\max}$, and the staffing of each partition $\{n_\ell\}_{\ell=0}^{L}$ to maximize

$$\pi(z) \;=\; p\,z^{1-\nu}\left(\sum_{\ell=0}^{L} n_\ell\,\bar{T}_\ell\right)^{\nu} \;-\; \omega\sum_{\ell=0}^{L} n_\ell\,\bar{T}_\ell \;-\; c_{\text{prop}}\,a^{\max} \;-\; L\cdot I, \tag{4}$$

where p is the price of output (normalized to one in equilibrium) and ω is the market wage per efficiency unit of labor. Let $L^*(z)$ denote the optimal number of specialized partitions chosen by a firm of productivity z. Figure 1 illustrates two cases: a low-productivity firm with $L^*(z) = 0$ (no specialized partitions; all workers are non-specialists) and a high-productivity firm with $L^*(z) = 3$.

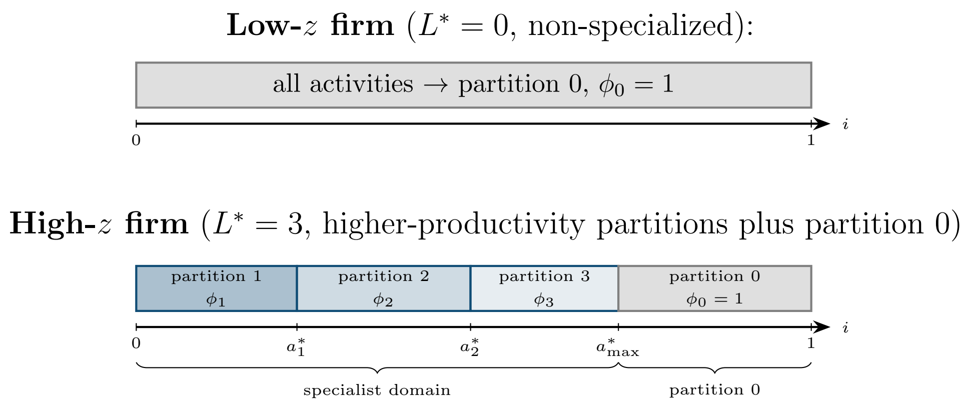

Figure 1: Two examples of the firm's chosen partition structure over the activity space $[0, 1]$. Lower-index activities are more productive and command higher specialization premia.

2.3 Worker's Time Allocation

A worker assigned to partition ℓ takes ω, ϕ_ℓ, and $\mathcal{A}_\ell$ as given and chooses the time profile $\{x_i\}_{i \in \mathcal{A}_\ell}$ to maximize her contribution to output:

$$\max_{\{x_i \geq 0\}} \ \omega\, \phi_\ell \int_{\mathcal{A}_\ell} \tilde{g}(x_i, i)\, f(i)\, di \qquad \text{s.t.} \qquad \int_{\mathcal{A}_\ell} x_i\, f(i)\, di \ \leq \ 1. \tag{5}$$

The Lagrangian is

$$\mathcal{L} \ = \ \omega\, \phi_\ell \int_{\mathcal{A}_\ell} \tilde{g}(x_i, i)\, f(i)\, di - \lambda_\ell \left(\int_{\mathcal{A}_\ell} x_i\, f(i)\, di - 1 \right),$$

and the first-order condition for x_i is

$$\omega\, \phi_\ell\, \tilde{g}_x(x_i^*, i) \ = \ \lambda_\ell, \tag{6}$$

where λ_ℓ is the shadow cost of time in partition ℓ. Because $\tilde{g}_{xx} < 0$, the function $\tilde{g}_x(\cdot, i)$ is strictly decreasing in x and hence invertible for each i; the optimal profile is

$$x_i^* \ = \ \tilde{g}_x^{-1}\left(\frac{\lambda_\ell}{\omega\, \phi_\ell};\, i \right), \tag{7}$$

with λ_ℓ pinned down by the binding time budget.[4]

Because $\tilde{g}_{xi} < 0$, the right-hand side of the FOC is strictly decreasing in i for fixed $c = \lambda_\ell/(\omega\phi_\ell)$, so x_i^* is strictly decreasing in i: workers spend more time on more productive activities within their partition. A first implication is that x_i^* is invariant to all firm- and market-level objects, *conditional on the partition $\mathcal{A}_\ell$ being held fixed*:

Proposition 1 (Conditional invariance of time allocation)**.** *For a given partition $\mathcal{A}_\ell$ and technology $\tilde{g}$, the optimal time profile $\{x_i^*\}_{i\in\mathcal{A}_\ell}$ is independent of the employer's productivity z, the economy-wide wage ω, and the partition premium ϕ_ℓ.*

Proof. The FOC equates $\tilde{g}_x(x_i^*, i) = c$ for all $i \in \mathcal{A}_\ell$, where $c \equiv \lambda_\ell/(\omega\phi_\ell)$. Since $\tilde{g}_{xx} < 0$, the function $\tilde{g}_x^{-1}(c;\, i)$ is strictly increasing in c for each i, so the left-hand side of the time constraint $\int_{\mathcal{A}_\ell} \tilde{g}_x^{-1}(c;\, i)\, f(i)\, di = 1$ is strictly increasing in c; hence c is pinned down uniquely as a function of $(\mathcal{A}_\ell, \tilde{g})$ alone. Since z, ω, and ϕ_ℓ do not appear in this system, x_i^* is invariant to all three. When ϕ_ℓ rises, λ_ℓ adjusts proportionally, leaving c and x_i^* unchanged. □

The conditionality on $\mathcal{A}_\ell$ is important. The partition is itself endogenous: the outer boundary $a^{\max *}$ is determined by a separate optimality condition (equation (10)) in which ω enters through its effect on optimal firm staffing. Interior boundaries $\{a_\ell^*\}$ are pinned down by equation (9), which is ω-invariant within a partition regime because the variable cost $c_{\text{prop}} \cdot a^{\max}$ depends only on the *total* specialist domain, not on how it is divided internally, and ω cancels symmetrically from the revenue terms on both sides. The outer boundary, by contrast, equates the marginal revenue from extending the specialist domain against the marginal variable cost c_{prop}, which is fixed: as ω rises, optimal staffing scales down and the revenue gain from expanding $a^{\max}$ falls below c_{prop} at a lower value of $a^{\max}$. General-equilibrium comparative statics of x_i^* therefore operate through the endogenous adjustment of $a^{\max *}(\omega)$, as characterized in the following proposition.

Proposition 2 (General-equilibrium comparative statics)**.**

(a) Outer boundary. $\partial a^{\max *}/\partial\omega < 0$*: a rise in the equilibrium wage contracts the specialist domain.*

(b) Time allocation in partition 0. *A fall in $a^{\max *}$ widens partition 0, reducing the ratio $c^* = \lambda_0/(\omega\phi_0)$ and hence x_i^* for every $i \in [a^{\max *}, 1]$: workers in the non-specialized*

[4] It is straightforward to show that the firm's problem of assigning a time profile to each worker in partition ℓ — subject to the worker's participation constraint and her ability to deliver the profile — is equivalent to letting the worker choose her profile freely; the Lagrange multiplier plays the role of an implicit transfer price of time, and the boundary of the partition plays the role of an incentive-compatibility constraint. We keep the worker-choice formulation throughout for transparency.

partition spread their time budget more thinly when wages rise.

(c) Time allocation in specialist partitions. *Interior boundaries* $\{a_\ell^*\}_{\ell=1}^{L-1}$ *are* ω*-invariant within a partition regime, so the width of each specialist partition is unchanged by a rise in* ω*, and* x_i^* *for* $i \in \mathcal{A}_\ell$*,* $\ell \geq 1$*, is invariant to* ω *within a regime.*

(d) Discrete reorganization at staircase thresholds. *A sufficiently large rise in* ω *crosses a staircase threshold* $\bar{z}_\ell(\omega)$*, eliminating the marginal specialist partition and discretely widening partition* 0*. This causes a discrete fall in* c^* *and in* x_i^* *for all remaining partition-*0 *workers.*

Proof in Appendix E.2.

Because firm revenue scales proportionally with z, the profit gain from adding one more partition is increasing in z while the fixed cost I is flat. This tension delivers the model's signature result — the specialization staircase — which we characterize next.

2.4 Firm's Optimality and the Specialization Staircase

The firm's first-order condition for staffing in partition ℓ is

$$p\,\nu\,z^{1-\nu}\left(\sum_k n_k\,\bar{T}_k\right)^{\nu-1}\bar{T}_\ell \;=\; \omega\,\bar{T}_\ell, \tag{8}$$

which implies that the marginal revenue product of a unit of labor is equalized across partitions. Hence more productive firms hire more workers in total.

Interior partition boundaries $\{a_\ell^*\}_{\ell=1}^{L-1}$ are determined by a first-order condition for boundary placement. Shifting boundary a_ℓ by da reassigns the activity interval $[a_\ell,\, a_\ell + da]$ from partition $\ell+1$ to partition ℓ, changing effective labor by $\big(n_\ell\phi_\ell - n_{\ell+1}\phi_{\ell+1}\big)\tilde{g}(x^*_{a_\ell}, a_\ell)\, f(a_\ell)\, da$. Dividing by $\tilde{g}(x^*_{a_\ell}, a_\ell)\, f(a_\ell) > 0$ and setting to zero yields:

$$n_\ell^*\,\phi_\ell \;=\; n_{\ell+1}^*\,\phi_{\ell+1}. \tag{9}$$

Because $\phi_\ell > \phi_{\ell+1}$, this implies $n_\ell^* < n_{\ell+1}^*$: more productive (head) partitions operate with fewer workers, consistent with narrow specialist roles at the head of the organization.

The outer boundary $a^{\max *}$ equates the marginal revenue gain from expanding specialization to the marginal partitioning cost:

$$\big(n_L^*\,\phi_L - n_0^*\big)\,\tilde{g}(x^*_{a^{\max}}, a^{\max})\,f(a^{\max}) \;=\; c_{\text{prop}}. \tag{10}$$

Since $\phi_L > 1 = \phi_0$, the firm strictly prefers $a^{\max *} > 0$: assigning the boundary activity to partition L rather than partition 0 strictly raises effective labor. Equation (10) is the first-order condition characterising the interior optimum. Since $\tilde{g}$ is decreasing in i and $\tilde{g}_{xi} < 0$, the gain from expanding $a^{\max}$ is diminishing; and because n_L^* rises with z while c_{prop} is flat, $a^{\max *}(z)$ is strictly increasing in z: more productive firms specialise a larger share of the activity space.

Finally, the firm adds another partition only if the profit gain covers the fixed cost I. Value functions satisfy

$$V_\ell \;=\; \max\!\left\{\, \pi_\ell^*,\; \pi_{\ell+1}^* - I \,\right\}, \tag{11}$$

and because $\pi_{\ell+1}^* - \pi_\ell^*$ scales proportionally with z while I is flat, there is a unique threshold $\bar{z}_\ell$ at which the gain exactly covers I, yielding a monotone sequence $\bar{z}_1 < \bar{z}_2 < \cdots$ at which the firm jumps up by one partition. The resulting optimal number of partitions $L^*(z)$ is a non-decreasing step function of z — the *specialization staircase*, illustrated in Figure 2.

Proposition 3 (Specialization staircase)**.** *(a) $L^*(z)$ is non-decreasing in z, with unique thresholds $\bar{z}_1 < \bar{z}_2 < \cdots$ at which the firm optimally adds one partition. (b) $a^{\max *}(z)$ is strictly increasing in z.*

Proof in Appendix E.2.

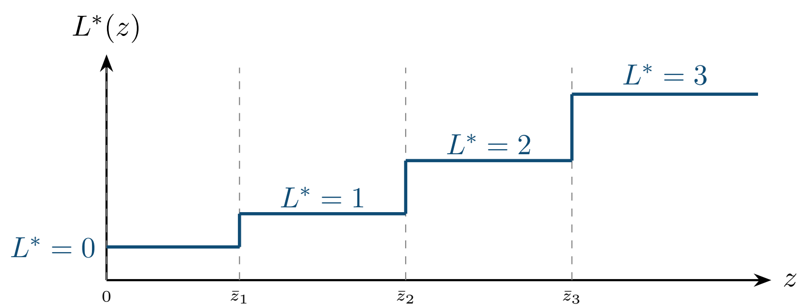

Figure 2: The specialization staircase. More productive firms operate more partitions. Thresholds $\bar{z}_\ell$ at which the firm jumps up by one partition are pinned down by equating the profit gain from adding partition $\ell + 1$ to the fixed cost I.

2.5 Wages

Each worker in partition ℓ is paid her contribution to effective labor, so

$$w_\ell \;=\; \omega\,\bar{T}_\ell, \tag{12}$$

and taking logs yields a clean three-way decomposition:

$$\ln w_\ell \;=\; \underbrace{\ln \omega}_{\text{market}} + \underbrace{\ln \phi_\ell}_{\text{specialization}} + \underbrace{\ln \int_{\mathcal{A}_\ell} \tilde{g}(x_i^*, i)\, f(i)\, di}_{\text{time allocation}} . \tag{13}$$

Three channels drive wages: a common market term, a specialization premium that is highest in the head partition, and a time-allocation term that depends on the width of the partition. Firm productivity z does *not* enter the wage directly. It enters only through the endogenous partition structure $L^*(z)$: more productive firms run more partitions, and their workers sort into higher-$\bar{T}_\ell$ roles. This is the model's mechanism for the firm-size wage premium. The competitive labor market prices efficiency units at rate ω: a worker assigned to partition ℓ earns $w_\ell = \omega \bar{T}_\ell$, the market value of her output. Wage differences within and across firms reflect differences in effective labor productivity $\bar{T}_\ell$ across partition assignments. The firm decides which worker fills each partition slot; since workers are homogeneous and interchangeable, this assignment is without loss of generality, and equilibrium ω is determined by aggregate labor market clearing as in (15).[5]

Proposition 4 (Wage gradient)**.** (*a*) *For any fixed partition structure* $\{\mathcal{A}_\ell, \phi_\ell\}$*, the wage ratio* $w_\ell / w_0 = \bar{T}_\ell / \bar{T}_0$ *is independent of the employer's productivity* z *and of the economy-wide wage* ω*.* (*b*) *At each regime threshold* $\bar{z}_L$ *where* $L^*(z)$ *increases, the domain of partition* 0 *shifts to a strictly higher-index activity set and* $\bar{T}_0$ *falls strictly.*

Proof in Appendix E.2.

For a fixed partition structure, $\bar{T}_\ell$ depends only on the boundaries $\mathcal{A}_\ell$ and the time profiles $\{x_i^*\}$; since x_i^* is invariant to z and ω (Proposition 1), the wage ratio is pinned solely by the partition technology, not by demand conditions. Note that as z increases within a partition regime (fixed L^*), $a^{\max *}(z)$ continues to expand (Proposition 3(b)), so the boundaries shift and wages adjust; the z-independence in part (a) is a cross-firm comparison among firms sharing the same partition structure, not a within-regime time-series statement. Across regime thresholds, $a^{\max *}(z)$ expands, reassigning partition-0 workers to a progressively higher-index activity set $[a_+^{\max}, 1] \subsetneq [a_-^{\max}, 1]$. Since $\tilde{g}$ is strictly decreasing in i, the removed activities are strictly more productive than those remaining; $\bar{T}_0$ falls unambiguously via this composition channel. The wage ratio $w_\ell / w_0 = \bar{T}_\ell / \bar{T}_0$ therefore tends to rise across regime thresholds: the model predicts that the within-firm wage spread is wider at firms that operate more partitions.

[5] Our empirical specifications control for college degree, experience, and occupation fixed effects, which attenuates — though does not eliminate — the confound between worker quality and partition assignment.

2.6 Equilibrium

Let $\bar{z} > 0$ denote the participation threshold. A competitive equilibrium consists of a wage ω, a price p (normalized to one), a participation threshold $\bar{z}$, a partition structure $\{L^*(z), \mathcal{A}_\ell^*(z), n_\ell^*(z), a^{\max *}(z)\}$ for each active firm $z \geq \bar{z}$, and an optimal time profile $\{x_i^*\}_\ell$ such that (i) firms solve (4) given ω, (ii) workers' time profiles solve (5) given their partition assignment, (iii) the labor market clears, $\int_{\bar{z}}^\infty N^*(z)\, dG(z) = \bar{N}$, where $N^*(z) = \sum_\ell n_\ell^*(z)$, and (iv) the free-entry condition holds: $\pi^*(\bar{z}; \omega) = 0$, so that the marginal firm earns zero profit net of all costs, with firms below $\bar{z}$ choosing not to operate.

Let $\bar{S} \equiv \bar{N}^{-1} \int_{\bar{z}}^\infty N^*(z)\, S(z)\, dG(z)$ denote the employment-weighted average of the firm-level organisational term $S(z) \equiv \sum_\ell (n_\ell(z)/N(z)) \bar{T}_\ell$. Organisational costs $c_{\text{prop}} a^{\max *}(z) + L^*(z) I$ are denominated in units of the numeraire good, and the income identity

$$Y \;=\; \omega \bar{N} \bar{S} \;+\; \int_{\bar{z}}^\infty \big[c_{\text{prop}} a^{\max *}(z) + L^*(z) I\big]\, dG(z) \;+\; \int_{\bar{z}}^\infty \pi^*(z)\, dG(z) \tag{14}$$

holds by construction. With one market, Walras' law implies goods market clearing follows from labor market clearing.

Substituting $N^*(z) = \Lambda^*(z)/S(z;\omega)$ and $\Lambda^*(z) = (p\nu/\omega)^{1/(1-\nu)} z$ into $\int_{\bar{z}}^\infty N^*(z)\, dG(z) = \bar{N}$ yields the fixed-point equation for the equilibrium wage:

$$\omega \;=\; p\nu \left[\frac{1}{\bar{N}} \int_{\bar{z}}^\infty \frac{z}{S(z;\omega)}\, dG(z)\right]^{1-\nu}. \tag{15}$$

The term $S(z;\omega)$ depends on ω through the staircase thresholds $\bar{z}_\ell(\omega)$: a higher ω raises the cost of adding partitions, shifting each threshold up and reducing S for marginal firms.

Proposition 5 (Equilibrium existence and uniqueness)**.** *Suppose $G(z)$ has finite mean and $S(z;\omega)$ is bounded and bounded away from zero uniformly in z and ω. Then at least one equilibrium wage $\omega^* > 0$ satisfying* (15) *exists. If additionally $\Phi(\omega) \equiv p\nu\big[\bar{N}^{-1} \int_{\bar{z}} z/S(z;\omega)\, dG(z)\big]^{1-\nu}$ is strictly decreasing in ω — that is, the direct wage elasticity of aggregate labor demand dominates the indirect wage elasticity operating through the staircase thresholds — then the equilibrium is unique.*

Proof in Appendix E.2.

Aggregating output across the firm productivity distribution $G(z)$ yields a Hopenhayn-style aggregation with an additional organizational channel:

$$Y \;=\; \bar{N}^\nu \cdot \left(\int_{\bar{z}}^\infty z\, dG(z)\right)^{1-\nu} \cdot \bar{S}^\nu. \tag{16}$$

The first two terms are standard. The third — the *organizational term* $\bar{S}$ — is new. It is the economy-wide employment-weighted average of $S(z)$, which is non-smooth in z, jumping at each productivity threshold $\bar{z}_\ell$ at which a firm adds a partition. A fall in the fixed partitioning cost I (or in the variable cost c_{prop}) lets more firms climb the staircase, raising $\bar{S}$ and thus aggregate output — with firm-level productivity z unchanged. We refer to this margin for aggregate TFP growth as *organizational deepening*. The derivation of (16) is in Appendix E.1.

2.7 Testable Implications

In equilibrium, more productive firms are also larger (employ more workers), since total employment $N^*(z) = \Lambda^*(z)/S(z)$ is proportional to z within a partition regime. We therefore use firm size as a proxy for productivity z throughout, recognising that the proportionality holds within regimes and is monotone but not exact across regime thresholds.

Implication 1: Specialization staircase. Proposition 3 establishes that $L^*(z)$ is non-decreasing in firm productivity and $a^{\max *}(z)$ is strictly increasing. In the cross section this translates into two predictions: the number of distinct partitions per firm is non-decreasing in firm size, and the share of the activity space assigned to specialist workers rises with firm size.

Implication 2: Within-firm wage premium and cross-firm wage spread. Proposition 4(a) establishes that, for two firms sharing identical partition boundaries $\{\mathcal{A}_\ell, \phi_\ell\}$, the wage ratio $w_\ell/w_0 = \bar{T}_\ell/\bar{T}_0$ is the same regardless of any productivity difference between them. Note that within a partition regime $[\bar{z}_L, \bar{z}_{L+1})$, $a^{\max *}(z)$ varies continuously with z (Proposition 3(b)), so two firms in the same regime do not generally share the same boundaries; the invariance prediction applies to firms with literally identical partition structures, not merely the same count L^*. Because specialist partitions carry premium $\phi_\ell > 1$ and cover lower-index (more productive) activities $\mathcal{A}_\ell \subset [0, a^{\max}]$, while partition 0 covers the higher-index activities $[a^{\max}, 1]$ with $\phi_0 = 1$, the condition $\bar{T}_\ell > \bar{T}_0$ holds whenever the productivity and premium advantages of specialist partitions dominate differences in partition width; under this condition, specialist workers earn a strictly positive within-firm wage premium. Proposition 4(b) adds a cross-firm prediction: at each regime threshold where $L^*(z)$ increases, $\bar{T}_0$ falls strictly as partition 0 absorbs only the least-productive activities; the wage ratio w_ℓ/w_0 therefore widens. Firms that operate more partitions — larger, more productive firms — exhibit greater within-firm wage dispersion. In the cross section this translates into two predictions: the specialist-to-non-specialized wage ratio is the same across firms with identical partition structures regardless of their productivity, and within-firm wage inequality rises with firm size.

Implication 3: Focus premium conditional on firm size. Firms with $L^*(z) = 0$ operate a single partition spanning $[0, 1]$: all workers are non-specialists earning the same wage. Firms above threshold $\bar{z}_1$ assign some workers to specialist partitions, and the maintained condition $\bar{T}_\ell > \bar{T}_0$ implies those workers earn a positive premium. The model therefore predicts that the wage premium for working in a specialist partition is absent at small firms and positive at large firms. In the cross section this translates into a single interaction prediction: individual task concentration is uncorrelated with wages at small firms and positively rewarded at large firms.

Implication 4: Time concentration, firm size, and partition structure. Two distinct channels both predict higher time-use concentration at larger, more productive firms.

Composition channel (Proposition 3): $L^*(z)$ is non-decreasing, so a larger firm assigns a greater share of its workforce to narrow specialist partitions. Workers in specialist partitions cover a strict subset of the activity space $\mathcal{A}_\ell \subsetneq [0, 1]$ and, by the binding time constraint (Proposition 1), allocate their entire unit budget within that subset: concentration is mechanically higher in narrower partitions.

*Within-partition-*0 *channel* (Propositions 3(b) and 1): Within any partition regime, $a^{\max *}(z)$ is strictly increasing, so partition 0 narrows as z rises: its domain $[a^{\max *}, 1]$ shrinks. By the mechanism of Proposition 1 applied to the endogenously narrowing domain established by Proposition 3(b), the shadow cost c_0^* is pinned by the time constraint over this domain; a narrower domain forces the same unit budget onto fewer activities, raising c_0^* and hence x_i^* for every $i \in [a^{\max *}, 1]$. Workers in partition 0 at larger firms within a given regime are therefore more focused, not less.

For firms above the first staircase threshold $\bar{z}_1$, both channels reinforce each other: average time-use concentration is strictly higher at larger firms through both composition and the within-partition-0 effect. At $L^*(z) = 0$ firms, all workers span the full activity space $[0, 1]$ — the widest possible partition — yielding the lowest time concentration in the economy. In the cross section this translates into a monotone prediction: average time-use concentration rises with firm size, both because more workers are assigned to narrow specialist partitions and because partition-0 workers themselves become more focused as the firm grows.

Implication 5: Asymmetric focus response to wages. Proposition 2 characterises how the equilibrium wage ω shapes partition boundaries and, through them, individual time profiles. Part (a) establishes that a higher ω contracts the specialist domain ($\partial a^{\max *}/\partial \omega <$

0); part (b) shows that this widens partition 0, lowers the shadow cost c_0^*, and reduces x_i^* for all partition-0 workers — they spread their time more thinly. Part (c) establishes the complementary result: interior boundaries are ω-invariant within a partition regime, so the time profiles of specialist workers ($\ell \geq 1$) are unaffected by a change in ω within a regime.

These results deliver a differential prediction: in higher-wage environments — whether across countries, sectors, or over time — partition-0 worker time concentration falls while specialist time concentration is unchanged, provided the comparison holds the partition regime fixed. Equivalently, any wage-increasing shock that shifts the equilibrium ω^* reduces focus for partition-0 workers but leaves specialist time profiles intact within their regime. The differential response between worker types (specialist versus partition-0) is the distinguishing testable content of Proposition 2: a uniform wage shock produces heterogeneous focus effects depending on partition assignment. Proposition 2(d) sharpens this into a nonlinearity: at wage levels where a sufficiently large rise in ω crosses a staircase threshold, the marginal specialist partition is eliminated and partition 0 widens discretely, causing a discrete rather than continuous fall in partition-0 worker focus. This predicts kinks in the focus-wage gradient at staircase reorganisation thresholds, testable as a nonlinearity in the cross-sectional relationship between industry or local wage levels and partition-0 worker time concentration. Across sectors or economies with different wage levels, this translates into a differential prediction: partition-0 workers are less focused in high-wage environments while specialist workers' time allocation is unchanged, and the gap between specialist and partition-0 worker focus is wider in low-wage contexts.

3 Data and Methodology

To study time allocation and specialization, we developed a novel survey of workers and managers, which we implemented in Lima, Peru, between September 2024 and March 2025. The goal is to capture fine-grained information about workers' time allocation across tasks for jobs in mid-skill business support occupations, which exist across many industries (manufacturing, services, retail). Front-line service workers are the fastest-growing occupational category worldwide (International Labour Organization, 2025), making fine-grained evidence on their time allocation particularly valuable.

Our measurement strategy addresses three challenges. First, jobs are heterogeneous across firms, industries, and countries. This makes cross-unit comparisons difficult without a common taxonomy. To alleviate this problem, we anchor our survey on the World Bank's Firm-level Adoption of Technology (FAT) taxonomy, which provides a common set of business functions that applies across sectors and firm sizes (Cirera, Comin and Cruz, 2022). Second, even within a firm, standard surveys ask what workers *do* but not how they allocate their time *across* activities; for us, this finer-grained information is necessary to recover partition boundaries and we collect it using a digital calendar-style survey instrument. Third, detailed time-allocation data for front-line service workers are not widely available: existing evidence is concentrated on managers and CEOs (Bandiera, Prat, Hansen and Sadun, 2018, 2020). We, therefore, target a variety of white-collar roles, such as human resources professionals, supply chain managers, or administrative workers, across eight business functions (Business Administration, Operations Planning, Quality Control, Marketing, Payment Methods, Supply Chain Management, Sales Methods, and Financial Services).

3.1 Survey Design and Implementation

Data collection proceeded in three stages to address key measurement challenges in understanding work organization and technology adoption in Peru.

First, establishment managers completed a questionnaire on organizational characteristics, and identified core activities for different worker categories.[6] Our activity classification adopts the taxonomy from the World Bank's Firm-level Adoption of Technology (FAT) survey (Cirera, Comin and Cruz, 2022).[7] Second, we surveyed workers using a digital calendar-style digital platform that captured time allocation at three different frequencies. Workers reported daily activities through an hourly time diary, weekly activities through a half-day diary, and monthly activities over the four weeks of a typical month. This multi-frequency approach allowed us to capture both unspecialized daily tasks and periodic activities that occur less frequently.

Workers reported time allocation at the sub-activity level across predefined activity categories (customized based on manager input), with the option to specify overlapping activities when multiple tasks occurred simultaneously.[8] Workers also provided individual characteristics including wages, education, and tenure. Figure 3 shows an example of the diary interface.

[6] The instrument also covered firm size (employment headcount) and sector of operations.

[7] This taxonomy provides distinct operational domains that firms perform regardless of industry or size. Managers confirmed that 95% of worker time in Peruvian service sector firms could be classified within the existing taxonomy. Mapping each firm's specific set of activities from the FAT framework and validating this set with the relevant manager, we were able to tailor subsequent worker questionnaires to include only tasks actually performed within each firm.

[8] Our measurement strategy also identifies scheduled work time that cannot be assigned to specific primary activities, potentially representing transitions between tasks or periods of reduced activity. This measurement contributes micro-level evidence to debates about bottlenecks and idle time in developing country workplaces (Bassi, Lee, Peter, Porzio, Sen and Tugume, 2023; Walker, Shah, Miguel, Egger, Soliman and Graff, 2024).

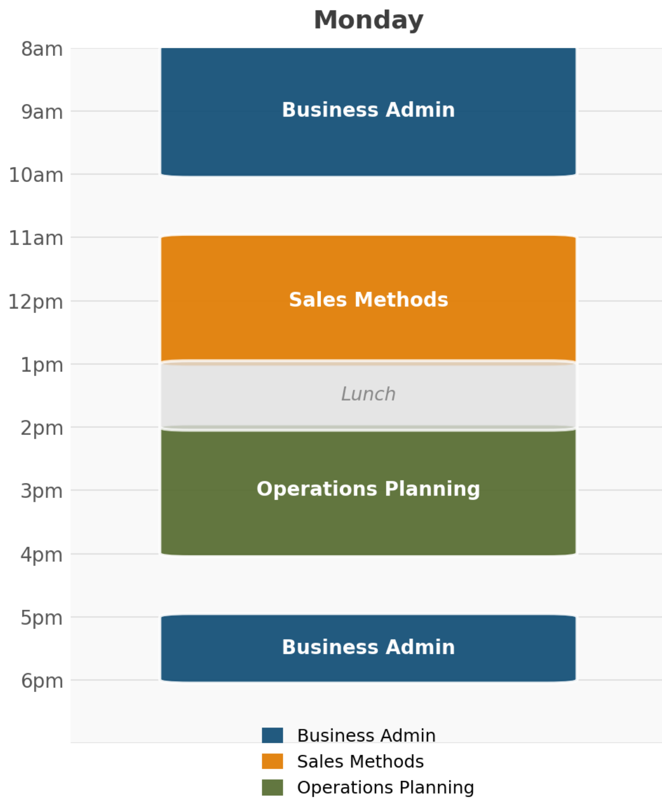

Figure 3: Example of the digital calendar–style worker time-diary interface. Workers entered activities at hourly granularity, with support for overlapping activities and scheduled-but-unallocated time.

Third, and more exploratory, we collected photographic documentation of individual workspaces to identify the equipment workers used in different functions, classified using automated image recognition.[9]

Our multi-stage approach offers several advantages over traditional surveys, but also has limitations. On the one hand, tailoring the worker questionnaire to each firm's specific activities enables workers to report time allocation at the sub-activity level for only the tasks they actually perform, increasing measurement precision and reducing survey costs. On the other hand, this level of granularity was cost-effective because the surveys were primarily conducted through video calls, an innovation that would obviously be inappropriate in contexts where basic technology tools for phone calls are lacking.

[9] This workspace-photo component is a pilot feature of our instrument rather than a central source of the results reported in this paper; we describe it here for completeness.

3.2 Sample Description

Our sample focuses on workers involved in administrative and operational support tasks, including business administration, supply chain management, accounting, finance, human resources, and back-office operations. The final dataset contains 162 observations at the worker-activity level. These observations come from 66 workers across 18 firms in Lima's service sector, spanning four broad occupational categories: Back Office (22.7%), Commercial (13.6%), Operations (36.4%), and Sales and Marketing (27.3%).

Table 1 presents summary statistics for the workers in our sample. The average worker is 33.5 years old, with 14.8 years of total work experience and 3.2 years of tenure at their current firm. Women comprise 41% of the sample, and 56% of workers hold college degrees. The mean monthly wage is 4,128 soles (approximately USD 1,086), though wages exhibit substantial variation with a standard deviation of 2,891 soles.[10]

[10] We constructed standardized wage measures across workers with different reporting frequencies, computing monthly wages by converting weekly and hourly wages to a common scale. Age and tenure were computed as of January 2025 to ensure temporal consistency.

Table 1: Summary Statistics: Worker Characteristics

Variable	Mean	SD	Min	Max
Demographics				
Age (years)	33.5	9.2	18	51
Female (%)	41	-	0	100
Human Capital				
College degree (%)	56	-	0	100
Work experience (years)	14.8	9.5	1	48
Tenure at employer (years)	3.2	2.8	0	12
Work Characteristics				
Hours per week	45.2	8.1	20	60
Monthly wage (soles)	4,128	2,891	0.8	18,000
Task Allocation				
Activities per worker	2.2	1.1	1	6
Task concentration (ITC; 1 = full concentration)	0.42	0.27	0.01	1.00

Notes: N=66. ITC is the worker-level Herfindahl of time across activities, the same measure used in Section 4.2. Two sub-minimum-wage entries (¡10 soles per hour) are data-entry errors; they are excluded from the wage ratios in Section 4.1.1 and the partition-wage appendix.

Workers engage with an average of 2.2 distinct activities. The eight main activity categories show different levels of involvement: Business Administration and Operations Planning are the most common, with 48.5% and 43.9% of workers reporting them on their calendar, respectively. Sales Methods and Financial Services, instead, are reported by a smaller portion of workers (15.2% and 7.6%).[11]

Operational frequencies vary. Core activities are typically reported daily, while support functions like marketing and payment processing weekly. Supply chain activities are often reported and performed at monthly frequency.

[11] Percentages may exceed 100% due to multiple activities per worker.

3.3 Measuring Specialization and Focus

3.3.1 Recovering Empirical Firm-Specific Partitions

To confront the model's partition structure with the data, we recover an empirical per-firm partition map. Our approach is intuitive: the more often workers jointly perform a set of activities, the more likely these activities are to belong to the same partition. We recover partitions from patterns of joint activity at the worker level; as the set of activities assigned to workers varies across firms, so do partitions. Firms need not explicitly identify partitions, nor are these necessarily tightly linked to occupations *per se*.

For each firm f in the survey, we observe the set of activities performed by each worker i, denoted $\mathcal{A}_{if} \subseteq \{a, b, \ldots, h\}$, across the eight FAT business functions. For activities j and k at firm f, let N_{jk}^f denote the number of workers who perform both activities and U_{jk}^f the number who perform at least one:

$$N_{jk}^f = |\{i : j \in \mathcal{A}_{if} \text{ and } k \in \mathcal{A}_{if}\}|$$

$$U_{jk}^f = |\{i : j \in \mathcal{A}_{if} \text{ or } k \in \mathcal{A}_{if}\}|$$

The (Jaccard) co-occurrence measure is

$$S_{jk}^f = \begin{cases} N_{jk}^f / U_{jk}^f & \text{if } U_{jk}^f > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Note that co-occurrence is computed from individual worker behavior rather than firm-level aggregates, preserving within-firm variation in how workers combine activities.

Specialist partitions We treat two functions as belonging to the same specialist partition when their co-occurrence is sufficiently high, $S_{jk}^f \geq \tau$. Forming the graph on the active functions with an edge whenever $S_{jk}^f \geq \tau$, a set of functions constitutes a specialist partition when *every* pair within it clears the threshold — that is, a maximal clique of this graph. When cliques overlap, each function is assigned to the clique with which it co-occurs most strongly, so that partitions are disjoint.[12] In addition, a single function constitutes its own one-activity partition when at least two workers devote their largest time share to it, so that

[12] Cliques are our baseline because they guarantee that all functions inside a partition are jointly performed above the threshold, ruling out "chaining" through a hub activity (e.g. administration or operations planning, which co-occur with almost everything). As robustness we also (i) apply average-linkage agglomerative clustering to $\mathbf{D}^f = \mathbf{1} - \mathbf{S}^f$ and cut the dendrogram at height $1 - \tau$, and (ii) take connected components of the thresholded graph. The first delivers partitions and moments nearly identical to the clique rule; the second is more permissive and we report it as a bound.

genuine dedicated roles are preserved as separate partitions rather than absorbed elsewhere.

We set the threshold to $\tau = 0.5$, the median of the distribution of strictly positive pairwise co-occurrences across multi-worker firms, and report robustness for $\tau \in \{0.33,\ 0.67\}$, its 25th and 75th percentiles.

The residual (non-specialized) partition. Functions that belong to no specialist partition — those that workers perform but not jointly enough to identify a partition, whether scattered one-by-one across workers or combined in idiosyncratic, varying ways — constitute the *residual* set. The residual is the empirical counterpart of the model's non-specialized partition, $\ell = 0$. The threshold rule promotes to a partition only those groups of functions that are jointly performed above τ, so a function is identified as part of a non-specialized partition precisely when its joint-performance pattern is too diffuse to define a specialized one.

Maximal cliques. We recover each firm's partitions from a co-occurrence criterion on the eight FAT business functions, as described above. Two functions are *linked* if $S^f_{jk} \geq \tau$. A *specialist partition* is a *maximal clique* of the resulting link graph: a maximal set of functions in which *every* pair co-occurs at or above τ. Requiring all internal pairs to be linked—rather than mere connectivity—ensures a partition is a tightly co-performed bundle of activities rather than a chain of loosely related tasks; when cliques overlap we make them disjoint by assigning each function to the largest, most tightly linked clique that contains it. Functions belonging to no clique form the residual, non-specialist partition, $\ell = 0$.

Robustness across the threshold τ and the maximal clique rule. The threshold τ controls how tightly functions must co-occur to count as a partition, so we do not fix it arbitrarily. We set the headline value to the median of the positive co-occurrence distribution across multi-worker firms and take the 25th and 75th percentiles as bounds, giving the grid $\tau \in \{0.33, 0.50, 0.67\}$ (headline $\tau = 0.50$). Because a higher τ makes links harder to form, it mechanically shrinks specialist partitions and enlarges the residual: the share of workers in the non-specialist partition rises from 0.11 at $\tau = 0.33$ to 0.57 at $\tau = 0.67$. Raising τ enlarges the $\ell = 0$ set through two channels: more function pairs fall below the link threshold and drop out of every clique, and more workers are reclassified as non-specialist because a worker is assigned to $\ell = 0$ whenever her largest specialist-partition time share falls below τ. Lowering τ works in the opposite direction, absorbing all but the most idiosyncratic functions into specialist partitions and driving the residual toward empty.

We also vary the grouping rule itself, replacing maximal cliques with (i) an average-linkage hierarchical cut at height $1 - \tau$ and (ii) the connected components of the link graph—a more permissive rule that admits chains of pairwise-linked functions. Across the full grid of thresholds and all three rules the qualitative conclusions are unchanged, so we report the clique rule at $\tau = 0.50$ as the headline and the remaining 3×3 combinations as robustness.

Assigning workers and ranking partitions. Finally, we assign each worker to the partition in which she concentrates the largest share of her time. A worker is classified as non-specialized — partition $\ell = 0$ — when the residual is her largest time bucket, or when her largest specialist-partition time share falls below τ; otherwise she is a specialist in her top partition.[13] Finally, within each firm we order the specialist partitions by their mean wage, so that $\ell = 1$ denotes the head (highest-paid) specialist partition, $\ell = 2$ the next, and so on, with $\ell = 0$ the non-specialized partition. We use these firm-specific partition maps as the observational counterpart of the model's $\{\mathcal{A}_\ell\}_{\ell=0}^{L}$.

3.4 Individual Task Concentration

Our primary measure of individual task focus is the Individual Task Concentration index, which captures the concentration of a worker's time allocation across activities:

$$ITC_i = \sum_{k=1}^{K} s_{ik}^2 \tag{17}$$

where s_{ik} represents the share of worker i's time allocated to activity k, with $\sum_k s_{ik} = 1$. The ITC ranges from $1/K$ (equal allocation across K activities) to 1 (complete specialization).

The ITC_i addresses several measurement challenges inherent in multi-frequency time diary data. First, we account for overlapping activities when workers report multiple tasks in the same time period by fractionally allocating time: if a worker reports n activities in a given time slot, each receives $1/n$ of that period's time. Second, we harmonize across reporting frequencies (daily, weekly, monthly) by calculating frequency-specific ITC_i^{freq} and then creating a weighted average based on the fraction of activities reported at each frequency.

4 Empirical Patterns in Organizational Specialization and Worker Focus

In this section, we investigate specialization of organizations, functions, and workers, and we make use of the measures in the previous section to describe how work is organized.

[13]Single-worker firms cannot reveal a co-occurrence pattern, so we treat such a worker's entire activity set as a single partition. There are three such occurrences.

4.1 Partitioning Patterns across Firms

We use the threshold procedure described in Section 3.3.1 to recover a partition map for each of the firms in the sample. The recovered partitions describe how each firm bundles its eight FAT business functions — Business Administration, Operations Planning, Supply Chain Management, Marketing, Sales Methods, Payment Methods, Quality Control, and Financial Services — into specialist teams, with any residual functions delivered by non-specialists.

Most firms operate few partitions. Excluding the residual non-specialist set, the number of specialist partitions per firm ranges from one to three, with a mean of 1.56 (median 1.0, SD 0.70). Ten firms operate a single specialist partition, six operate two, and two operate three. Counting the residual non-specialized set as an additional partition where it is non-empty (six of the eighteen firms) raises the mean to 1.89 (median 2.0). Among the multi-partition firms, the head partition typically groups the most prevalent business functions (administration, operations planning, or payments), while secondary partitions isolate complementary or support activities such as supply-chain management, marketing, or sales.

Functions per partition are concentrated. Across the 28 specialist partitions, the average partition contains 2.21 functions (median 2.0, SD 0.74). A typical specialist partition, therefore, groups two functions that workers at the firm tend to perform jointly. The distribution is right-skewed: four partitions consist of a single function, fifteen of two, eight of three, and only one partition counts four activities. Including the residual sets, partitions average 2.06 functions (median 2.0). Table 2 summarizes our findings, while the full firm-level partition map is reported in Appendix C.1.

Table 2: Partition Structure across Firms ($N = 18$)

	Mean	**Median**	**SD**	**Range**
Partitions per firm (excluding residual)	1.56	1.0	0.70	$[1, 3]$
Partitions per firm (including residual)	1.89	2.0	0.76	$[1, 3]$
Functions per partition (excluding residual)	2.21	2.0	0.74	$[1, 4]$
Functions per partition (including residual)	2.06	2.0	0.78	$[1, 4]$

Partitions are obtained from the threshold procedure of Section 3.3.1 (clique grouping, $\tau = 0.5$) applied to worker-activity co-occurrence within each firm. The residual non-specialized set ($\ell = 0$) is non-empty in six firms; the full firm-level partition map appears in Appendix C.1.

Partitions are not occupation labels. A natural concern is that the recovered partitions might mechanically reproduce occupational categories. They do not. Even within a single occupation, time is spread across six to eight FAT functions, and no occupation maps cleanly onto one. Table 3 reports the distribution of activities for workers in each of the four occupational categories used elsewhere in the paper (Back Office, Commercial, Operations, Sales & Marketing): each row sums to one and the cell entries are activity shares. The top three FAT activities account for between 72% and 79% of each occupation's total worker-activity time (Table 4). An alternative occupational classification using the firm-survey job-family categories yields the same conclusion and is reported in Appendix Table C.2. The empirical partitions are therefore *within-firm* groupings of activities recovered from co-occurrence patterns in worker time use, distinct from any pre-existing occupational classification.

Table 3: Occupation × FAT Activity — Row Shares of Worker-Activity Observations

	FAT business function							
Occupation	Bus. Adm.	Ops. Plan.	Supply Chain	Mktg Mktg	Sales Sales	Pay-ments	QC	Fin. Svcs.
Back Office	.217	.304	.043	.043	.043	.174	.087	.087
Commercial	.400	.067	.067	.067	.133	.267	.000	.000
Operations	.226	.302	.132	.038	.000	.132	.170	.000
Sales & Marketing	.216	.135	.054	.243	.243	.054	.054	.000

Top-5 most frequent specialist partitions across firms (from threshold map):

1. Business Administration alone
2. Marketing + Sales Methods
3. Business Administration + Payment Methods
4. Operations Planning + Supply Chain Management
5. Supply Chain Management + Payment Methods

$N = 128$ worker-activity observations. Each row reports the share of an occupation's worker-activity observations falling in each of the eight FAT business functions. A worker reporting time on k activities contributes k rows. The bottom block lists the five most common specialist partitions across firms recovered by the threshold procedure of Section 3.3.1.

Table 4: Activity Concentration within Occupation: Time Share on Top-k Most Common FAT Activities

Occupation	N_{obs}	N_{acts}	**Top Activity**	**Top-1 %**	**Top-2 %**	**Top-3 %**
Back Office	23	8	Operations Planning	31.5	53.2	72.1
Commercial	15	6	Business Administration	39.3	64.3	78.6
Operations	53	6	Operations Planning	28.7	54.3	72.9
Sales & Marketing	37	7	Business Administration	26.8	52.4	73.2

Top-k % is the fraction of an occupation's total worker-activity time falling in its k most common FAT activities. N_{obs} is the number of worker-activity observations within the occupation; N_{acts} is the number of distinct FAT activities reported.

Table 5: Occupation × Partition Rank — Row Shares of Workers

		Specialists			**Non-specialized**
Occupation	N	ℓ=1	ℓ=2	ℓ=3	ℓ=0
Back Office	15	.667	.133	.000	.200
Commercial	8	.250	.250	.250	.250
Operations	24	.833	.042	.000	.125
Sales & Marketing	18	.444	.333	.000	.222
All	65	.615	.169	.031	.185

$N = 65$ workers with at least one activity assigned. Each row reports the share of an occupation's workers assigned to each partition rank under the threshold procedure of Section 3.3.1 (clique grouping, $\tau = 0.5$). Within each firm, specialist partitions are ranked by mean wage, so ℓ=1 is the head (highest-paid) specialist partition; ℓ=0 is the non-specialized/residual partition. Each worker is assigned to a single partition, so rows sum to one.

4.1.1 Partitions and wages

A natural follow-up question is whether the recovered partitions track within-firm wage differentials. We compute two reduced-form moments, both used as calibration targets in Section 5: the head-specialist-to-non-specialized wage ratio w_1/w_0, which disciplines the specialist premium ϕ_{sp}, and the consecutive partition wage step w_2/w_1, which disciplines the geometric decay ξ.

Wages by partition rank. Within each firm we rank partitions by mean monthly wage, so that $\ell = 1$ is the head (highest-paid) specialist partition, $\ell = 2$ the next, and $\ell = 0$ the non-specialized (residual) partition (Section 3.3.1). Because the threshold rule leaves the residual populated, $\ell = 0$ now contains workers and the specialist premium is identified from our own data rather than only from external sources. Mean wages decline in rank away from the head: non-specialists earn 3,604 soles per month on average, head specialists 5,566, and second-rank specialists 4,902 (Table 6, Panel A). Within firms, the head specialist partition pays on average 1.90 times the non-specialized partition (median 2.00; $N = 5$ firms with both a head specialist partition and non-specialized workers), and each subsequent specialist partition pays about 0.58 of the one above it (median 0.49; $N = 6$). The partition wage step closely matches the geometric-decay parameter $\xi = 0.566$ recovered in the calibration. Both ratios are robust to the partition-grouping rule: at $\tau = 0.5$, w_1/w_0 ranges 1.64–1.90 and w_2/w_1 0.51–0.58 across the clique, average-cut, and connected-component rules. Per-firm detail and a regression-based decomposition (Method C) are in Appendix D and Appendix D.1.

Table 6: Partition Wage Moments by Rank

Panel A: Mean monthly wage by partition rank (soles, pooled)

Rank	**Mean wage**	N_{workers}
Non-specialized ($\ell = 0$)	3,604	12
Head specialist ($\ell = 1$)	5,566	40
Other specialist ($\ell = 2$)	4,902	11

Panel B: Within-firm wage ratios

Moment	**Mean**	**Median**	**SD**	N_{firms}
w_1/w_0 (identifies ϕ_{sp})	1.90	2.00	1.14	5
w_2/w_1 (identifies ξ)	0.58	0.49	0.30	6

Partitions are recovered by the threshold procedure of Section 3.3.1 (clique grouping, $\tau = 0.5$) and ranked within firm by mean monthly wage; $\ell = 0$ is the non-specialized (residual) partition. w_1/w_0 is the head specialist partition's mean wage divided by the non-specialized partition's mean wage, over firms with both; w_2/w_1 is the ratio of consecutive specialist partitions' mean wages. Robustness at $\tau = 0.5$: $w_1/w_0 = 1.88$ (average-cut), 1.64 (components); $w_2/w_1 = 0.56$ (average-cut), 0.51 (components).

4.1.2 Working in a single-activity partition vs. a multi-activity partition

The recovered partitions let us distinguish three margins of specialization at the worker level. First, a worker may sit in a *single-activity partition* (a dedicated one-function role) or a *multi-activity partition* (a bundle of two or more jointly performed functions). Second, a worker may be a *single-activity* or *multi-activity worker*, depending on how many functions she herself performs. Third, her reported activities may stay within a *single partition* or *span multiple partitions*. Table 7 summarizes how workers sort along all three.

The three margins are quantitatively distinct. Most specialist partitions bundle functions: only four of the 28 specialist partitions (covering five workers) are single-activity, while the rest group two-to-four functions. At the worker level, 24.6% perform a single function and the rest several; and while three-quarters of workers stay within one specialist partition, 21.5% spread their activities across two or more.

Table 7: Specialization Margins: Worker Sorting across Partition and Activity Breadth

	N	**Median wage**	**Mean ITC**	**Mean # activities**
Panel A: Specialist workers, by partition type				
Multi-activity partition	48	3,400	0.52	2.52
Single-activity partition	5	3,200	1.00	1.00
Panel B: All workers, by own activity breadth				
Multi-activity worker ($\geq$ 2 functions)	49	3,500	0.44	2.73
Single-activity worker (1 function)	16	3,450	0.88	1.00
Panel C: All workers, by partitions spanned				
Multi-partition ($\geq$ 2 partitions)	14	3,750	0.36	3.21
Single-partition (1 partition)	48	3,350	0.58	2.12
Residual-only	3	3,500	1.00	1.00

N = 65 workers with at least one activity assigned (53 specialists in Panel A). Wages in soles/month; medians are reported because the wage mean is sensitive to a few high earners. ITC is the worker's Herfindahl of time across activities. Partitions are recovered by the threshold procedure of Section 3.3.1 (clique grouping, $\tau = 0.5$).

Time concentration, not pay, is what separates these groups. Multi-partition workers are the least focused (mean ITC 0.36) and the broadest (3.2 functions on average), whereas single-partition workers are more focused (0.58) and single-activity workers are mechanically concentrated (0.88–1.00). Median wages, by contrast, are remarkably flat—between roughly

3,200 and 3,750 soles across every category—so spreading across functions or partitions carries no raw wage penalty; if anything multi-partition workers have the highest median. Within specialists, single-activity partitions are the lower-ranked, lower-paid roles (mean rank $\ell = 1.8$ versus 1.2 for multi-activity partitions; median wage 3,200 versus 3,400): they are isolated marginal functions, such as administration or quality control performed alone, rather than the head of the organization.

Mapping to the model. In the model each worker belongs to exactly one partition ℓ and performs the activities in that partition's interval of the activity space, so the relevant heterogeneity is the *width* of a worker's partition and its rank, not the number of partitions she touches. Partition 0 (the non-specialized partition) covers the widest, least-productive range of activities, while specialist partitions $\ell \geq 1$ are narrower and sit on more-productive (lower-index) activities, with the head $\ell = 1$ narrowest and most productive. The model thus predicts that non-specialized workers spread time over many activities and are weakly focused, while specialists concentrate—exactly the ITC ordering in Panel C (0.36 for multi-partition/non-specialized workers versus 0.58 for single-partition specialists). The empirical multi-partition workers, who have no literal counterpart in a model where each worker occupies a single partition, are best read as the data analogue of the wide partition-0 role: broad, low-focus, non-specialized time use. Finally, the data caution against equating "narrow" with "head": the model's head premium comes from *which* high-value functions a partition occupies, not from how few it bundles, so the small single-activity partitions we observe are low-rank isolated functions rather than the productive apex.

4.2 Returns to Individual Focus

A long literature on the division of labor asks whether workers who specialize on a narrow set of tasks earn more than workers whose time is dispersed across many activities (**??**). The answer is not a priori obvious: time concentration may reflect human-capital-intensive expertise, but it may also indicate routinized work with limited returns to skill. We argue, building on our model's intuition, that the trade-off plausibly depends on firm scale, as larger firms allow finer divisions of labor and richer internal hierarchies (**??**). This section examines, therefore, whether the cross-sectional association between time-use concentration and wages varies with firm size.

We measure within-worker time-use concentration with the ITC, i.e. a worker-level Herfindahl index using the shares of each worker's scheduled time devoted to each activity. When a worker reports multiple activities in the same time cell, each is assigned $1/k$ credit, with k the number of co-occurring activities. The index is bounded by $[0, 1]$, with 1 corresponding to a worker who spends all scheduled time on a single activity and lower values to more dispersed allocations. In our sample the distribution of ITC_w is mildly left-skewed, with mean 0.41, median 0.42, and 75th percentile 0.51.

We estimate by OLS

$$\ln w_i \;=\; \alpha \,+\, \beta_1\, ITC_i \,+\, \beta_2\, L_{f(i)} \,+\, \beta_3\, (ITC_i \times L_{f(i)}) \,+\, \mathbf{X}_i'\boldsymbol{\gamma} \,+\, \varepsilon_i, \tag{18}$$

where ITC_i is the individual task concentration measure (continuous or thresholded at the median or 75th percentile), $L_{f(i)}$ is an indicator for firm f having above-median size, and $\mathbf{X}_i$ collects worker controls: indicators for being young ($\leq$ 26 years), female, college-educated ($\geq$ 9 years of schooling), and high tenure ($\geq$ 8 years on the current job). The coefficient of interest is β_3, which tests whether the wage–concentration gradient differs systematically across firm-size strata. Standard errors are clustered at the worker level.

Table 8 reports the estimates. The first specification makes use of the continuous ITC measure and is illustrated in column (1). Our preferred specification is column (2), which uses an indicator for above-75th-percentile ITC_i. Panel A offers estimates of various controls, and shows negative effects of being young or female, and a positive effect of working in a larger firm and of college education (as expected). Panel B shows the marginal effects of being a high-focus worker in firms of varying size. The focus premium is small and statistically indistinguishable from zero in small firms. In large firms, instead, the effect of task concentration is large and significant at the 5% level ($\widehat{dy/dx} = 0.70$, p = 0.02 for the discrete specification and 0.73, p = 0.07 for the continuous one). High-focus workers earn approximately 70 log points more than otherwise-comparable workers, a premium equivalent to approximately 100%.[14] The predicted log wages in Panel C make the asymmetry concrete: small firms display a small concentration gradient (8.04 versus 8.29), while large firms show a more significant premium (8.22 versus 8.92).

4.3 Reduced-form tests of the model's comparative statics

The model delivers two qualitative predictions that can be confronted with the survey directly. First, the number of specialist partitions is non-decreasing in z, firm-level productivity, and, second, the *shape* of a worker's time profile across the activities of her partition is

[14] The marginal effect of ITC in large firms is $\hat\beta_1 + \hat\beta_3$, with variance $\widehat{\mathrm{Var}}(\hat\beta_1 + \hat\beta_3) = \widehat{\mathrm{Var}}(\hat\beta_1) + \widehat{\mathrm{Var}}(\hat\beta_3) + 2\,\widehat{\mathrm{Cov}}(\hat\beta_1, \hat\beta_3)$. The main-effect and interaction estimates are mechanically negatively correlated in regressions of this form (in column (2), the implied correlation between the two estimates is approximately -0.64.). Therefore, the standard error of their sum is generally smaller than the standard error of either component. Significance tests in interaction models should therefore be conducted on the marginal effect itself, as we do (Imbens and Wooldridge, 2009).

Table 8: Individual Task Concentration and Log Monthly Wages

	(1) Continuous	(2) Above 75th
Panel A: Regression coefficients		
Large firm	0.010	0.178
	(0.306)	(0.202)
Young ($\leq$ 26)	-0.678^{**}	-0.554^{*}
	(0.306)	(0.277)
Female	-0.328	-0.313
	(0.211)	(0.194)
College ($\geq$ 9 yrs educ)	0.341	0.438^{*}
	(0.287)	(0.241)
Panel B: Marginal effects (dy/dx)		
Small firm, high ITC	-0.087	0.245
	(0.483)	(0.238)
Large firm, high ITC	$\mathbf{0.726}^{*}$	$\mathbf{0.700}^{**}$
	(0.388)	(0.284)
Panel C: Predicted log wages		
Small firm, low ITC	—	8.043
		(0.114)
Small firm, high ITC	—	8.288
		(0.208)
Large firm, low ITC	—	**8.221**
		(0.166)
Large firm, high ITC	—	**8.921**
		(0.222)
Observations	128	128
Workers	60	60
R^2	0.301	0.358

Notes. OLS estimates of equation (18). The dependent variable is the log of monthly wages. Column (1) uses ITC_i as a continuous regressor; column (2) replaces it with an indicator equal to one if worker-level ITC_i is at or above the sample 75th percentile. "Large firm" equals one if firm size is above the sample median. Panel B reports the average marginal effect of the ITC regressor on log wages by firm-size stratum. Panel C reports predicted log wages from the estimated model at each combination of firm size and ITC bin; predicted wages for the continuous specification are omitted. Standard errors clustered at the worker level in parentheses. * $p < 0.10$, ** $p < 0.05$.

invariant to firm productivity (conditional on partition rank).

We proxy productivity z three ways: total payroll of surveyed workers, employment (both

Table 9: Reduced-form tests

	Prediction	Headline statistic
P1	Partitions per firm rise in z	ρ = +0.42 (payroll, p=.08); +0.48 (employment, p=.05)
P2	Time profile invariant to z	$b_{\log z}$ = −0.005 (payroll, p=.84); +0.011 (emp., p=.76)

Notes. Headline partitioning: maximal-clique grouping at τ=0.50. 18 firms; P2 uses the 53 workers with $\geq$ 2 activities inside their partition.

Table 10: Threshold robustness (maximal-clique grouping)

τ	ℓ=0 share	P1: $\rho(z, L)$ payroll	P2: $\rho(z, a^{\max})$ payroll	P2: $\rho(z, \text{emp. sh.})$ employment	P2: $b_{\log z}$ payroll
0.333	0.11	+0.50	−0.15	−0.38	+0.005
0.500	0.19	+0.42	−0.27	−0.51	−0.007
0.667	0.57	+0.37	−0.32	−0.50	−0.008

Notes. ρ are Spearman correlations across 18 firms; $a^{\max}$ is the share of the activity space in specialist partitions; "emp. sh." is the share of workers in specialist partitions. P2 is the coefficient on log(payroll) from a regression of the within-partition time HHI on $\log z$ and rank dummies (firm-clustered, $\geq$ 12 clusters). All values reproducible and seed-independent.

the surveyed headcount and register employment from SUNAT firm characteristics), and labor productivity (register sales per worker). We report Spearman rank correlations for the first prediction (P1) and, for the second (P2), the coefficient on $\log z$ in a regression of within-partition time concentration on $\log z$ and partition-rank dummies, with standard errors clustered by firm. Table 9 summarizes; the patterns are robust across partitioning rules and thresholds (Table 10).

The data support the two predictions the calibration leans on. More productive firms operate more specialist partitions (P1): the rank correlation of the partition count with payroll is +0.42 (p=.08) and with surveyed employment +0.48 (p=.05), positive on every proxy and stable across partitioning rules. And conditional on rank, the within-partition time profile does not vary with firm productivity (P2): the coefficient on $\log z$ is economically and statistically zero for payroll (−0.005, p=.84), surveyed employment (+0.011, p=.76), and register employment (−0.001, p=.79). Both are exactly the model's signature that z acts on wages and organization only through the discrete partition margin.

Robustness. Table 10 traces the headline clique rule across the data-driven threshold grid $\tau \in \{0.333, 0.500, 0.667\}$; the average-linkage and connected-component rules give the same qualitative pattern. P1 stays positive across the grid and P2 remains a precise zero on the high-cluster proxies.

5 Calibration and Quantitative Exercises

This section presents the empirical implementation and calibration of the model. Section 5.1 sets up the parametric specifications adopted in the calibration. Section 5.2 describes the calibration strategy, Section 5.3 explain the identification of main parameters and Section 5.4 reports the parameter estimates and moment fit. Sections 5.5 and 5.6 present the two counterfactual exercises.

5.1 Parametric Specifications

We adopt the following parametric forms for the model's primitive functions:

$$\begin{aligned} h(i) &= e^{-\alpha i}, \\ f(i) &= (1+\delta)(1-i)^{\delta}, \\ g(x) &= x^{\beta}, \quad \beta \in (0,1), \\ \phi &= \phi_{\text{sp}} \cdot \xi^{\,\ell-1}, \\ &\phi_{\text{sp}} > 1,\ \xi \in (0,1). \end{aligned}$$

The composite technology $\tilde{g}(x_i, i) = h(i)\, g(x_i) = e^{-\alpha i} x_i^{\beta}$ combines the cross-activity productivity gradient α (head activities are more productive) with the within-activity concavity β (diminishing returns to time on a single task). The task-composition density f front-loads the job onto productive activities, with sharpness governed by δ.

Across partitions within a firm, the specialization premium ϕ decays geometrically at rate ξ, so the head partition ($\ell = 1$) earns the largest premium.

Under these forms the worker's optimal time profile is closed-form, $x_i^* = C_\ell\, e^{-\theta i}$ with $\theta \equiv \alpha/(1-\beta)$ and C_ℓ a partition-specific constant pinned down by the binding time budget. The profile depends on α and β only through the composite θ, so α and β are not separately identified. We therefore fix $\beta = 0.5$ and calibrate θ, recovering $\alpha = (1-\beta)\theta$.

Finally, firms draw productivity z from a Pareto distribution $G(z)$ with tail exponent κ. The tail exponent is central to the model's predictions: it governs the distribution of firm sizes and the prevalence of high-productivity firms that can afford to adopt multiple specialist partitions. A lighter tail (lower κ) concentrates more mass on highly productive firms, raising the mean number of partitions and specialist employment; a heavier tail (higher κ) concentrates mass on smaller firms with fewer partitions.

5.2 Calibration Strategy

We calibrate the model's parameters in three groups, drawing every target except the wage-size elasticity directly from our firm-worker time-use survey. The elasticity comes from the *Encuesta Económica Anual* (EEA), Peru's annual establishment survey.

The Pareto tail κ is set fixed set near the Zipf benchmark for firm size (Axtell, 2001; Gabaix, 2009). Two more are set analytically: the geometric premium decay ξ from the partition wage step and the wage-size elasticity ν from the EEA. The within-activity concavity β is also fixed, as it is not separately identified from the cross-activity gradient α; only the composite $\theta = \alpha/(1-\beta)$ enters. The remaining five $(\theta, \delta, \phi_{\text{sp}}, I_{\text{part}}, c_{\text{prop}})$ are estimated jointly, in general equilibrium, by matching model moments to the survey with a Nelder-Mead unspecialized (1,500 evaluations, equal weights across moments). Each candidate vector is solved for the market-clearing wage ω^* holding the labor endowment $\bar{N}$ fixed, so that ω^* responds to the structural parameters both during estimation and in the counterfactuals. Table 11 summarizes.

Table 11: Calibration strategy: parameters, targets, and sources

Step	**Param.**	**Target / moment**	**Source**
Fixed	κ	Pareto tail, near Zipf benchmark	Axtell (2001); Gabaix (2009)
Pre-set	ξ	Partition wage step $w_2/w_1 = 0.58$	Our survey
	ν	Wage-size elasticity $= 0.37$	EEA
Joint	θ	Median individual task concentration $\text{ITC}_{\text{med}} = 0.42$	Our survey
	δ	Activity-prevalence Herfindahl $\sum_k p_k^2 = 0.16$	Our survey
	ϕ_{sp}	Head-to-partition-0 wage ratio $w_1/w_0 = 1.90$	Our survey
	I_{part}	Avg. partitions per firm (incl. partition 0) $= 1.89$	Our survey
	c_{prop}	Specialization depth vs. firm productivity $= 0.106$	Our survey
Additional moments (untargeted checks)			
	—	Partition-count distribution (1 / 2 / 3 partitions)	Our survey
	—	Median ITC by partition	Our survey
	—	Worker share (head vs. partition 0)	Our survey

Notes: All survey moments use the threshold partitions of Section 3.3.1 (clique, $\tau = 0.5$). Partition 0 is the residual non-specialist partition covering high-index activities $[a^{\max}, 1]$.

The calibration targets map to the model as follows. The premium decay ξ is set to the within-firm wage step between consecutive specialist partitions, $w_2/w_1 = 0.58$ (median 0.49, $N = 6$ firms); this own-data value is essentially identical to the 0.566 implied by external occupation wages, so re-sourcing leaves the decay parameter virtually unchanged. The wage-size elasticity ν is set to the establishment-level gradient of 0.37 from the EEA. Among the jointly estimated parameters, θ governs how sharply a worker concentrates time on her most productive activities and is disciplined by the median individual task-concentration index ($\text{ITC}_{\text{med}} = 0.42$), the same focus measure used in the wage regressions of Section 4.2; δ governs the prevalence of activities across the (common) activity space and targets the pooled activity-prevalence Herfindahl ($\sum_k p_k^2 = 0.16$); the head premium ϕ_{sp} targets the head-specialist-to-partition-0 wage ratio $w_1/w_0 = 1.90$ (median 2.00, $N = 5$); the fixed cost I_{part} matches the mean number of partitions per firm including partition 0 (1.89); and the marginal cost c_{prop} matches the elasticity of specialization depth with respect to firm productivity. The fit of all six matched moments is reported in Table 13.

5.3 Identification

Although the five parameters are estimated jointly, the mapping from parameters to targeted moments is close to one-to-one, for structural reasons. The optimal time profile is $x_i^* \propto e^{-\theta i}$, so the concentration of a worker's time across activities (the median ITC) is a monotone function of θ alone, independent of every other parameter; it identifies θ. The activity-prevalence Herfindahl depends only on the task-composition density f and is increasing in δ, so it isolates δ. The proportional cost c_{prop} governs the specialist domain $a^{\max *}(z)$ through the outer-boundary condition and is identified from how that domain widens with firm productivity.

The two organizational parameters are pinned jointly by two moments. The fixed cost I_{part} sets the productivity thresholds at which firms add partitions, and hence the *level* of the mean number of partitions per firm. The head premium ϕ_{sp} governs the size of the productivity gain from specialization, and hence how strongly partitioning rises with firm size: the *co-variation* captured by the correlation between payroll and the number of partitions. Together, the partition level and its correlation with size separate the cost from the premium.

The within-firm wage ratio w_1/w_0 is reported as an over-identifying check rather than a target. Because the model routes all cross-firm organizational variation through a single deterministic channel, heavily specialized firms leave a thin, low-productivity residual partition (partition 0), so the model overstates the head-to-non-specialist wage gap relative to the data.

5.4 Calibration Results

Estimated parameter values and the moment fit are reported in Tables 12 and 13. The model matches the time-use and activity-composition margins closely: the activity-prevalence concentration is 0.17 in the model versus 0.16 in the data, and the median individual task concentration is 0.34 versus 0.42. The organizational moments are matched less tightly. The mean number of partitions per firm (including the residual non-specialist partition) is 1.30 in the model against 1.89 in the data, and the correlation between payroll and the number of partitions is 0.80 versus 0.42: the model's deterministic staircase ties specialization more tightly to firm size than the data do. The specialization-depth gradient is 0.244 versus 0.106. As an untargeted check, the within-firm head-to-non-specialist wage ratio is 8.78 in the model, above the survey maximum of 3.35. Because the model channels all cross-firm organizational variation through firm productivity, heavily specialized firms leave a thin, low-productivity residual partition and the implied wage gap is overstated. We view the remaining gaps as targets for further refinement.

Table 12: Calibrated parameter values

Parameter	**Interpretation**	**Value**	**How set**
ν	Returns to scale (wage-size elasticity)	0.370	fixed / external (EEA)
κ	Pareto tail exponent (firm size)	1.200	fixed (Zipf benchmark)
ϕ_{scale}	Depth scaling of head premium	1.000	fixed
ξ	Premium decay across partitions	0.580	analytic (survey, w_2/w_1)
θ	Time concentration, $\alpha/(1-\beta)$	4.185	joint (survey, median ITC)
δ	Activity-prevalence sharpness	0.995	joint (survey, $\sum_k p_k^2$)
ϕ_{sp}	Head partition premium	3.147	joint (survey, corr. payroll-partitions)
c_{prop}	Marginal partitioning cost	1.000	joint (survey, $a^{\max}$ gradient)
I_{part}	Fixed partitioning cost	0.489	joint (survey, partitions/firm)

ν, κ, and ϕ_{scale} are fixed; $\beta = 0.5$ is normalized (only the composite $\theta = \alpha/(1-\beta)$ is identified); ξ is set analytically to the within-firm partition wage step w_2/w_1 (Table 11). The remaining five are estimated jointly in general equilibrium. Equilibrium wage $\omega^* = 0.346$.

Table 13: Model fit

Moment	Data	Model
Targeted		
Individual task concentration (median)	0.42	0.34
Activity-prevalence concentration	0.16	0.17
Corr(payroll, partitions)	0.42	0.80
Mean partitions per firm (incl. residual)	1.89	1.30
Specialization depth vs. productivity	0.106	0.244
Untargeted checks		
Head/non-specialist wage ratio w_1/w_0	3.35	8.78
Partition shares (1/2/3 partitions)	.56/.33/.11	.70/.30/.00

Targeted moments are matched jointly in general equilibrium; the two checks are computed but not targeted. The w_1/w_0 benchmark is the maximum head-to-non-specialist ratio observed in the survey.

5.5 Counterfactual 1: Development and Organizational Deepening

Our reduced-form view is that technology shocks map to *shifts in the model's primitives and costs* rather than to direct shocks on worker output. A shift in $\tilde{g}(\cdot, i)$ changes the productivity of activity i; a change in $f(i)$ changes the prevalence of activities; an expansion of the activity interval $[0, 1]$ captures technology-driven creation of new activities; and a fall in fixed organizational costs I_{part} (better management, digital collaboration tools) lets firms climb the staircase.[15]

The first counterfactual builds on this framework and traces a development path by reducing the fixed partitioning cost I_{part}. The interpretation is that richer economies have deeper markets, better management practices, and lower coordination frictions that reduce the per-partition fixed cost. We consider three scenarios: I_{part} reduced by 25%, 50%, and 75%. Aggregate output gains are entirely organizational: firm-level productivity z is held unchanged.

The mechanism is straightforward. As I_{part} falls, the cost of adopting specialist partitions declines, and the threshold productivity $\bar{z}_L$ (Proposition 3) shifts downward. More firms cross these thresholds and climb the specialization staircase. The expanding specialist domain

[15] Several extensions are left for future work: endogenizing the partition premia ϕ_ℓ (for example as a function of a complementary capital technology), allowing the activity space to expand endogenously with technology, and validating the empirical partitions against manager-survey measures of role structure.

within each firm creates demand for specialist workers, who earn a wage premium relative to partition-0 workers. However, as specialist employment expands, more workers fill these premium-wage roles, compressing the average specialist premium through composition effects.

Table 14 quantifies these effects across all three scenarios. Figure 4 reveals the distributional consequences at the 50% reduction level. The left panel shows that output gains concentrate in the mid-productivity range (z around 1), where firms have room to add specialist partitions while remaining below the highest productivity tiers. The right panel displays the wage distribution shift under this scenario (I_{part} falls by 50%): the wage CDF shifts rightward across all quantiles, but the shift is most pronounced at the lower tail where partition-0 wages rise and specialist wages compress, narrowing the wage gap between the bottom and top of the distribution.

Table 14: Counterfactual 1: Fall in I_{part} and organizational deepening

Scenario	**Mean partitions**	**Output**	**Wage premium**	**Specialist emp.**
Baseline	1.30	—	—	28%
I_{part} −25%	1.36	+2.6%	−4.6%	+6 pp
I_{part} −50%	1.45	+5.4%	−9.7%	+13 pp
I_{part} −75%	1.62	+9.2%	−18.5%	+24 pp

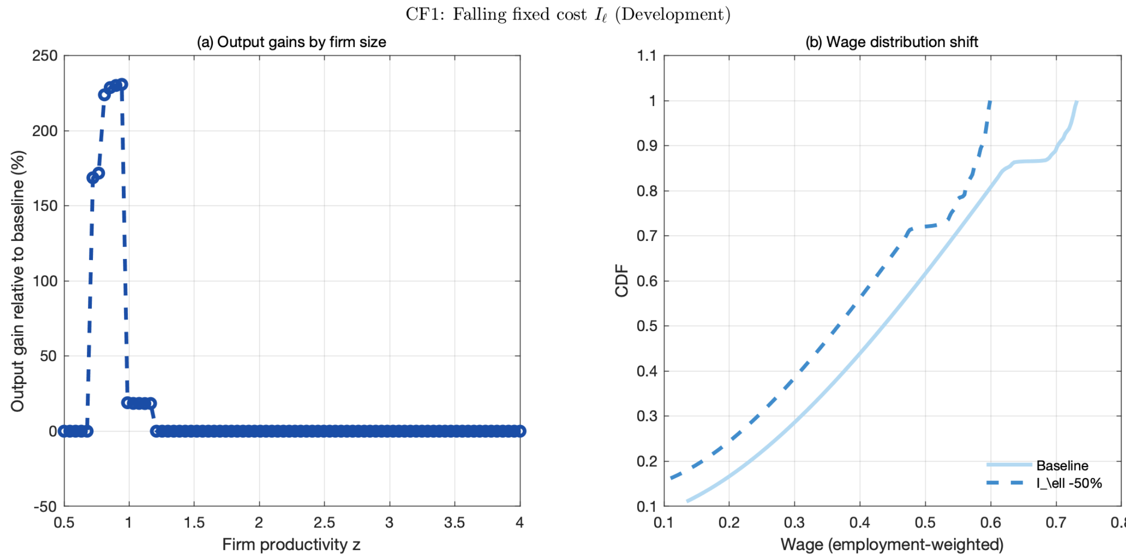

Figure 4: Counterfactual 1. Left panel: distribution of firm-level output gains as fixed costs I_{part} fall by 50%. Gains concentrate at mid-range productivity firms that have room to climb the specialization staircase. Right panel: shift in the employment-weighted wage cumulative distribution under the 50% cost reduction. The distribution shifts rightward, with the largest gains at the lower tail as partition-0 wages rise and the specialist premium compresses.

5.6 Counterfactual 2: Technology Shocks

The second counterfactual considers technological shocks that displace activities at either end of the activity space. unspecialized-task displacement, an automation-like shock that removes high-index (lower-productivity) activities such as assembly-line automation, tractors, or standard administrative workflows, reduces the measure of partition-0 activities and induces firms to *partition up*, expanding the specialist domain. Specialist-task displacement, an expert-system-like shock that removes low-index (higher-productivity) activities such as AI-assisted diagnosis or expert-system automation of skilled-technician work, reduces the measure of specialized activities and induces firms to *de-partition*, collapsing specialist partitions back into partition-0.

Figure 5 summarizes the two mechanisms. In the unspecialized-displacement case (left panel), aggregate output rises as the specialist domain expands and absorbs higher-productivity activities. However, this expansion masks a striking compositional shift in employment: partition-0 employment *falls* as unspecialized activities are automated away, while specialist employment *rises* to fill the expanded specialist domain. This presents an apparent paradox: technology is displacing workers from partition-0 roles, yet aggregate output and specialist employment both increase.

The resolution lies in recognizing that the displaced unspecialized activities are low-productivity.

The firms' response is to reallocate workers from diminished partition-0 pools into specialist partitions where they achieve higher productivity. Figure 6 zooms in on this employment decomposition for unspecialized task displacement across three displacement levels (10%, 20%, 30%). The figure shows the diverging trajectories starkly: as unspecialized tasks are automated, partition-0 shrinks (dark blue bars, left side) while specialist employment expands (light blue bars, right side). The dashed output line reveals the net effect: despite partition-0 losses, aggregate output (right y-axis) rises monotonically from 0.07% at 10% displacement to 0.73% at 30% displacement. This is not job destruction in the conventional sense, but rather the rational reallocation of workers away from low-productivity unspecialized activities toward higher-productivity specialist roles.

In the specialist-displacement case (right panel of Figure 5), the specialization wage premium compresses and specialist employment falls as firms collapse specialist partitions in response to the narrowed specialist domain. The aggregate impact of a task-displacing innovation therefore depends not only on how many activities are displaced but on where in the activity space they lie: a dimension that cannot be captured without explicit within-job time-allocation data. Within our framework, technology acts by shifting the primitives of the model ($\tilde{g}(\cdot, i)$, $f(\cdot)$, the activity range, or the costs I_{part} and c_{prop}), and the counterfactual in this section implements the simplest versions: truncations at the low- or high-i end of the activity space.

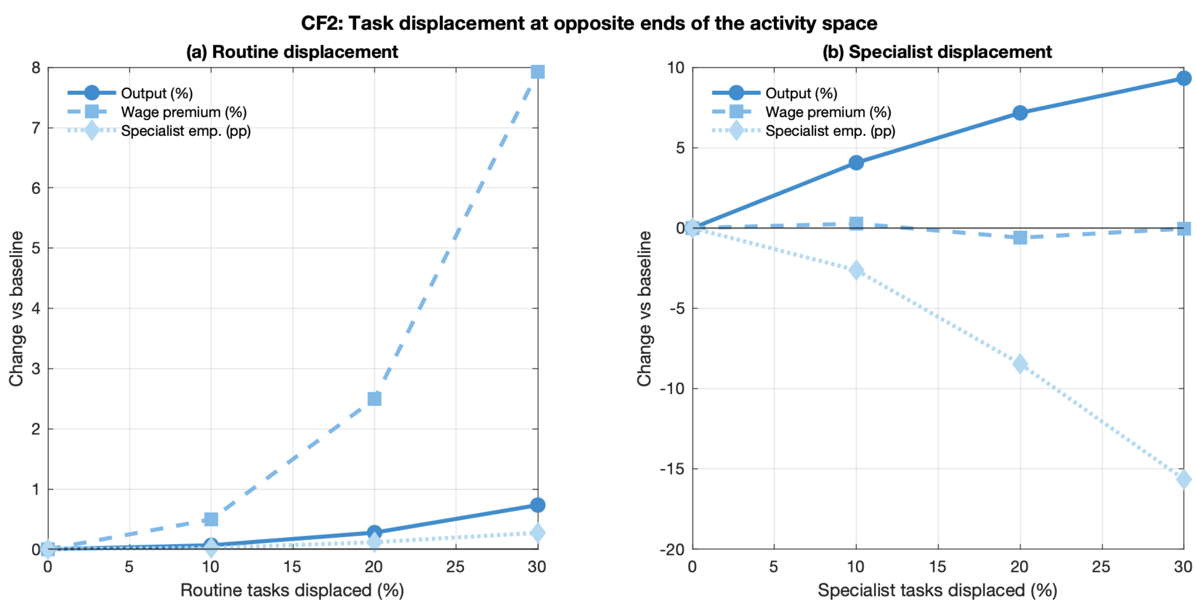

Figure 5: Counterfactual 2. Left panel: technology substitutes for low-productivity (high-index) activities, inducing firms to partition up. Right panel: technology substitutes for high-productivity (low-index) activities, inducing firms to de-partition.

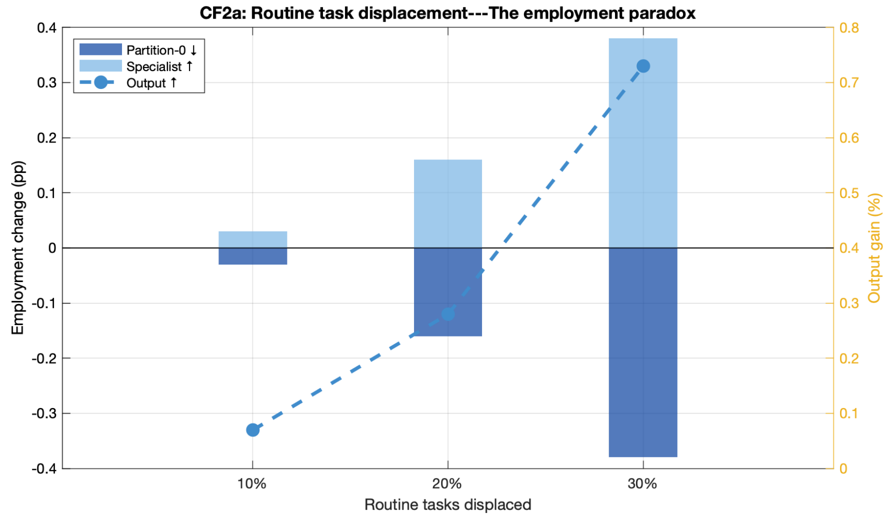

Figure 6: Counterfactual 2a: The employment paradox. unspecialized-task displacement removes low-productivity (high-index) activities. Partition-0 employment falls (dark blue, left bars) while specialist employment rises (light blue, right bars) and aggregate output gains (dashed line, right axis). Workers reallocate from shrinking partition-0 pools to higher-productivity specialist partitions.

6 Conclusion

This paper develops a theory of endogenous organization in which firms partition a continuum of job activities, workers allocate time within partitions, and firm productivity jointly determines firm size, specialization structure, and workers' time allocation. More productive firms pay the fixed cost of creating additional partitions, assign a larger share of activities to specialists, and run narrower specialist partitions. This "specialization staircase" implies that within-job time concentration rises with firm scale, something we verify in the data.

To do so, we collect novel data on workers' time allocation within jobs across a variety of service-sector roles. We field a pilot survey of 66 workers across 18 firms in Lima, Peru, using a digital calendar-style platform, and recover each firm's partition structure directly from workers' joint time use.

We find three new facts on the organization of production. First, firms operate few partitions — about two on average, including the non-specialist pool — and these do not coincide with occupational titles. Second, the partitions carry a within-firm wage gradient: the head specialist partition pays about 90% more than non-specialists, and each lower-ranked partition roughly 0.58 of the one above it. Third, the return to individual *focus* depends sharply on scale: high-focus workers earn roughly 70 log points (about double) more in large firms, while the premium is statistically indistinguishable from zero in small firms.

We discipline a calibrated version of the model using moments from our data, concerning the specialist wage premium, the partition wage step, the degree of individual focus, and the mean number of partitions per firm. We then use the calibrated model for two counterfactuals: first, *development counterfactual* that lowers the fixed partitioning cost. This decline induces firms to climb the specialization staircase, raising aggregate output through organizational deepening, while firm-level productivity is held fixed and the specialization wage premium decreases. Second, a *technology counterfactual* that shows the aggregate effect of various task-displacing innovation. The aggregate consequences depend on *where* technology-driven displacement bites: displacing unspecialized activities prompts firms to partition up, while displacing specialist activities prompts them to de-partition.

References

Acemoglu, Daron and Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

__, Ufuk Akcigit, and Simon Johnson, "Technology and Economic Development," Working Paper 34730, National Bureau of Economic Research 2026.

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst, "Leisure Luxuries and the Labor Supply of Young Men," *Journal of Political Economy*, 2021, *129* (2), 337–382.

Atencio-De-Leon, Andrea, Munseob Lee, and Claudia Macaluso, "Does Turnover Inhibit Specialization? Evidence from a Skill Survey in Peru," *American Economic Review: Insights*, 2025, *7* (1), 56–70.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller, "New Frontiers: The Origins and Content of New Work, 1940–2018," *The Quarterly Journal of Economics*, 2024, *139* (3), 1399–1465.

Autor, David H. and Neil Thompson, "Expertise," Working Paper 33941, National Bureau of Economic Research 2025.

Axtell, Robert L., "Zipf Distribution of U.S. Firm Sizes," *Science*, 2001, *293* (5536), 1818–1820.

Bandiera, Oriana, Andrea Prat, Stephen Hansen, and Raffaella Sadun, "CEO Behavior and Firm Performance," Working Paper 23248, National Bureau of Economic Research 2018.

__, __, __, and __, "The Allocation of Authority in Organizations: A Field Experiment with Bureaucrats," *The Quarterly Journal of Economics*, 2020, *135* (4), 2195–2242.

Bassi, Vittorio, Jung Hyuk Lee, Alessandra Peter, Tommaso Porzio, Ritwika Sen, and Esau Tugume, "Self-Employment within the Firm," Working Paper 31740, National Bureau of Economic Research 2023.

Becker, Gary S. and Kevin M. Murphy, "The Division of Labor, Coordination Costs, and Knowledge," *The Quarterly Journal of Economics*, 1992, *107* (4), 1137–1160.

Bick, Alexander, Nicola Fuchs-Schündeln, and David Lagakos, "How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications," *American Economic Review*, January 2018, *108* (1), 170–199.

Bonney, Kathryn, Cory Breaux, Emin Dinlersoz, Lucia Foster, John Haltiwanger, and Aditya Pande, "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks," Working Paper CES-WP-26-25, U.S. Census Bureau, Center for Economic Studies April 2026.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *The Quarterly Journal of Economics*, 2002, *117* (1), 339–376.

Caliendo, Lorenzo and Esteban Rossi-Hansberg, "The Impact of Trade on Organization and Productivity," *The Quarterly Journal of Economics*, 2012, *127* (3), 1393–1467.

Caunedo, Julieta and Namrata Kala, "Mechanizing Agriculture: Impacts on Labor and Productivity," 2022. Working Paper.

Cirera, Xavier, Diego Comin, and Marcio Cruz, *Bridging the Technological Divide: Technology Adoption by Firms in Developing Countries*, Washington, DC: World Bank Publications, 2022. World Bank Productivity Project, Volume 7.

Freund, Lukas B. and Lukas Mann, "Job Transformation, Specialization, and the Labor Market Effects of AI," 2025. Working paper.

Gabaix, Xavier, "Power Laws in Economics and Finance," *Annual Review of Economics*, 2009, *1* (1), 255–294.

Garicano, Luis, "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 2000, *108* (5), 874–904.

_ , Jin Li, and Yanhui Wu, "Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries," 2026. Working paper.

Guadalupe, Maria and Julie Wulf, "The Flattening Firm and Product Market Competition: The Effect of Trade Liberalization on Corporate Hierarchies," *American Economic Journal: Applied Economics*, 2010, *2* (4), 105–127.

Imbens, Guido W. and Jeffrey M. Wooldridge, "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature*, March 2009, *47* (1), 5–86.

International Labour Organization, *World Employment and Social Outlook: Trends 2025*, Geneva: ILO, 2025.

Smith, Adam, *An Inquiry into the Nature and Causes of the Wealth of Nations*, London: W. Strahan and T. Cadell, 1776.

Walker, Michael W., Nachiket Shah, Edward Miguel, Dennis Egger, Felix Samy Soliman, and Tilman Graff, "Slack and Economic Development," Working Paper 33055, National Bureau of Economic Research 2024.

A Data and Sample

A.1 FAT Taxonomy

Our activity classification is based on the World Bank's Firm-level Adoption of Technology (FAT) survey (Cirera, Comin and Cruz, 2022). FAT was developed under the World Bank's Productivity Project as a harmonised instrument for measuring technology use inside firms across countries at very different levels of development; it has been deployed across more than 35 countries (in Africa, Latin America, Eastern Europe, and Asia) and covers establishments ranging from micro-firms to large multinational subsidiaries.

A central design choice in FAT is to anchor measurement at the level of *business functions* — high-level activities that any firm must perform regardless of sector or scale — rather than at the level of products or industries. The taxonomy decomposes firm activity into eight core business functions: (a) Business Administration, (b) Operations Planning, (c) Supply Chain Management, (d) Marketing, (e) Sales Methods, (f) Payment Methods, (g) Quality Control, and (h) Financial Services. Within each function, FAT distinguishes between *general* (cross-cutting) and *sector-specific* sub-tasks, and asks managers about the technology used to execute each one. This structure makes it possible to compare, e.g., how a Peruvian retail firm processes payments to how a Vietnamese manufacturing firm does so, despite the differences in the underlying product or service.

We adapt the FAT taxonomy as our backbone for two reasons. First, the eight functions provide a common activity space across the heterogeneous firms in our Lima pilot (manufacturing, retail, professional services), which is essential for partitioning the activity space in a way the model can interpret. Managers in our sample confirmed that 95 % of worker time falls within the existing FAT taxonomy, and we tailored each firm's worker questionnaire to the subset of FAT functions that the manager identified as relevant. Second, anchoring on a publicly documented international classification makes the time-allocation moments we report below directly comparable to FAT-derived measures from other countries, which is a precondition for any future cross-country extension of this work.

Figure A.1 reproduces the taxonomy.

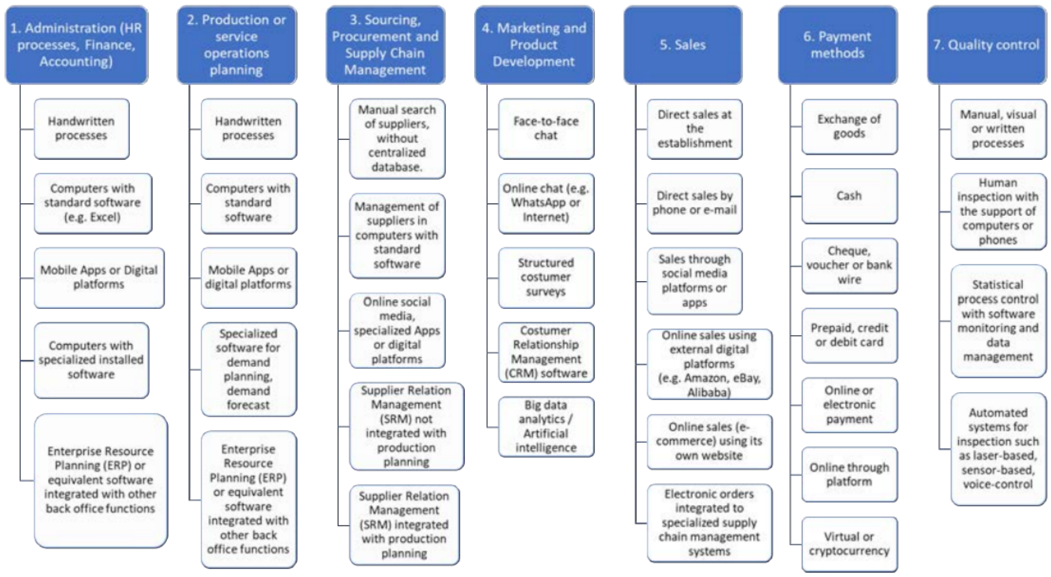

Figure A.1: World Bank FAT taxonomy of eight core business functions. Source: Cirera, Comin and Cruz (2022).

A.2 Sample Composition

Table A.1 reports the distribution of workers and firms across sectors and firm-size bins; this is the descriptive backbone of the survey sample referred to in Section 3.

Table A.1: Sample Distribution by Sector and Firm Size

	Small Firms		Large Firms	
	Workers	Firms	Workers	Firms
Manufacturing	8	2	17	3
Retail	9	3	14	2
Professional Services	6	2	11	2
Other	1	1	1	0
Total	24	8	43	7

A.3 Concentration Measures

Extreme Concentration. To capture the intensity of concentration, we calculate the Cubic Concentration Index (CCI), which applies stronger penalties to extreme concentration:

$$CCI = \sum_{k=1}^{K} s_k^3 \tag{19}$$

where s_k represents the share of time (or workforce) allocated to unit k. The CCI penalizes dominant concentrations more heavily than the HHI. For example, a distribution of (0.5, 0.3, 0.2) yields HHI = 0.38 and CCI = 0.152, but a more concentrated distribution of (0.7, 0.2, 0.1) yields HHI = 0.54 and CCI = 0.351. The CCI increases by 131% while the HHI increases only by 42%. The CCI/HHI ratio provides a normalized measure of concentration intensity, where higher ratios indicate the presence of dominant time slots or activities that capture disproportionate shares.

Entropy. We use a measure of entropy to gauge uniformity of distribution across time slots:

$$H = -\sum_{i=1}^{T} p_i \log_2(p_i) \tag{20}$$

where p_i is the proportion of activity in time slot i. Higher entropy indicates more uniform distribution, with values ranging from 0 (complete concentration) to $\log_2(T)$ (perfectly uniform). For our 10-hour analysis window, maximum entropy is $\log_2(10) \approx 3.32$.

Dispersion. The Coefficient of Variation (CV) measures consistency of labor engagement across time:

$$CV = \frac{\sigma}{\mu} \tag{21}$$

where σ is the standard deviation and μ is the mean of hourly staffing levels. Lower CV values indicate consistent workforce engagement across time periods, while higher values indicate variable staffing with concentration in some periods and reduced presence in others.

Extreme Temporal Peaks. Extreme temporal concentration patterns are captured through the Peak-to-Average Ratio:

$$\textit{Peak-to-Average} = \frac{\max_t(N_t)}{\text{mean}_t(N_t)} \tag{22}$$

where N_t represents staffing or person-hours at time t. Ratios exceeding 2.0 suggest deliberate concentration in specific periods, while ratios near 1.0 indicate a more uniform distribution.

Gap Analysis. Gap Analysis counts time periods with zero staffing:

$$\textit{Gaps} = \sum_{t=1}^{T} \nVdash[N_t = 0] \tag{23}$$

Strategic gaps indicate selective presence rather than continuous coverage of various activities. We distinguish between activities with zero gaps (continuous presence), few gaps (1-3 periods), and many gaps (4+ periods).

Temporal Clustering. Temporal Clustering measures whether work concentrates in continuous blocks or not:

$$Clustering = \mathrm{Corr}(N_t, N_{t+1}) \tag{24}$$

This measures the correlation of staffing levels in adjacent time slots. High positive correlation indicates work clusters in continuous shifts; low correlation indicates dispersed presence across isolated periods.

B Time-of-Day Concentration

Note. The indices in this appendix — the Activity Concentration AC_t and the Time-Slot Concentration TSC_k — measure *when* activities are performed across the hours of the workday (temporal batching). They are distinct from the cross-activity moments used in the calibration: the headline values $\overline{AC} = 0.123$ and $\overline{TSC} = 0.174$ reported here are **not** the calibration targets for δ and θ (which use the activity-prevalence Herfindahl $\sum_k p_k^2 = 0.16$ and the individual task concentration $\overline{\mathrm{ITC}} = 0.42$; see Table 11).

We measure how a firm's workforce is distributed across activities and how each activity is distributed across time using two Herfindahl-style indices. The *Activity Concentration* (AC) index captures how concentrated the workforce is across the eight FAT business functions at a given hour of the workday. For each hour t,

$$AC_t \;=\; \sum_{k=1}^{K} \left(\frac{N_{kt}}{N_t}\right)^2, \tag{25}$$

where N_{kt} is the number of workers performing activity k at time t and N_t is total workers at time t. Higher AC_t means workers cluster on a few activities at hour t; lower AC_t means workers are spread across many activities. We report the workforce-weighted average $\overline{AC} = \sum_t (N_t/N) \cdot AC_t$, where $N = \sum_t N_t$. Under uniform allocation across $K = 8$ FAT activities, $\overline{AC} = 1/K = 0.125$. The *Time-Slot Concentration* (TSC) index captures the complementary question: for each activity k, how concentrated is its execution across hours of the workday? For each activity k,

$$TSC_k \;=\; \sum_{t=1}^{T} \left(\frac{H_{kt}}{H_k}\right)^2, \tag{26}$$

where H_{kt} is total person-hours worked on activity k at time t and $H_k = \sum_t H_{kt}$. Higher TSC_k means activity k batches into a few specific hours; lower TSC_k means activity k is spread evenly across the day. We report the person-hour–weighted average $\overline{TSC} = \sum_k (H_k/H) \cdot TSC_k$, where $H = \sum_k H_k$. Under uniform allocation across $T = 10$ hours of the workday, $\overline{TSC} = 1/T = 0.100$. The two indices view the same data from orthogonal perspectives: AC fixes the time slot and looks across activities; TSC fixes the activity and looks across time slots. The contrast between them quantifies how much firms batch activities temporally (high $\overline{TSC}/0.100$) without changing the workforce composition at any given hour (low $\overline{AC}/0.125$).

Firms Concentrate Business Activities Temporally We find strong evidence that workers organize their time by clustering activities into specific hours rather than spreading activities uniformly throughout the day. Using the AC_t and TSC_k measures defined above, we compare concentration per hour across activities and per activity across time. The comparison reveals significant temporal batching. The weighted AC_t equals 0.123, indicating that at any given hour, workers are spread across diverse activities (equal distribution would be 0.125 across 8 activities). In contrast, the weighted TSC_k is 0.174, showing that activities cluster into specific time periods (an equal allocation would be 0.100). Relative to their uniform benchmarks, the two concentration indexes tell different stories. The AC_t is essentially at its uniform benchmark, indicating that at a given hour workers are spread across the full range of activities. The TSC_k , by contrast, is 74% above its uniform benchmark, indicating that activities cluster into specific hours rather than being spread evenly throughout the workday. The gap is evidence of intentional temporal specialization. Figure B.1 illustrates this temporal concentration through an example: hour-by-hour tracking of back-office workers during core business hours (8am-5pm). The heatmap reveals which activities occur when, with darker shades of purple indicating higher workforce involvement in specific activity-time combinations.[16]

[16]Unsurprisingly, the workforce exhibits a pronounced daily rhythm. Staffing ramps up gradually to peak operations around 10-11am, with a clear lunch-hour contraction in staffing levels. The afternoon then maintains steady staffing levels through 5pm. This pattern reflects deliberate scheduling that concentrates work in morning hours while maintaining thinner crews during low-activity periods such as lunchtime.

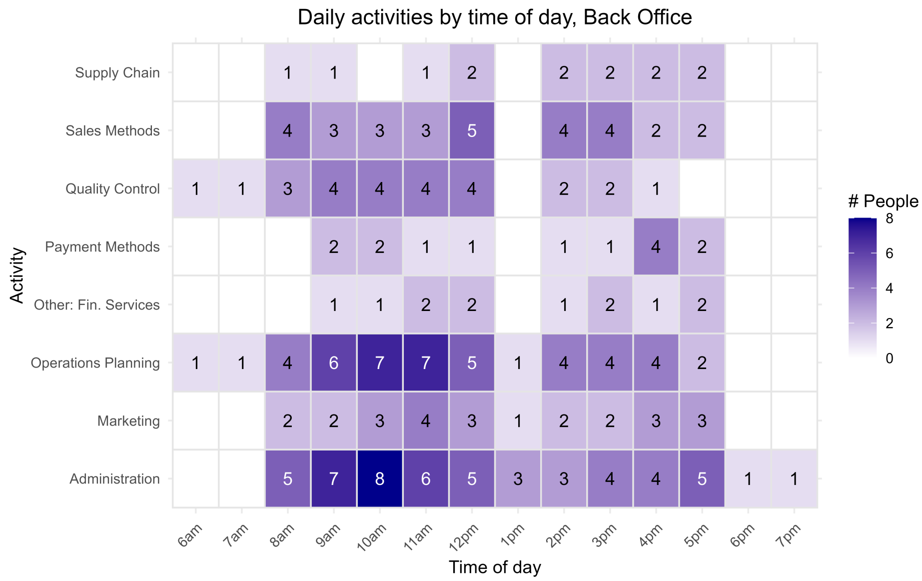

Figure B.1: Distribution of back-office workers across activities by hour of day.

Table B.1 details how workforce composition varies across the day, showing both total staffing and the distribution of workers across activities for each hour. The TSC_k at the bottom quantifies concentration at each hour, revealing when the workforce is most concentrated versus most diversified across activities.

Table B.1: Workforce Composition by Time of Day

Time	8am	9am	10am	11am	12pm	1pm	2pm	3pm	4pm	5pm
Total Staff	19	26	28	28	27	5	19	21	21	18
Activity Share (%)										
Administration	26.3	26.9	28.6	21.4	18.5	60.0	15.8	19.0	19.0	27.8
Operations	21.1	23.1	25.0	25.0	18.5	20.0	21.1	19.0	19.0	11.1
Sales Methods	21.1	11.5	10.7	10.7	18.5	0.0	21.1	19.0	9.5	11.1
Quality Control	15.8	15.4	14.3	14.3	14.8	0.0	10.5	9.5	4.8	0.0
Payment Methods	0.0	7.7	7.1	3.6	3.7	0.0	5.3	4.8	19.0	11.1
Financial Services	0.0	3.8	3.6	7.1	7.4	0.0	5.3	9.5	4.8	11.1
Supply Chain	5.3	3.8	0.0	3.6	7.4	0.0	10.5	9.5	9.5	11.1
Marketing	10.5	7.7	10.7	14.3	11.1	20.0	10.5	9.5	14.3	16.7
TSC_k	0.197	0.178	0.194	0.168	0.150	0.440	0.152	0.147	0.152	0.167

Notes: Percentages show proportion of workforce engaged in each activity at specified time. TSC_k measures concentration of workforce across activities at each hour. Weighted average $TSC_k = 0.174$. Uniform benchmark $TSC_k = 0.100$.

The temporal patterns plausibly reflect strategic organizational choices. The 1pm lunch hour shows high concentration, with fewer workers mostly engaged in administration, plausibly maintaining essential functions while most workers break. This contrasts sharply with the 3pm slot, the most balanced hour of the day, where all activities receive relatively similar staffing. Peak operational hours are, on average, in mid-morning and combine high staffing levels with moderate diversity in activities. The end-of-day period shows increased concentration of workers in payment processing and financial services, likely reflecting closing procedures and daily reconciliation requirements.

Heterogeneous time concentration across functions Are some functions more likely to be time-concentrated across all firms? To understand whether temporal heterogeneity across functions is salient we classify whether observed temporal patterns represent deliberate operational strategies or arbitrary scheduling. To detect intentional patterns, we compare TSC_k to random benchmarks generated through 10,000 Monte Carlo simulations, and further calculate more concentration measures, such as entropy, coefficient of variation, peak-to-average ratios, periods with zero staffing, and correlation of activity levels in adjacent time slots, as well.[17] Table B.2 presents concentration measures for each activity, revealing three distinct patterns. The table reports, in column 2, the TSC_k based on HHI (as in equation 26); in column 3, the TSC_k based on CCI, a concentration index that penalizes deviations from uniformity more than the HHI does; in column 4, the ratio between the two; and pattern classifications and details in columns 5 and 6.

[17]Activities must satisfy multiple criteria to be classified as intentionally batched or distributed, with statistical validation using chi-square tests for deviation from uniform distribution and Kolmogorov-Smirnov tests for deviation from random allocation.

Table B.2: Heterogeneous Temporal Organization by Activity

Activity	HHI	CCI	CCI/HHI	Pattern	Peak Time
Payment Methods	0.163	0.034	0.205	Batched	4pm (28.6%)
Quality Control	0.142	0.022	0.152	Batched	9am-12pm
Financial Services	0.139	0.021	0.150	Mixed	Multiple peaks
Supply Chain	0.136	0.020	0.144	Mixed	Afternoon
Sales Methods	0.120	0.015	0.128	Mixed	12pm (16.7%)
Operations Planning	0.118	0.015	0.129	Distributed	10am-12pm
Marketing	0.110	0.013	0.119	Distributed	11am (16.0%)
Administration	0.110	0.013	0.119	Distributed	Throughout day

Notes: HHI and CCI calculated based on hourly distribution of activities. Equal allocation baseline HHI = 0.100. CCI/HHI ratio indicates presence of dominant time slots; higher values suggest greater concentration. Pattern classification based on temporal clustering analysis comparing observed to 10,000 random allocations.

We find that firms opt for intentional batching for Payment Methods and Quality Control, which exhibit concentrated temporal patterns that significantly exceed random allocation. On the other hand, there is intentional distribution for Administration, Marketing, and Operations Planning, that all show concentration values significantly below their random expectation. These activities maintain a continuous presence throughout the workday with no gaps in coverage and likely serve coordinating or support functions that require ongoing availability. Other activities, like Finance and Supply Chain Management, are mixed and their organization differs significantly across firms that they resist a unique characterization.[18]

Discussion Temporal batching is consistent with our model's prediction that firms organize business functions into partitions to maximize productivity gains from specialization. The gap between a near-uniform AC_t and a +74% TSC_k suggests that the costs of creating partitions — both fixed setup costs and variable costs proportional to the measure of

[18] Among batched functions, Payment Methods shows the most extreme concentration with a peak-to-average ratio of 2.86, allocating 28.6% of all activity to 4pm. Quality Control clusters 67% of its activity between 9am and 12pm, perhaps when most workers' attention levels are highest for detail-oriented work. Both activities show strategic gaps (2 time periods each with zero staffing), rather than continuous coverage. The high CCI/HHI ratio for Payment Methods (0.205) confirms the presence of dominant time slots. On the other hand, Administration shows the highest entropy (0.98) and lowest CV (0.31), reflecting deliberate distribution to ensure continuous support across all business hours. The low CCI/HHI ratio (0.119) indicates genuinely uniform distribution rather than artificially balanced schedules. Statistical validation confirms these patterns are intentional. All activities deviate significantly from uniform distribution (chi-square test, $p < 0.05$) and from random allocation (Kolmogorov-Smirnov test, $p < 0.05$).

activities — are sufficiently low relative to productivity gains ($\phi_\ell > 1$) to justify the chosen organizational structure. The heterogeneous temporal strategies across activities also align with our model's prediction that partition boundaries are determined by the density of time requirements. Activities with higher time intensity are more likely to be batched, as the productivity gains from specialization outweigh the costs of creating separate partitions. The mixed strategies we observe suggest firms face binding constraints on the number of partitions they can profitably maintain, consistent with the model's fixed cost I per partition. In general, we see our evidence as corroborating the hypothesis that the degree of time concentration within jobs reflects firms' optimal responses to the trade-off between specialization gains and partitioning costs. More productive firms find it profitable to create specialized roles where workers focus intensively on narrower activity sets, generating the observed variation in temporal patterns across activities and firms.

Table B.3: Activity Distribution Across Time - Row Proportions

Activity	**8am**	**9am**	**10am**	**11am**	**12pm**	**1pm**	**2pm**	**3pm**	**4pm**	**5pm**	**N**
Supply Chain	.077	.077	.000	.077	.154	.000	.154	.154	.154	.154	13
Sales Methods	.133	.100	.100	.100	.167	.000	.133	.133	.067	.067	30
Quality Control	.125	.167	.167	.167	.167	.000	.083	.083	.042	.000	24
Payment Methods	.000	.143	.143	.071	.071	.000	.071	.071	.286	.143	14
Financial Services	.000	.083	.083	.167	.167	.000	.083	.167	.083	.167	12
Operations	.091	.136	.159	.159	.114	.023	.091	.091	.091	.045	44
Marketing	.080	.080	.120	.160	.120	.040	.080	.080	.120	.120	25
Administration	.100	.140	.160	.120	.100	.060	.060	.080	.080	.100	50

Table B.4: Weighted Concentration Analysis

Dimension	Unweighted HHI	Weighted HHI	Weighted CCI	Weight Share	Interpretation
By Activity					
Supply Chain	0.136	0.008	0.001	6.1%	Moderate concentration
Sales Methods	0.120	0.017	0.002	14.2%	Low concentration
Quality Control	0.142	0.016	0.002	11.3%	Morning batching
Payment Methods	0.163	0.011	0.002	6.6%	Highest concentration
Financial Services	0.139	0.008	0.001	5.7%	Multiple peaks
Operations	0.118	0.024	0.003	20.8%	Distributed pattern
Marketing	0.110	0.013	0.002	11.8%	Most uniform
Administration	0.110	0.026	0.003	23.6%	Continuous presence
Total	-	**0.123**	**0.017**	100%	Low overall concentration
By Time Slot					
8am	0.197	0.018	0.004	9.0%	Morning ramp-up
9am	0.178	0.022	0.005	12.3%	Building diversity
10am	0.194	0.026	0.006	13.2%	Peak hour
11am	0.168	0.022	0.004	13.2%	Balanced peak
12pm	0.150	0.019	0.003	12.7%	Most diverse
1pm	0.440	0.010	0.005	2.4%	Lunch skeleton
2pm	0.152	0.014	0.002	9.0%	Afternoon recovery
3pm	0.147	0.015	0.002	9.9%	Lowest concentration
4pm	0.152	0.015	0.003	9.9%	Payment processing
5pm	0.167	0.014	0.003	8.5%	End-of-day
Total	-	**0.174**	**0.037**	100%	Moderate concentration

C Empirical Partitions

This appendix collects the detailed firm-level material underlying the partitioning patterns reported in Section 4.1. The eight FAT business functions are coded (a) Business Administration, (b) Operations Planning, (c) Supply Chain Management, (d) Marketing, (e) Sales Methods, (f) Payment Methods, (g) Quality Control, (h) Financial Services.

C.1 Per-Firm Partition Map

Table C.1 reproduces the partition assignment for each of the 18 firms, obtained from the threshold procedure of Section 3.3.1 (clique grouping, $\tau = 0.5$). P1, P2, … denote specialist partitions (ordered by mean wage); R denotes the residual non-specialist set, which—unlike under a forced clustering—is populated by performed-but-scattered functions. Single-worker firms (whose activity set is treated as one partition) are flagged.

Table C.1: Partition Map by Firm (threshold method; specialist partitions and residual)

Firm	**1 worker**	**Partitions and residual**
201013945741		P1:{a,b} | R:{f}
201017925451		P1:{a,b,g} | P2:{c,f}
201715456771		P1:{b,c,f,g} | P2:{d,e} | P3:{a}
203284954741		P1:{a,e,f} | R:{b}
203297254311		P1:{a,b,f} | P2:{c,g}
203524668173	✓	P1:{b,c}
203524668175	✓	P1:{a}
204526147671		P1:{b,c} | P2:{a} | R:{d,e}
204526147672		P1:{d,e}
204526147674		P1:{a,f}
205530148411		P1:{b,d,g} | P2:{a} | P3:{c,f}
206015943241		P1:{b,c,g} | P2:{a,f} | R:{d}
206027584871		P1:{b,c,g} | R:{a,f}
206040903511	✓	P1:{a,f}
206040903512		P1:{a,b,d} | R:{g}
206040903513		P1:{a,b,c}
206071670961		P1:{b,h} | P2:{d,e}
206095632661		P1:{b,d} | P2:{e,f}

C.2 Activity Distribution by Job Family (CM Classification)

The CM job-family classification provides a finer occupation-to-activity mapping than the four-category occupational scheme used in the main text (Table 3). Table C.2 reports the analogous row shares. Several CM categories show tighter concentration than the four-category scheme (e.g., *Gerente operaciones* on Operations Planning, *Ventas y marketing* on Marketing and Sales), but most still touch multiple FAT functions, reinforcing the conclusion that the recovered partitions are not occupation labels.

Table C.2: Job Family (CM Classification) × FAT Activity — Row Shares

	FAT Activity Code							
Job family (CM)	a	b	c	d	e	f	g	h
(unclassified)	.278	.167	.056	.056	.167	.278	.000	.000
IT	.000	1.000	.000	.000	.000	.000	.000	.000
Arquitecto	.000	.333	.167	.167	.000	.000	.333	.000
Contador	.333	.250	.000	.000	.000	.333	.000	.083
Control riesgo	.250	.250	.000	.000	.250	.000	.000	.250
Gerente operaciones	.184	.368	.132	.053	.000	.105	.158	.000
Trab. comerciales y admin.	.389	.056	.111	.056	.056	.167	.167	.000
Ventas y marketing	.226	.129	.065	.258	.226	.032	.065	.000

N = 128 worker-activity observations. Row shares of worker-activity observations across the eight FAT business functions (a)–(h), defined as in Table 3. The CM classification is constructed from the firm-survey free-text job titles.

D Partition Wages

Table D.1 reports, for each firm, the mean monthly wage in the non-specialized partition ($\ell = 0$) and the head and second specialist partitions ($\ell = 1, 2$), together with the within-firm ratios w_1/w_0 and w_2/w_1 that enter the calibration. Partitions are recovered by the threshold procedure of Section 3.3.1 (clique grouping, $\tau = 0.5$) and ranked within firm by mean wage. The means across firms are $w_1/w_0 = 1.90$ (median 2.00, $N = 5$) and $w_2/w_1 = 0.58$ (median 0.49, $N = 6$), as reported in Table 6.

Table D.1: Per-Firm Mean Wage by Partition Rank (threshold method)

Firm	N	$\bar{w}_{\ell=0}$	$\bar{w}_{\ell=1}$	$\bar{w}_{\ell=2}$	w_1/w_0	lw_2/w_1
201013945741	5	1520	3925	–	2.58	l–
201017925451	4	–	1750	–	–	l–
201715456771	5	5000	10000	9000	2.00	l0.90
203284954741	4	–	3500	–	–	l–
203297254311	2	–	9999	4000	–	l0.40
203524668173[a]	1	–	2000	–	–	l–
203524668175[a]	1	–	2000	–	–	l–
204526147671	11	4174	14000	4541	3.35	l0.32
204526147672	2	–	29000	–	–	l–
204526147674	2	–	5150	–	–	l–
205530148411	5	–	3250	3200	–	l0.98
206015943241	5	–	5500	3200	–	l0.58
206027584871	4	4000	2167	–	0.54	l–
206040903511[a,b]	1	–	3	–	–	l–
206040903512	3	1752	1800	–	1.03	l–
206040903513	2	–	3175	–	–	l–
206071670961	5	–	6300	–	–	l–
206095632661	3	–	5000	1450	–	l0.29
Across firms (mean)					1.90	l0.58

Wages in soles/month. $\bar{w}_\ell$ is the mean wage of workers assigned to rank ℓ; "–" indicates no worker assigned to that rank in the firm. w_1/w_0 and w_2/w_1 are within-firm ratios, averaged across firms where defined: w_1/w_0 mean 1.90 (median 2.00, N=5); w_2/w_1 mean 0.58 (median 0.49, N=6).

[a] Single-worker firm: the worker's full activity set is treated as one partition.

[b] Reported `wage_monthly`=3 soles is almost certainly a data-entry error (Peru's minimum wage is ≈1,025 soles/month); the worker enters no ratio.

D.1 Method C: Regression-Based Specialist Premium

As a regression analogue we estimate the specialist log-wage premium, comparing specialists ($\ell \geq 1$) to non-specialists ($\ell = 0$). Because the rank ordering of specialist partitions is defined by wage, we use the specialist/non-specialist split (which is determined by time concentration, not wages) to avoid mechanical circularity:

$$\ln w_i = \alpha + \beta\,\mathbf{1}[\text{specialist}_i] + \mathbf{X}_i'\gamma\ (+\,\mu_{f(i)}) + \varepsilon_i,$$

with controls $\mathbf{X}_i$ for age, age^2, female, and firm size (between-firm), and firm fixed effects $\mu_{f(i)}$ (within-firm).

Table D.2: Method C: Specialist ($\ell \geq 1$) vs. non-specialist ($\ell = 0$) Log-Wage Premium

	(1) Between-firm	**(2) Within-firm (FE)**
Specialist ($\ell \geq 1$)	0.270	0.012
	(0.260)	(0.266)
	[$p = 0.30$]	[$p = 0.97$]
Implied premium e^{β}	1.31×	1.01×
Controls (age, age^2, female, firm size)	Yes	Yes
Firm fixed effects	No	Yes
N	60	60
R^2	0.285	0.625

Standard errors in parentheses. The pooled specialist premium is positive but imprecise between firms and essentially zero within firms; it is smaller than the raw head-to-non-specialist ratio $w_1/w_0 = 1.90$ (Table 6) because it averages all specialist ranks ($\ell \geq 1$) against non-specialists rather than contrasting the head partition alone, and because firm fixed effects absorb most cross-firm wage variation in a sample of 18 firms. Results are unchanged when the single data-error wage is dropped.

D.2 Wages by Partition Breadth: Per-Firm Detail

Single-partition workers report all activities within one specialist partition; multi-partition workers span two or more. Across the seven firms with both types, the mean multi/single wage ratio is 1.29 (median 1.26).

Table D.3: Within-Firm Wage Ratio: Multi- vs. Single-Partition Workers (threshold method)

Firm	N_{single}	N_{multi}	**Mean single**	**Mean multi**	**Ratio**
201017925451	3	1	1467	2600	1.77
201715456771	1	4	5500	9125	1.66
203297254311	1	1	9999	4000	0.40
204526147671	9	1	5538	1600	0.29
205530148411	2	3	2600	3267	1.26
206071670961	3	2	6000	6750	1.12
206095632661	1	2	1300	3300	2.54
Summary ($N_{\text{firms}} = 7$)			Mean = 1.29		Median = 1.26

Wages in soles/month. Partitions from the threshold procedure of Section 3.3.1. Restricted to firms with both single- and multi-partition workers.

E Model: Theory and Proofs

E.1 Aggregate Output: Hopenhayn Derivation

We derive the aggregate output decomposition in equation (16). Given the partition structure $\{\mathcal{A}_\ell\}$ and the optimal time profile $\{x_i^*\}_\ell$, let $\bar{T}_\ell$ be as defined in Section 2 and let $\Lambda \equiv \sum_\ell n_\ell \bar{T}_\ell$ be the firm's total effective labor. *Step 1: Firm's profit.* Taking the partition and $\bar{T}_\ell$ as given,

$$\pi_z \;=\; p\, z^{1-\nu}\, \Lambda^\nu \;-\; \omega\, \Lambda.$$

Step 2: FOC for Λ. Let $\Lambda^*(z) \equiv \sum_\ell n_\ell^*(z)\bar{T}_\ell$ denote total effective labour at a firm of type z. The first-order condition yields effective labor that scales linearly in z:

$$p\,\nu\, z^{1-\nu}\, \Lambda^{\nu-1} \;=\; \omega \qquad \Longrightarrow \qquad \Lambda^*(z) \;=\; \left(\frac{p\,\nu}{\omega}\right)^{1/(1-\nu)} z.$$

Step 3: Firm output. Substituting $\Lambda^*(z)$ into the firm's production function,

$$y^*(z) \;=\; p\left(\frac{p\,\nu}{\omega}\right)^{\nu/(1-\nu)} z \qquad \Longrightarrow \qquad Y \;=\; p\left(\frac{p\,\nu}{\omega}\right)^{\nu/(1-\nu)} \int_{\bar{z}} z\, dG(z).$$

Step 4: Workers per firm. Using $\Lambda = N \cdot \sum_\ell \frac{n_\ell}{N}\bar{T}_\ell$ with $\frac{n_\ell}{N}$ fixed within a partition band,

$$N^*(z) \;=\; \frac{\Lambda^*(z)}{\displaystyle\sum_\ell \frac{n_\ell}{N}\bar{T}_\ell}.$$

Step 5: Labor market clearing. Since $N^*(z) = \Lambda^*(z)/S(z)$ where $S(z) \equiv \sum_\ell (n_\ell^*/N^*)\bar{T}_\ell$ is the firm-level employment-weighted average of efficiency units, the clearing condition $\int_{\bar{z}} N^*(z)\,dG(z) = \bar{N}$ gives

$$\Big(\frac{p\,\nu}{\omega}\Big)^{1/(1-\nu)} \int_{\bar{z}} \frac{z}{S(z;\omega)}\,dG(z) \;=\; \bar{N},$$

which is the fixed-point equation for ω stated in (15). Define the economy-wide employment-weighted average $\bar{S} \equiv \bar{N}^{-1}\int_{\bar{z}} N^*(z)\,S(z)\,dG(z)$; since $N^*(z)\,S(z) = \Lambda^*(z) = (p\nu/\omega)^{1/(1-\nu)}z$, we have

$$\bar{S} \;=\; \frac{1}{\bar{N}}\Big(\frac{p\,\nu}{\omega}\Big)^{1/(1-\nu)} \int_{\bar{z}} z\,dG(z).$$

Substituting into Step 3 gives the aggregate output formula in (16):

$$Y \;=\; \bar{N}^{\,\nu} \cdot \Big(\int_{\bar{z}} z\,dG(z)\Big)^{1-\nu} \cdot \bar{S}^{\,\nu}. \tag{27}$$

The organizational term $\bar{S}$ is the economy-wide aggregate; the firm-level $S(z)$ is non-smooth in z (jumping at each staircase threshold), which is why $\bar{S}$ responds to changes in the partitioning costs I and c_{prop}.

E.2 Proofs

Proof of Proposition 3. Part (a). Equilibrium profit with ℓ active specialist partitions satisfies $\pi_\ell^*(z) = C_\ell(\omega)\cdot z$, where $C_\ell(\omega) > 0$ is independent of z (since $\Lambda^*(z) = (p\nu/\omega)^{1/(1-\nu)}z$ is linear in z). Since $\phi_\ell > 1$, specialist partitions strictly raise output per worker, so $C_{\ell+1}(\omega) > C_\ell(\omega)$. The gain from adding one partition, $[C_{\ell+1}(\omega) - C_\ell(\omega)]\cdot z$, is strictly increasing in z while the fixed cost I is flat; a unique finite threshold $\bar{z}_\ell = I/[C_{\ell+1}(\omega) - C_\ell(\omega)]$ exists at which the gain equals I. Firms with $z < \bar{z}_\ell$ optimally select $L^* = \ell$ and firms with $z > \bar{z}_\ell$ optimally select $L^* = \ell+1$. Since $\tilde{g}$ is strictly decreasing in i, each successive specialist partition added at a higher partition count covers a less productive range of activities (higher activity indices). Formally, $C_{\ell+1}(\omega) - C_\ell(\omega)$ equals the profit increment from assigning the marginal activity range to a specialist partition rather than partition 0; since this marginal range lies at higher i (less productive), the increment is strictly smaller for each additional partition: $C_{\ell+2} - C_{\ell+1} < C_{\ell+1} - C_\ell$. This induces the ordering $\bar{z}_1 < \bar{z}_2 < \cdots$ *Part (b).* Define $H(a^{\max}; z) \equiv (n_L^*(z)\phi_L - n_0^*(z))\,\tilde{g}(x^*_{a^{\max}}, a^{\max})\,f(a^{\max}) - c_{\text{prop}}$. Since worker shares within a regime are fixed (from the staffing FOC), both $n_L^*(z)$ and $n_0^*(z)$ are proportional to $N^*(z) \propto z$, so $\partial H/\partial z > 0$. Let $G(a^{\max}) \equiv \tilde{g}(x^*_{a^{\max}}, a^{\max})\,f(a^{\max})$. By the envelope theorem applied to the worker's problem: $\mathrm{d}[c \cdot x^*_{a^{\max}} - \tilde{g}(x^*_{a^{\max}}, a^{\max})]/\mathrm{d}a^{\max}$ has sign determined by $c \cdot x^*_{a^{\max}} - \tilde{g}(x^*_{a^{\max}}, a^{\max})$. Since $\tilde{g}_{xx} < 0$ and $\tilde{g}(0, i) \geq 0$, the concavity inequality $\tilde{g}(x, a^{\max}) > \tilde{g}_x(x, a^{\max}) \cdot x = c \cdot x^*_{a^{\max}}$ holds for all $x^*_{a^{\max}} > 0$, which implies $\partial G/\partial a^{\max} < 0$: increasing $a^{\max}$ reduces the marginal gain from the boundary activity. Therefore $\partial H/\partial a^{\max} < 0$, and by the implicit function theorem $\mathrm{d}a^{\max *}/\mathrm{d}z = -(\partial H/\partial z)/(\partial H/\partial a^{\max}) > 0$. □

Proof of Proposition 4. Part (a). The efficiency units produced by a worker in partition ℓ are $\bar{T}_\ell = \phi_\ell \int_{\mathcal{A}_\ell} \tilde{g}(x_i^*, i)\,f(i)\,di$. By Proposition 1, x_i^* depends only on $\mathcal{A}_\ell$ and $\tilde{g}$, not on z or ω. Hence for any fixed partition structure $\{\mathcal{A}_\ell, \phi_\ell\}$, the term $\bar{T}_\ell$ is constant across all z and ω. Since $w_\ell = \omega \bar{T}_\ell$, the ratio $w_\ell/w_0 = \bar{T}_\ell/\bar{T}_0$ is independent of z and of ω. (Within a partition regime, $a^{\max *}(z)$ varies continuously with z by Proposition 3(b), so the partition structure itself changes; part (a) applies to a cross-firm comparison among firms with identical partition boundaries.) *Part (b).* By Proposition 3(b), $a^{\max *}$ rises strictly at threshold $\bar{z}_L$: the lower bound of the partition-0 domain shifts from $a_-^{\max}$ to $a_+^{\max} > a_-^{\max}$. The integral defining $\bar{T}_0 = \int_{a^{\max}}^1 \tilde{g}(x_i^*, i)\,f(i)\,di$ loses the interval $[a_-^{\max}, a_+^{\max}]$, over which the integrand is strictly positive (since $\tilde{g} > 0$). Moreover, since $\tilde{g}(x, i)$ is strictly decreasing in i (the activity-ordering assumption in Section 2), the integrand on $[a_-^{\max}, a_+^{\max}]$ is strictly larger than on $[a_+^{\max}, 1]$: the removed activities are strictly more productive than those remaining. Therefore $\bar{T}_0$ falls strictly at every regime threshold. □

Proof of Proposition 2. Part (a): $\partial a^{\max *}/\partial\omega < 0$. Define

$$H(a^{\max}, \omega) \;\equiv\; (n_L^*\phi_L - n_0^*)\,\tilde{g}(x^*_{a^{\max}}, a^{\max})\,f(a^{\max}) - c_{\text{prop}}.$$

The equilibrium outer boundary satisfies $H(a^{\max *}, \omega) = 0$. By the implicit function theorem,

$$\frac{\partial a^{\max *}}{\partial \omega} \;=\; -\frac{\partial H/\partial\omega}{\partial H/\partial a^{\max}}.$$

Sign of $\partial H/\partial a^{\max}$. This is established in the proof of Proposition 3(b): since $\tilde{g}$ is strictly decreasing in i and $\tilde{g}_{xi} < 0$, the marginal gain from extending $a^{\max}$ is strictly diminishing, so $\partial H/\partial a^{\max} < 0$. *Sign of $\partial H/\partial\omega$.* From the firm's staffing FOC (8), optimal total efficiency units satisfy $\Lambda^*(z; \omega) = (p\nu z^{1-\nu}/\omega)^{1/(1-\nu)}$, which is strictly decreasing in ω. Within a given partition regime the staffing shares $s_\ell \equiv n_\ell^*/N^*$ are fixed (the interior boundary conditions (9) pin down the ratios $n_\ell^*/n_{\ell+1}^*$ independently of ω), so $n_\ell^* = s_\ell \Lambda^*/\bar{T}_\ell$ and each n_ℓ^* is proportional to $\Lambda^*(\omega)$. Therefore

$$n_L^*\phi_L - n_0^* \;=\; (s_L\phi_L/\bar{T}_L - s_0/\bar{T}_0)\,\Lambda^*(\omega) \;\equiv\; \kappa\,\Lambda^*(\omega),$$

where $\kappa > 0$ is required for the firm to find it profitable to maintain any specialist partition (otherwise $a^{\max *} = 0$). Since $\tilde{g}(x^*_{a^{\max}}, a^{\max})\, f(a^{\max}) > 0$ and $\partial\Lambda^*/\partial\omega < 0$, we have

$$\frac{\partial H}{\partial \omega} \;=\; \kappa\,\frac{\partial \Lambda^*}{\partial \omega}\,\tilde{g}(x^*_{a^{\max}}, a^{\max})\, f(a^{\max}) \;<\; 0,$$

where we have used the fact that $x^*_{a^{\max}}$ is invariant to ω conditional on $a^{\max}$ (Proposition 1), so $\tilde{g}$ does not move with ω holding $a^{\max}$ fixed. Combining signs:

$$\frac{\partial a^{\max *}}{\partial \omega} \;=\; -\frac{(-)}{(-)} \;<\; 0.$$

The intuition is direct: a higher wage reduces the firm's optimal labor force Λ^* proportionally, shrinking the absolute gain from assigning the boundary activity to the specialist partition. Since the marginal variable cost c_{prop} is fixed, the firm finds it optimal to contract the specialist domain. *Part (b): Effect on x_i^* in partition* 0. Partition 0 covers $[a^{\max}, 1]$. A fall in $a^{\max *}$ (from part (a)) widens this partition: the new domain $[a^{\max *}_{new}, 1] \supset [a^{\max *}_{old}, 1]$ includes additional activities in $[a^{\max *}_{new}, a^{\max *}_{old}]$. The ratio $c_0^* \equiv \lambda_0/(\omega\phi_0)$ is pinned by the time constraint

$$\int_{a^{\max *}}^{1} \tilde{g}_x^{-1}(c_0^*;\, i)\, f(i)\, di \;=\; 1.$$

Adding activities raises the left-hand side for any fixed c_0^* (since $\tilde{g}_x^{-1} > 0$). For the constraint to continue to hold, c_0^* must fall. Since $x_i^* = \tilde{g}_x^{-1}(c_0^*;\, i)$ is strictly increasing in c_0^* (because $\tilde{g}_{xx} < 0$), x_i^* falls for every $i \in [a^{\max *}_{old}, 1]$: workers spread their time budget more thinly across the now-wider partition. Formally, $\partial x_i^*/\partial\omega > 0$ for $i \in \mathcal{A}_0$: as the wage rises, partition 0 widens and focus within it falls. *Part (c): ω-invariance of specialist partitions within a regime.* Interior boundaries are pinned by (9): $n_\ell^*\phi_\ell = n_{\ell+1}^*\phi_{\ell+1}$. The variable cost $c_{\text{prop}} \cdot a^{\max}$ does not appear in this condition because moving an interior boundary a_ℓ reassigns an activity between two specialist partitions, leaving $a^{\max}$ and thus the total variable cost unchanged. The ω-dependence cancels because each n_ℓ^* scales proportionally with $\Lambda^*(\omega)$: the ratios $n_\ell^*/n_{\ell+1}^*$ are ω-invariant. Hence the interior boundaries $\{a_\ell^*\}$ are fixed within a partition regime, the widths $|\mathcal{A}_\ell|$ are unchanged, and by Proposition 1 so are $\{x_i^*\}_{i\in\mathcal{A}_\ell}$ for all $\ell \geq 1$. *Part (d): Discrete reorganization at staircase thresholds.* From the proof of Proposition 3(a), the threshold at which the ℓ-th specialist partition becomes unprofitable satisfies $\bar{z}_\ell(\omega) = I/[C_{\ell+1}(\omega) - C_\ell(\omega)]$, where $C_{\ell+1}(\omega) - C_\ell(\omega)$ is strictly decreasing in ω (a higher wage compresses per-unit profit increments). Hence $\bar{z}_\ell(\omega)$ is strictly increasing in ω:

as wages rise, the minimum productivity needed to sustain each partition rises. For a firm of fixed productivity z, a sufficient increase in ω eventually violates $z > \bar{z}_\ell(\omega)$, causing the firm to drop partition ℓ. The activities previously in that specialist partition collapse into partition 0, discretely widening $\mathcal{A}_0$, discretely reducing c_0^*, and discretely reducing x_i^* for all remaining partition-0 workers. $\square$

Proof of Proposition 5. Let $\Phi(\omega) \equiv p\nu\big[\bar{N}^{-1}\int_{\bar{z}} z/S(z;\omega)\,dG(z)\big]^{1-\nu}$. Since $S(z;\omega)$ is bounded above by $S_{\max} < \infty$ (finitely many partitions are ever profitable), $\Phi(\omega) \geq p\nu\big[\bar{N}^{-1}(\int_{\bar{z}} z\,dG)/S_{\max}\big]^{1-\nu} > 0$ for all ω. Hence $\Phi(\omega) > \omega$ for all ω sufficiently close to zero. Since $S(z;\omega)$ is bounded below by $S_{\min} > 0$ (partition 0 always exists with $\phi_0 = 1 > 0$), we have $\Phi(\omega) \leq p\nu\big[\bar{N}^{-1}(\int_{\bar{z}} z\,dG)/S_{\min}\big]^{1-\nu} < \infty$; thus $\Phi(\omega) < \omega$ for all ω sufficiently large. The function Φ is continuous in ω on each open interval between consecutive staircase threshold values of ω (since $S(z;\omega)$ is constant on such intervals). As ω crosses a staircase threshold, $S(z;\omega)$ falls for the affected firms, causing Φ to jump upward; between thresholds, Φ is continuous. Since $\Phi(\omega) - \omega$ starts positive (for small ω) and ends negative (for large ω), and is continuous on each inter-threshold interval, the intermediate value theorem applied to the last interval on which $\Phi - \omega$ transitions from positive to negative guarantees at least one zero. If Φ is strictly decreasing in ω, then $\Phi(\omega) - \omega$ is strictly decreasing and crosses zero at most once, so the fixed point is unique. $\square$

F Quantitative Appendix

F.1 Calibration Robustness: Partition Rule and Threshold

Table F.1 reports the partition moments under the three grouping rules at the headline threshold $\tau = 0.5$. The wage step w_2/w_1 is stable around 0.5–0.6 (close to $\xi = 0.566$) and the head premium w_1/w_0 around 1.6–1.9 across rules. Across $\tau \in \{0.33, 0.50, 0.67\}$ the structure moves monotonically: at $\tau = 0.33$ only $\sim$11% of workers are non-specialists, while at $\tau = 0.67$ the threshold is so strict that 50–60% collapse into the residual and partitions nearly vanish; $\tau = 0.5$ (the median positive co-occurrence) is the headline.

Table F.1: Calibration robustness across grouping rules ($\tau = 0.5$)

Grouping rule	**Mean # spec. part.**	ℓ=0 **share**	w_1/w_0	w_2/w_1
Clique (headline)	1.39	18.5%	1.90	0.58
Average-cut	1.33	16.9%	1.88	0.56
Components	1.17	15.4%	1.64	0.51

Headline: clique, $\tau = 0.5$. Full $\tau\times$rule grid in the replication output `rank_moments_threshold.csv`.