Working paper, August 2026. Comments welcome. The online appendix is available from the corresponding author on request.
Cite as: Aldunate, F., Enrione, A., Girardi, J., Opazo, C., Torres-Sahli, M., & Truffa, S. (2026). Hierarchies: New facts from administrative HR data. Working paper, August 2026. manueltorressahli.com/papers/hierarchies

Hierarchies: New Facts from Administrative HR Data

Felipe Aldunate*[a], Alfredo Enrione[a], Jacinta Girardi[b], Constanza Opazo[b], Manuel Torres-Sahli†[c], and Santiago Truffa[d]

[a]ESE Business School, Universidad de los Andes, Chile
[b]Research Department, Buk SpA, Chile
[c]Faculty of Social Sciences, Universidad de los Andes, Chile
[d]Business School, Universidad Adolfo Ibáñez, Chile

August 2026

Abstract

We map the internal hierarchy of 3,059 firms from administrative HR records in which at least 95% of employees resolve to a position in the reporting tree, reading the rungs of the job ladder off the reporting tree rather than inferring them from occupation codes, titles or job transitions. We document three regularities. Firms grow wider rather than deeper: doubling employment adds 0.51 layers. Pay tracks hierarchical position, and 8.3 of the 30.6 percentage points hierarchy adds to within-firm wage variance come from span and reach, which have no occupational analogue. Deeper firms do not replicate shallow ones: steps between neighbouring rungs compress as height rises even though the top-to-bottom gap widens, and at the apex the premium of tall firms turns out to be a premium for reach, not for rank. Observing the ladder also shows that about four fifths of the gender wage gap (79%) is pay within the same rung of the same firm, not sorting across rungs or firms.

Keywords: Organizational hierarchy, internal labor markets, job ladder, promotion, wage structure, span of control, administrative data

JEL Codes: D23, J24, J31, J33, J62, L22, M52

*Aldunate acknowledges funding from ANID FONDECYT Regular 1262214.
†Corresponding author: mtorressahli@gmail.com

1 Introduction

Hierarchies are one of the most fundamental yet least directly observable institutions for organizing production. Unlike markets, which allocate resources through prices, firms rely on internal structures to coordinate effort, allocate authority, and process information. Hierarchies are the backbone of this coordination: they embed who decides what, who supervises whom, and how knowledge and accountability flow. These internal architectures affect not only how firms function, but also how wages, promotions, and opportunities are distributed.

Empirical work on hierarchies has long been constrained by measurement. Because firms rarely disclose internal reporting structures in a systematic way, researchers have had to rely on partial or indirect proxies, including occupational categories in administrative data, executive-team structures in public disclosures, and survey-based organizational measures (e.g., Belloc et al., 2023; Caliendo et al., 2015; Guadalupe et al., 2013; Rajan & Wulf, 2006; Russo & van Houten, 2021; Spanos, 2022). A major recent advance is Ewens and Giroud (2025), who reconstruct hierarchical layers for roughly 3,100 U.S. publicly traded firms from the online résumés of about seven million employees, finding a mean firm height of 9.5 layers and pyramidal organizational structures. Our paper is complementary, but it draws on a different source of information: administrative records from a human resources software platform. Because the platform is used to manage payroll, reporting lines, and core HR workflows, it records direct supervisor–subordinate links together with compensation and worker outcomes across the workforce of participating firms, from CEOs to frontline workers; the analysis sample keeps the firm-months in which at least 95% of employees resolve to a position in the reporting tree. This allows us to observe the hierarchy itself rather than infer it from titles, functions, or occupations, to measure spans of control directly, and to track changes in the reporting tree over time. Our data cover 3,059 Chilean firms and approximately 10.1 million employee-months (70,753 firm-months) from 2021 to 2025, spanning a broad range of firm sizes and industries. We use these data to document descriptive facts on hierarchical structure, compensation, and the compression of both as firms deepen, together with evidence on industry heterogeneity and the gender wage gap. The records are administrative but the sample is not a census: the platform's clients are self-selected, and our estimation sample covers 3.3% of the 73,340 Chilean firms with at least twenty dependent workers, with coverage rising in firm size up to the 250-999 bracket and thinning again above it. A further hierarchy-quality filter, which requires a firm-month's reporting graph to resolve cleanly to a single apex, leaves, together with the data-quality screen, 28.2% of the employee-months in the window and binds hardest on the largest firms. We document both margins in Appendix A.4 and return to what they imply for external validity in the conclusion.

Hierarchical structure. Firms exhibit substantial variation in organizational depth, even among those of similar size. Firm height ranges from 2 to 9 layers, with height-4 being the most common. Each doubling of firm size is associated with only 0.51 additional layers on average – roughly one new layer for every quadrupling of employment. Spans of control are highly skewed: 21.4% of managers supervise exactly one direct report and 63% supervise four or fewer, yet a long right tail extends to a maximum of 1,779 direct reports. As hierarchies deepen, firms shift supervisory responsibility downward, producing a diamond shape, in which the CEO level accounts for less than 1% of headcount in tall firms while the bottom two layers contain the majority of employees.

Compensation. Pay is sharply convex in hierarchical position.[1] CEO-to-median pay ratios rise systematically with hierarchy height – from approximately 3× in height-3 firms to 6× in height-6 – and the overall CEO-to-median ratio is 4.2×, mirroring patterns documented for U.S. public firms. Most wage variation (68.6%) occurs *within* rather than between firms, and hierarchy variables substantially increase explained within-firm wage variance (Section 4.2). Salary volatility is lowest at the top and rises toward the bottom of the ladder: month-to-month pay risk falls disproportionately on workers at the bottom. Because the series is monthly payroll, it excludes the annual and deferred components that carry executive incentive pay, so it measures take-home risk rather than total compensation risk.

Compression. Deep firms are not shallow firms with more layers attached: as height rises, adjacent rungs become more alike at both ends of the ladder even as the distance from top to bottom widens. The wage step between neighbouring levels falls from -0.223 to -0.181 log points as firm height rises from three layers to six, an attenuation of about 19%. Among workers who supervise anyone at all, the level gradient in span falls from -0.518 to 0.055 (SE 0.012), crossing zero: in the tallest firms a supervisor one rung deeper commands a *larger* team, because spans widen into the middle of the hierarchy and ease only below it. Rank does not stop governing team size; the direction reverses. The pattern survives collapsing every chain of single-report managers, so it is not an artefact of titling, and documenting it requires observing depth and span for the same firm at the same time. Height also interacts with industry: the height–size semi-elasticity ranges from 0.90 in agriculture to 0.65 in administrative services across the 11 sectors estimated with precision, though the ordering is imprecise: 15 of the 136 pairs across all 17 sectors separate at the 5% level, and 11 of the

[1] The convex-pay pattern is consistent with at least four mechanisms: assignment-equilibrium with multiplicative spillovers (Rosen, 1982); tournament prize structures (Lazear & Rosen, 1981); supervision-loss vertical attenuation (Calvo & Wellisz, 1978); and team-production metering by a centralized residual-claimant monitor (Alchian & Demsetz, 1972). Prendergast (1999) surveys the principal-agent literature and reaches the same conclusion: cross-sectional pay-by-level patterns alone cannot adjudicate among them.

55 pairs among the 11 precisely estimated ones, so we read that variation as heterogeneity rather than as a ranking.

Observing the ladder also lets us split the gender wage gap along it. Of a raw gap of -0.188 log points, 12% is accounted for by the rung women occupy and 9% by the firm they work in; the remaining 79% is pay differences within the same rung of the same firm.[2] The glass ceiling is real — women are 39.6% of the workforce and 16.8% of CEO-level positions — but it is numerically small, because the levels where under-representation is most extreme hold few workers. Whether a gap of this kind reflects sorting or pay is a distinction that requires knowing the rung, and the reporting tree supplies it.

Table 1 summarises how six measurement approaches across five countries compare in terms of data, scope, and key estimates.

[2]The split depends on the weighting convention, and both readings tell the same story. Under the firm-month-equal weights used in every regression in this paper the raw gap is -0.188 and the shares are 12/9/79. Weighting each firm equally over the panel instead—the convention used for the descriptive exhibits—gives a raw gap of -0.182 and shares of 9/11/80. What is invariant is that roughly four fifths of the gap is pay within rung and firm.

Table 1: **Cross-country comparison of hierarchy measurement approaches.**

	Caliendo et al. (2015)	Tåg (2013)	Belloc et al. (2023)	Russo & van Houten (2021)	Ewens & Giroud (2025)	This paper
Country	France	Sweden	Europe (32)	EU (27/28)	United States	Chile
Measurement	Occupation-based (DADS categories)	Occupation-based (Swedish admin.)	Manager survey (ECS 2013)	Manager survey (ECS 2019)	Title/resume-based (web-scraped)	Reporting links (HR platform)
Sample	Manuf. census	Manuf. census	Establishments $\geq$10[e]	Establishments $\geq$10[e]	Public firms	Platform clients
N firms[a]	≈76,000	≈5,600	≈19,000	18,287	≈3,100	3,059
Mean firm size[b]	≈43	60.3	n/c	n/c	7,700	140
Mean height (total layers)[c]	≈2.6	2.4	≈3	3.06	9.5	4.1
Height–size elasticity[d]	—	—	—	—	≈ 2.2 layers/doubling	0.51 layers/doubling
Span distribution	Not reported	Not reported	Not reported	Not reported	Not reported	Directly observed

Notes: The six studies measure conceptually different objects and are not directly comparable. Differences in mean height likely reflect a combination of measurement definition, sample composition, and genuine cross-country differences in organizational depth, though the relative contribution of each cannot be determined from these comparisons. "n/c" denotes a quantity that is not directly comparable because of stratified sampling.

[a] Number of unique firms in the cross-section (averaged across panel years for Caliendo et al. and Tåg). For Belloc et al., the full-controls regression sample.

[b] Mean number of employees per firm. For Caliendo et al., we estimate this by dividing the average annual hours per firm (≈78,000) by a standard 1,800-hour work year. Belloc et al. and Russo & van Houten use employment-stratified establishment sampling that does not yield a single comparable mean.

[c] Total hierarchical layers, including the bottom layer of production workers. Caliendo et al. and Tåg report the number of *management* layers (excluding the worker layer), so we add 1 to their reported means of ≈1.6 and 1.4, respectively, for comparability.

[d] To compare height–size scaling on a common metric, we report *layers per doubling of firm size*. Ewens and Giroud (2025) estimate a log-log elasticity of 0.33 in their setting; at their mean depth of 9.5 layers, this translates to roughly 2.2 additional layers per doubling of employment. Our 0.51 is reported directly as a semi-elasticity (additional layers per doubling). Caliendo et al. and Tåg do not report a regression of layers on employment; their reported quantities are cross-sectional means by layer count. Belloc et al. and Russo & van Houten do not report height–size elasticities; their respective contributions concern the employee-representation channel and the incentive-substitution channel rather than the size–depth scaling per se. The remaining gap between ≈ 2.2 and 0.51 reflects, in part, the substantially deeper hierarchies of large U.S. public firms relative to our Chilean sample.

[e] Belloc et al. (2023) and Russo & van Houten (2021) both use establishment-level manager self-reports of hierarchical layer counts – the closest non-reporting-link approach to direct establishment-level depth measurement. Their independent ≈ 3-layer European means converge across two ECS waves conducted at different times by different consortia, providing cross-validation of the European depth benchmark against which our Chilean sample mean of 4.1 sits roughly one layer higher.

It is worth separating what is new from what is newly measured. The height–size scaling and the convexity of pay largely confirm patterns documented in other countries with other methods; what we add there is a cross-country benchmark on a measure that does not depend on occupation coding. Three further regularities — the span distribution, in which 21.4% of managers supervise exactly one person; the month-to-month pay-risk gradient, lowest at the top and peaking below the middle; and the asymmetry of height changes, where increases involve reporting-chain restructuring and decreases are boundary churn — have been conjectured or inferred before but not observed directly, because inferring a span from a title is not the same as reading it off a reporting link. Two results we have not found elsewhere at all, and they are the ones the next paragraph takes up.

Two of the facts above are, as far as we can establish, new. The first is compression: as firms add layers, the rung-to-rung differences that define a hierarchy get smaller. The wage step falls by 19% between height-3 and height-6 firms, the per-level gradient in the probability of supervising anyone runs from -0.476 to -0.184 — a margin whose movement is close to what the range restriction alone requires, so we report it for completeness rather than as economic content — and the level gradient in team size goes from -0.518 to 0.055 — it crosses zero, so among supervisors in the tallest firms a manager one rung deeper runs a *larger* team. Deep firms are not shallow firms with more of the same on top; the ladder they build has finer treads, and at the top it stops running in the same direction. The second is that *reach* — the size of the subtree below a manager — explains wages beyond the number of direct reports, contributing 8.3 of the 30.6 percentage points — a sequential attribution, and because span and reach enter last it is the conservative one — that hierarchy adds to within-firm wage variance. Position in the chain of authority is priced by how much of the organisation sits underneath it, not by how many people report to it directly.

Neither fact is recoverable from the data economists usually have. Occupation codes and job titles order workers vertically but do not say who reports to whom, so span and reach have no analogue in them, and a rung in one firm cannot be matched to a rung in another. What makes these facts observable is that the reporting link itself is recorded — not inferred from titles or surveys, but written down in the systems firms use to run payroll. That is the warrant for the exercise rather than its result, and we treat it that way: the remaining regularities we document — the span distribution, the pay gradient, the volatility profile, the mechanics of height change — confirm and sharpen patterns others have found by other means, and we report them as such.

We also link hierarchical position to outcomes measured on the same records: compensation, turnover, absence, leave, and internal mobility. Because the hierarchy and the outcome come from one administrative source, the link does not depend on matching across datasets, and it covers a cross-section of firms varying in size, industry, and organisational form.

Section 3 describes the data source, institutional context, measurement approach, and height dynamics. Section 4 presents the main empirical findings, organized around three core fact groups — hierarchical structure, compensation, and compression — followed by industry heterogeneity and the gender gap. Section 5 discusses implications and concludes.

2 Related Literature

This paper sits between two literatures that rarely meet: the theory of knowledge hierarchies, which predicts how depth, span and pay should co-move as firms grow, and the empirical study of internal wage structure, which documents how careers unfold inside firms. What keeps them apart is largely a matter of what the data record. Tests of the theory have mostly been run on layer counts built from occupational categories or establishment surveys; studies of internal wage structure locate a worker on a ladder defined by the content of the job — its complexity, autonomy and responsibility — rather than by whom the worker reports to. Reporting links are seldom in the file.

The theoretical importance of hierarchy is well established along several distinct lines. One tradition derives hierarchical form from the architecture of decentralized information processing under capacity constraints (Bolton & Dewatripont, 1994; Radner, 1993); a related but distinct line treats hierarchy as the efficient way to organize knowledge acquisition and the matching of problems to those who can solve them (Garicano, 2000; Garicano & Rossi-Hansberg, 2015). A second family rationalizes hierarchical structure through incentive alignment and the limits of supervision – through team-production monitoring (Alchian & Demsetz, 1972; Calvo & Wellisz, 1978), rank-order tournament prize structures (Lazear & Rosen, 1981), or assignment of heterogeneous talent to ranks with multiplicative-productivity spillovers (Rosen, 1982); the broader principal-agent / incentive-provision literature is surveyed in Prendergast (1999), with bounded rationality and cumulative loss of control imposing a quasi-static limit on firm size (Williamson, 1967). A third family treats the firm as a dedicated hierarchy organized around the management of access to non-contractible critical resources (Rajan & Zingales, 2001).

Empirically, pioneering insider-data studies of internal labor markets (Baker et al., 1994) demonstrated the value of personnel records for understanding wage dynamics within firms. Caliendo et al. (2015) relate inferred layers to productivity and wages in French firms, documenting that adding a management layer is associated with falling average wages at all pre-existing layers; Caliendo et al. (2018) extend this framework to show how firms reorganize their management hierarchies as they enter export markets. Rajan and Wulf (2006) document a thirteen-year flattening of senior U.S. corporate hierarchies (CEO span rose from 4.46 to 6.79; depth fell by more than 25%; intermediate positions, especially the COO, were eliminated) and argue that this combination of wider direct CEO contact with delegated authority and incentive-based pay cannot be reduced to either centralization or decentralization; Guadalupe et al. (2013) extend this evidence to show that approximately three-quarters of CEO span growth is attributable to functional rather than general managers, characterizing the emerging hybrid as "functional centralization." Tåg (2013) replicates these patterns for Swedish manufacturing, confirming that firms with more layers are larger and pay higher wages; Spanos (2022) extends the Caliendo et al. framework from productivity to wage inequality, documenting that firm organization mediates the density-driven component of within-firm wage disparities in French employment areas; Belloc et al. (2023) document, across 32 European countries, that establishments with institutionalized employee representation have deeper hierarchies, attributing this to skip-level reporting mechanisms that reduce communication costs; and Russo and van Houten (2021) document, across 18,287 EU establishments, that broader adoption of complex job design and autonomous teams is associated with fewer hierarchical layers, attributing this to incentive substitution: when jobs themselves provide motivation, hierarchical mobility is less needed as an incentive mechanism. Cross-country evidence further documents large variation in delegation of decision rights (Bloom et al., 2012) and in management practices (Bloom & Van Reenen, 2007). In a complementary stream, Card et al. (2013) apply the AKM framework of Abowd et al. (1999) to West German administrative data, documenting that between-firm wage differentials account for a large share of the *rise* in overall wage inequality; Song et al. (2019) extend the framework to U.S. SSA full-population data and show that between-firm sorting and segregation explain essentially the entire rise in between-firm inequality over 1978–2013. AKM-tradition decompositions thus characterize *between-firm* wage differentials as a major component of inequality and its rise; the present paper sits at the intersection of that tradition with the within-firm-hierarchy literature, using direct reporting links to measure how within-firm hierarchical position structures pay in a setting where the between-firm AKM decomposition is unavailable but the within-firm hierarchy structure is directly observed. These studies establish important regularities, but they often rely on indirect measures, focus on upper layers, or lack linked worker-level outcomes. The three closest comparison studies – Caliendo et al. (2015), Tåg (2013), and Ewens and Giroud (2025) – measure conceptually different objects. Caliendo et al. (2015) define layers from occupational categories in French administrative data (cadres, techniciens, employés, ouvriers) – a coarser measure that conflates organizational hierarchy with job classification conventions. Ewens and Giroud (2025) infer hierarchy from observed promotion flows over web-scraped resume data, which captures the implied transition graph but not actual supervisory relationships. Our reporting-link measure directly records who reports to whom, avoiding both the occupational conflation of the French approach and the inference required of the promotion-flow approach.

Closest to our exercise is the work of Bayer and Kuhn (2023) and Bayer and Kuhn (2019), who show that *job levels* — a graded measure of the complexity, autonomy and responsibility of a job, distinct from occupation — account for most observed wage differences and for a large share of life-cycle wage growth, and who read the gender wage gap as a gap in career progression across those levels. Their ladder and ours are not the same object. A job level describes the content of a role and is comparable across employers; our level is depth in an observed chain of authority within one firm. The two are complementary, and the distinction matters for the gender results in Section 4.5: when the rung is defined by the reporting tree rather than by job content, sorting across rungs accounts for a small share of the raw gap. Progression in the sense of Bayer and Kuhn (2019) therefore takes place largely *within* hierarchical layers, on a margin that the reporting tree does not resolve and job levels do. Jung and Kuhn (2019) similarly link mobility to life-cycle earnings using worker histories rather than firm structure.

Compensation patterns also speak to the broader literature on wage inequality. While that literature has emphasized differences across firms, our results show that a substantial share of pay dispersion operates within firms and is structured by hierarchical position. Hierarchy thus mediates the firm-level patterns documented in AKM-style decompositions (Card et al., 2013), pointing to organizational structure as a channel through which between- and within-firm inequality interact.

A complementary literature links organizational structure to firm performance. Bloom and Van Reenen (2007) develop management-practice measures (operations, monitoring, targets, incentives) showing that better-managed firms are more productive, profitable, and more likely to survive — though hierarchical structure itself is not among the eighteen practices their instrument scores. Bloom et al. (2012) extend the empirical program to organizational structure proper, showing that decentralization of decision rights varies systematically with institutional trust and complements information-technology investment to raise within-firm productivity (their analysis finds no direct productivity effect of decentralization absent IT). On the inverse causal direction, Guadalupe and Wulf (2010) use the 1989 Canada–U.S. Free Trade Agreement as a quasi-natural experiment to show that increased product market competition causally flattens hierarchies, with simultaneous changes in firm scope, division-manager pay, and incentive intensity — illustrating that hierarchies adjust as a complementarities-bundled system to environmental shocks. These studies establish that hierarchy is not merely an internal labor market feature but a strategic variable with market consequences.

A separate but complementary literature on team-based incentives and peer effects examines within-firm productivity and pay dynamics under different incentive arrangements. Kandel and Lazear (1992) formalize peer pressure as an endogenous social-preference component of worker utility, deriving conditions under which it raises equilibrium effort above the partnership free-rider level; their analysis applies in the small-group regime where mutual monitoring is effective. Hamilton et al. (2003) document, in a U.S. garment plant's transition from individual to group piece rates, an approximately 14% productivity gain net of worker fixed effects, with high-ability workers self-selecting into early voluntary teams despite earnings losses. Bandiera et al. (2010) provide field-data causal identification, in a UK soft-fruit-picking setting under individual piece rates, that workers' own productivity varies by approximately ±10% with the presence of friends on a given work day; Bandiera et al. (2013) extend this with a within-subject field experiment showing that team-incentive design changes endogenous team composition in ways that simple effort-only models do not capture. Although these settings differ substantially from ours — multi-layer Chilean firms with hundreds to thousands of employees, no measurement of peer ties, team membership, or experimental incentive variation — we cite this literature to flag that within-firm productivity and pay variability depend on team-composition and peer-effects mechanisms that operate beneath the formal reporting structure and that are conceptually distinct from the hierarchical-depth dimension we measure. We do not test these mechanisms in our setting; we flag them as competing-explanation candidates for the salary-volatility-by-level and cross-industry patterns we document.

The empirical study of hierarchies sits inside a longer tradition of business-historical and organization-studies scholarship that documents how organizational form responds to strategic, technological, and environmental imperatives. Chandler (1962) argued that organizational structure follows from strategic choice across DuPont, General Motors, Standard Oil of New Jersey, and Sears Roebuck. Lawrence and Lorsch (1967) independently established, through a comparative empirical study of three U.S. industries (plastics, food, containers) varying in environmental uncertainty, that effective organizational design is contingent on environmental demands – articulating the differentiation/integration framework that became the canonical organization-studies foundation for cross-industry organizational variation. Chandler (1977) extended the strategy-structure tradition with a comprehensive historical demonstration that the modern multiunit business enterprise displaced traditional small-firm structures across U.S. industry between 1840 and 1920 because administrative coordination by salaried managerial hierarchies proved operationally more efficient than market coordination at managing high-volume goods flows. Rumelt (1986) operationalized the structure-follows-strategy proposition empirically across 246 Fortune 500 corporations sampled at 1949, 1959, and 1969, classifying firms by the type rather than degree of diversification and showing that organizational form predicts performance differentials – Related-Constrained firms, which diversified only into businesses sharing a central skill with the core, earned the highest returns on capital and equity, while firms diversifying without such relatedness performed worst. We treat this combined business-historical and organization-studies tradition as the institutional and theoretical foundation for our empirical project: the hierarchies we measure are the descendants of the managerial enterprise these accounts historicized, and the cross-industry heterogeneity we document echoes the contingency claim that hierarchical complexity is environment- and technology-specific.

The dynamic dimension of hierarchy has received increasing attention. Caliendo et al. (2015) document that when French manufacturing firms add a management layer, average wages at all pre-existing layers fall – the new layer absorbs complex tasks previously handled below, shifting knowledge boundaries downward. Caliendo et al. (2018) show that firms entering export markets add layers and reorganize wages across all pre-existing layers. Friedrich (2022) provides complementary evidence from Danish administrative data, showing that trade shocks affect within-firm wage inequality through hierarchy changes: firms that add layers in response to trade exposure experience increased wage dispersion, establishing a direct link between organizational restructuring and inequality. Our monthly panel data allow us to characterize height dynamics at finer temporal resolution than annual administrative data, revealing that within-firm height variation involves boundary-layer shifts with asymmetric consequences: routine churn when firms shrink, but concurrent reporting-chain adjustment when firms grow taller.

3 Data and Measurement

Our data come from Chile, classified as a high-income economy by the World Bank and an OECD member since 2010, with labor markets that are relatively flexible by Latin American standards. We use administrative records from a Chilean human-resources software platform whose payroll and workflow systems serve a broad range of industries and firm sizes. The data were provided under a formal research collaboration with the platform operator. All data were delivered in anonymized form: the working dataset contains no direct identifiers (names, tax identifiers, or addresses) for workers or firms, and geographic detail was limited to the regional level to protect client confidentiality. The platform is used to manage payroll, reporting lines, and core HR processes, and it captures detailed employer–employee interactions at high frequency. Because the platform is used for payroll processing, reporting-link fields are updated regularly as part of the payroll cycle; in practice, HR administrators typically update reporting lines immediately when organizational changes occur, since correct supervisor assignment is required for leave approval and performance review workflows.

Each organizational unit in our data corresponds to a client account subscribed to the HR platform. Client accounts are typically identified by a tax identifier (*Rol Único Tributario*, or RUT), but multi-establishment groups may appear as a single client account or subscribe separately per subsidiary, so the organizational unit in our data should be understood as the platform client rather than necessarily a legally or economically distinct entity.[3]

Several institutional features of the Chilean labor market are relevant for interpreting our data. Dismissal of workers on indefinite contracts (*contrato indefinido*) triggers severance pay (*indemnización por años de servicio*) of approximately one month's wages per year of service, which may contribute to lower turnover at higher tenure levels. The minimum wage rose from CLP 460,000 per month (July 2023) to CLP 500,000 (July 2024), CLP 510,636 (January 2025), and CLP 529,000 (May 2025), roughly USD 530 at current exchange rates, providing context for bottom-level compensation in our data.[4]

All monetary values are expressed in UF-adjusted January 2025 Chilean pesos (CLP), deflated using the Unidad de Fomento (UF), a daily inflation-adjusted unit of account maintained by the Central Bank of Chile.

The underlying database consists of linked employer–employee records at the employee-month level. For each worker and month, we observe: (i) the worker's direct supervisor, via a unique manager identifier, which allows us to reconstruct reporting graphs; (ii) gross monthly wages, base salary, and net pay;[5] (iii) contract type and employment status; (iv) separation events, including codes distinguishing voluntary from involuntary turnover; (v) absence days, including sick leave and other authorized absences; (vi) formal leave spells, including vacation and non-maternity leave; and (vii) basic demographics such as gender and age. For firms, we observe industry classification and, via linkage to tax data, annual

[3] Each organizational unit in our data corresponds to a client account subscribed to the HR platform, identified by a tax identifier (RUT). Multi-establishment groups may subscribe per subsidiary, in which case our data treats each subsidiary as a distinct firm; this could understate group-level hierarchical depth. The external-manager anchoring rule in our hierarchy construction algorithm flags roughly 10% of employees in the underlying data whose reported supervisor is not in the firm's roster, providing a lower-bound proxy for boundary cases.

[4] Minimum wage figures are reported in nominal terms. All other monetary values in this paper are expressed in January 2025 CLP, deflated using the Unidad de Fomento (UF).

[5] Our wage measure is gross monthly salary (*sueldo bruto mensual*) as recorded in the payroll system. This includes the monthly component of the legally mandated profit-sharing payment (*gratificación legal*, Articles 47–52 of the Código del Trabajo) when employers opt for the monthly payment modality. Lump-sum annual gratificaciones, if paid separately, would appear as salary volatility in the within-person CV measures.

sales tier categories.

Platform adoption is not random, and Table 19 benchmarks our analysis sample against the universe of Chilean firms recorded by the Servicio de Impuestos Internos (SII), the Chilean tax authority. We restrict the comparison to firms with at least 20 dependent workers – the segment large enough to sustain the multi-layer hierarchies we study – a universe of 73,340 firms. Of our 3,059 sample firms, 2,407 fall in this range; they cover 3.3% of that universe by firm count. Coverage rises with firm size up to the 250-999 bracket, from 1.8% of firms with 20–49 workers to 6.7%, and then thins again among the very largest, where multi-entity groups are more likely to run payroll on systems of their own (Table 19). The industry and region comparisons that follow are drawn against SII firms with at least five dependent workers, a broader universe than the twenty-worker one behind the coverage table, so they describe how our sample differs from formal Chilean employers generally rather than from firms of comparable size. The sample's geographic footprint is markedly more concentrated in the capital: 83.0% of sample firms are in the Metropolitan Region, against 48.8% in the universe. Its industry mix, likewise, tilts toward knowledge- and service-intensive sectors: information and communications, professional services, and finance are over-represented (10.4%, 10.9%, and 5.6% of sample firms versus 2.8%, 7.2%, and 1.6% in the universe), while construction (6.9% versus 13.9%), agriculture (3.6% versus 6.9%) and above all education (0.2% versus 3.4%) are under-represented (Table 20). Manufacturing, by contrast, is represented almost exactly in proportion (9.6% versus 9.7%). The patterns we document are therefore best read as characterizing formalized, medium-to-large enterprises, most representative of knowledge-intensive and service sectors concentrated in the capital region.

These patterns should not be read as facts about Chilean firms in general, much less about firms in developing economies more broadly. The hierarchy quality filter that defines the analysis sample is the dominant selection step; together with the data-quality screen it retains 28.2% of the employee-months in the analysis window. This internal selection differs in character from the platform-versus-universe comparison above: it works most strongly against the *largest* firms, whose reporting graphs are the least likely to resolve cleanly into a single org chart, as external secondments, dotted-line matrices, and holding-company arrangements leave reporting links unresolved. Firm-months that the filter retains and drops are otherwise well balanced – wages, age, tenure, contract mix, gender composition, and geography all differ by less than a quarter of a pooled standard deviation – so the binding gaps are firm size, where dropped firms are substantially larger, together with separation rates. The analysis sample thus describes the organizational architecture of firms with well-defined reporting structures, and the documented patterns may be specific to this selected population. Appendix A.4 reports the full selection ladder, the criterion-by-criterion balance tables, and the binding-constraint taxonomy; comparisons with the French administrative census of Caliendo et al. (2015) or the U.S. public-firm data of Ewens and Giroud (2025) involve different populations, and cross-country differences in magnitudes should be interpreted with this caveat in mind.

Table 2: **Retained versus dropped firm-months at the hierarchy-quality boundary.**

Variable	Retained	Dropped	Std. diff.
Firm size	150.89	284.03	-0.20
Median wage	1,790,039.46	1,667,731.18	+0.09
Mean age	39.36	39.29	+0.01
Mean tenure (yrs)	3.96	3.61	+0.13
Share female	0.38	0.39	-0.00
Share open-ended	0.86	0.84	+0.11
Share in Metro region	0.83	0.81	+0.06
Turnover rate	0.14	0.17	-0.18

Standardised mean differences between the 70,753 firm-months retained by the hierarchy-quality filter and the 88,018 it drops. The comparison is at the firm-month level, not the firm level: most dropped firm-months belong to firms that survive in other months. Every covariate falls inside the conventional quarter-of-a-standard-deviation band, and the two that come closest to it are firm size (-0.20) and the turnover rate (-0.18): the filter drops larger firm-months with somewhat higher separation rates, which is the direction matrix and holding structures would predict.

Our main analysis sample covers 3,059 firms, 70,753 firm-months, and approximately 10.1 million employee-months – the subset of platform-observed firms with reporting graphs that resolve cleanly enough to permit hierarchical analysis (full sample-flow accounting in Appendix A.4).[6] The sample spans manufacturing, retail, professional services, technology, and other sectors, with firm sizes from fewer than 50 to more than 1,000 employees, providing substantial variation in organizational structure across which to document empirical regularities.

Our analysis window begins in January 2021, after the acute phase of COVID-19 disruptions in Chile. However, the early months of the sample may still reflect recovery dynamics: Chile experienced extended lockdowns through much of 2020, enacted the *Ley de Protección al Empleo* (a furlough law that remained active through mid-2021), and allowed three rounds

[6]Construction applies, in sequence, basic data-quality filters, the analysis-window restriction (2021–2025), and the hierarchy quality filter. A firm-month passes the hierarchy quality filter when it meets three criteria: a hierarchy-resolution rate of at least 95% (at least 95% of employees assigned a hierarchical level), an identifiable unique organizational apex, and no anomalous team structures. The hierarchy quality filter is by far the largest source of firm-month attrition, and concentrates disproportionately in a small number of large firms whose reporting trees do not resolve under our criteria; we present the full selection ladder, the binding-criterion taxonomy, and the size-attribution in Appendix A.4.

of pension fund withdrawals (*retiros de AFP*) in 2020–2021, which reshaped household liquidity and labor supply. Our regression specifications include period fixed effects at the calendar-month level (one indicator per year-month) that absorb common time shocks, but we acknowledge that early-sample patterns – particularly in turnover and compensation – may differ from steady-state behavior. In robustness checks (Appendix B.2), we verify that our main findings are not sensitive to excluding 2021 observations. Splitting the window into 2021–2022 and 2023–2025 yields nearly identical height–size semi-elasticities across the two halves (0.754 vs. 0.737) and wage–level gradients (-0.282 vs. -0.268).

To ensure reliable measurement of hierarchy, we construct a sample using a series of filters that address data quality concerns. These include checks for valid employment relationships, consistency in compensation records, and completeness of reporting links. We assign hierarchical levels by tracing each employee's reporting path to a unique organizational apex, and retain only firm-months with high hierarchy-resolution rates (at least 95% of employees assigned a hierarchical level) and no anomalous team structures. Full details on sample construction are provided in Appendix A.1. The filtering process reduces the raw data (41.6 million employee-months) to 9.6 million regression-ready observations; the hierarchy quality filter accounts for approximately three-quarters of this attrition, with the remaining reduction due to data quality filters, analysis window restrictions, and FTE wage requirements. Appendix A.4 provides a complete stage-by-stage decomposition, the criterion-by-criterion balance comparison of retained and excluded firm-months, and a taxonomy of which criterion binds for each excluded firm; the hierarchy quality filter binds most strongly on the largest firms, whose reporting graphs are least likely to resolve cleanly into a single org chart.

We construct a set of hierarchy, compensation, and internal outcome measures using the administrative reporting links and payroll records in the HR platform. Hierarchical position is defined from the reporting graph, allowing us to measure employee level, firm height, span of control, and the reach of managerial responsibility. Compensation measures distinguish base and variable pay and are standardized to full-time equivalent workers. We also observe a range of internal outcomes – including turnover, absences, and leave – as well as team-level aggregates and standard demographic and firm controls. The construction and validation of these variables, along with the hierarchy assignment algorithm, are described in detail in Appendix A.2.

We characterize each firm-month organizational chart using a small set of interpretable measures:

- **Firm height:** the number of hierarchical levels from the CEO (level 1) down to the deepest level of employees in the reporting graph.

- **Employee level:** an employee's distance from the top of the hierarchy (CEO = 1; larger numbers indicate lower rank).
- **Span of control:** the number of direct reports, capturing direct supervisory responsibility.
- **Reach:** the size of an employee's subordinate subtree (the number of employees below them in the reporting graph), summarized as $\log(\text{Reach} + 1)$.

Figure 1 illustrates how these measures are defined on an actual reporting graph observed in the HR system.

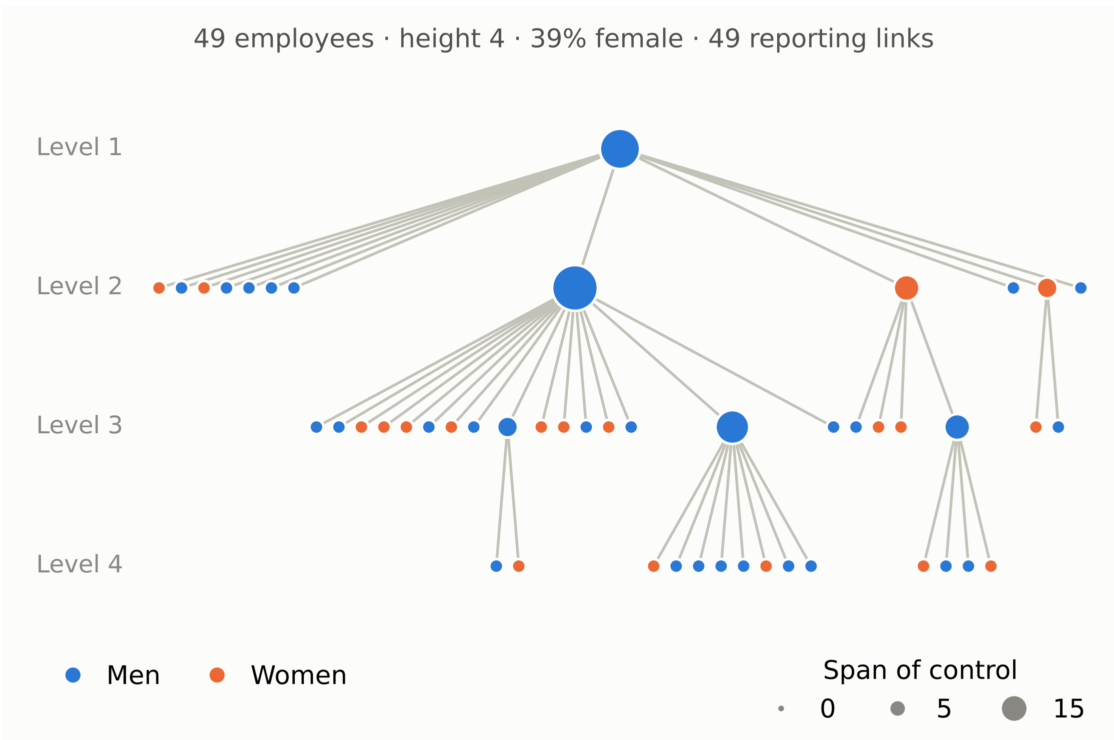

Figure 1: **Example organizational chart.** A representative reporting structure observed in the administrative HR platform. Nodes are employees and edges denote direct reporting links. The figure illustrates how we define firm height (maximum level), employee level (distance from the CEO), and span of control (out-degree, shown here as node area); node colour marks gender.

Table 3 summarises the estimation sample. The average firm has 140 employees (median 41), with substantial heterogeneity: the 95th percentile is 559 employees and the maximum is 9,417.

Throughout the analysis, we present height-disaggregated results for firms with 3 to 6 hierarchical levels, which account for approximately 85% of firm-months in the analysis sample. Firms with fewer than 3 levels have insufficient vertical differentiation for meaningful

Table 3: **Summary statistics for firms, teams, and workers.**

	N	Mean	SD	Q1	Median	Q3
Panel A: Firm-Level Statistics						
Firm Size (employees)	70,753	139.8	352.0	19	41	117
Firm Height (levels)	70,753	4.1	1.4	3	4	5
Panel B: Team-Level Statistics						
Team Size (incl. manager)	1,285,243	8.68	16.15	3.0	5.0	9.0
Subtree Reach (raw)	1,285,243	26.14	110.00	3.0	6.0	16.0
log(Subtree Reach + 1)	1,285,243	2.19	1.21	1.39	1.95	2.83
Manager Age	1,285,234	43.69	9.79	36.0	43.0	50.0
Manager Tenure (years)	1,285,243	7.34	6.88	2.2	5.2	10.4
Manager Wage (CLP '000)	1,280,388	4,325.4	2,926.6	2,140	3,437	5,583
Manager Wage (USD)	1,280,388	4,325.4	2,926.6	2,140	3,437	5,583
Panel C: Worker-Level Statistics (Firm-Weighted)						
Employee Level	9,948,827	3.16	1.19	2.00	3.00	4.00
Age	9,948,567	39.28	11.63	30.00	37.00	47.00
Tenure (years)	9,948,827	3.63	5.05	0.50	1.75	4.42
Wage (CLP '000)	9,646,426	2,167.1	2,073.4	915	1,411	2,551
Wage (USD)	9,646,426	2,167.1	2,073.4	915	1,411	2,551

[a] Monthly gross salary, full-time-equivalent sample only, deflated to 2025 pesos and winsorized at the 1st/99th percentiles within period. [b] USD conversion at 1,000 CLP/USD for approximate reference; at this round rate the USD figures restate the thousands-of-pesos row one for one. *Notes:* Statistics correspond to the estimation sample (2021–2025). Panel A: N = 70,753 firm-months across 3,059 unique firms; firm size is the number of employees assigned a hierarchical level in the firm-month. Panel B: N = 1,285,243 team-months across 88,352 unique managers, defined as workers with at least one direct report; 34.6% of manager-months are female. Teams are unweighted, so the panel describes the typical team rather than the team of the typical firm. Panel C: firm weighted — each worker enters with weight $1/(n_{ft}T_f)$, so every firm contributes equally regardless of size or of how long it is observed. This is the convention used for descriptive exhibits throughout, and differs from the firm-month weights used in the regressions; Table 30 restates the panel on that basis. N = 9,948,827 worker-months across 3,059 firms; wage statistics use the FTE sample (9,646,426 worker-months, 3,058 firms, one firm short because it has no full-time worker with a valid wage). Age rows drop 260 worker-months and 9 team-months whose recorded age falls outside [15, 100] years.

height comparisons, while firms with 7 or more levels represent a small fraction of the sample (146 firms by modal height, 5.4% of firm-months) and exhibit greater measurement variability due to limited cross-sectional coverage. Summary statistics for these extreme heights are available from the authors.

3.1 Height Dynamics and Measurement Stability

Within-firm height variation comes from two distinct sources, which we want to separate explicitly. Some of the variation reflects genuine organizational change – firms adding or removing layers as they grow, restructure, or respond to environmental shocks. Other variation reflects routine personnel turnover at the firm's lowest level: when the most junior employees leave or are added, the firm's measured height shifts by one even though no organizational restructuring has occurred. We treat both sources as contributing to within-firm height dispersion; for our headline within-firm coefficients, the conservative interpretation is that boundary-layer turnover attenuates rather than biases estimates, yielding lower-bound coefficients. We address this with a comprehensive height dynamics analysis, summarized here with full details in Appendix B.4.

Firm height is primarily a between-firm characteristic: the intraclass correlation coefficient (ICC) is 0.892, meaning approximately 89% of height variance lies between firms. Among the 61% of firms (1,868 of 3,059) that change height at some point during the panel, variation is modest – the median range among changers is one level, and transitions occur in only 7.6% of consecutive *observed* firm-month pairs, or 6.29% of strictly adjacent calendar-month pairs.

The source and persistence of these transitions reveal a consistent pattern. Of the 4,017 observed height transitions, 54.3% are persistent (the new height is maintained for at least three consecutive months) while 45.7% are transient, not holding the new height for three consecutive months. All 3,853 unit (± 1) height changes involve the boundary layer, which is mechanical rather than a finding[7] – the appearance or disappearance of workers at the deepest tier of the reporting tree.

However, the consequences of height transitions for existing workers differ markedly by direction. Tracking employees present both before and three months after height transitions that hold for two months (349,328 worker-events across 2,329 events), we find that height *decreases* are consistent with routine boundary departures: 82% of workers maintain their hierarchical level, and the 14% who shift upward are concentrated at the vanishing bottom layers. Height *increases* differ: 77% of workers maintain their level, but 20% are displaced downward.[8] Among these displaced workers, 62% retained their direct supervisor, suggesting that much of the level shift reflects mechanical relabeling as the tree deepens beneath them – a necessary but not sufficient condition for pure relabeling, since some supervisor-preserving

[7] Of the 4,017 total height transitions, 3,853 are unit (± 1) changes, of which 2,260 are increases and 1,593 decreases. Because firm height is defined as the maximum employee level, a one-layer change must by construction occur at the deepest tier; we report the split for completeness.

[8] The tracking analysis requires the new height to hold for the two months following the transition – a slightly weaker criterion than the three-consecutive-month persistence rule of Table 27, which classifies 2,180 transitions as persistent; the tracked-event count is therefore larger. The three-month comparison window captures some level changes unrelated to the height transition itself. The direction-dependent asymmetry (20% displaced downward during increases against 4% during decreases) is too large to reflect normal organizational flux alone.

changes may involve genuine shifts in responsibility. The remaining 38% experienced both a level change and a manager change, consistent with concurrent reporting-chain adjustment – particularly at Level 2, where 78% of displaced workers also changed managers.

These findings suggest a two-regime interpretation. Height decreases are predominantly routine personnel churn at the organizational periphery. Height increases, while triggered at the boundary, are accompanied by broader organizational adjustment: a mix of mechanical relabeling (the tree growing deeper beneath stable positions) and concurrent restructuring (reporting-chain changes at upper-middle levels). The contrast with Caliendo et al. (2015), where "adding a layer" means hiring the first worker in a new occupational category (a qualitatively different event), underscores that measurement definition shapes what counts as organizational adjustment. Additionally, annual data may mask precisely the kind of monthly churn documented here; some of what Caliendo et al. classify as "reorganization" could be the cumulative effect of many monthly boundary shifts.[9]

Contract composition is consistent with this interpretation, though the gradient is modest: permanent contracts (*contrato indefinido*) account for 99.3% of Level-1 employee-months, declining to 77–81% at the deepest levels, with fixed-term and project-based (*obra*) contracts correspondingly more common toward the bottom of the hierarchy. The deepest tiers thus carry a higher share of workers on temporary arrangements, consistent with the interpretation that height fluctuations reflect routine hiring and separation of peripheral workers rather than structural reorganization.

This bottom-layer churn has important implications for estimation. Because within-firm height variation is primarily measurement noise from personnel turnover at the boundary, it attenuates within-firm estimates toward zero rather than introducing systematic bias. Significant coefficients in our firm-fixed-effects specifications should therefore be interpreted as conservative lower bounds on the true height–outcome relationship. A stable-subset check is consistent with this interpretation, but the two headline facts move in opposite directions and both movements are informative. Restricting to firms whose height range over the panel is ≤ 1 (2,435 of 3,059 firms, 79.6%) yields a steeper wage gradient (-0.282 vs. -0.265), as attenuation predicts: there level is a regressor, and classical measurement error in a regressor biases its coefficient toward zero. The height–size semi-elasticity instead flattens (0.723 vs. 0.741). This is not a contradiction. In Fact 1 height is the *dependent* variable, so measurement error enters the residual and inflates its variance without attenuating the slope; the small decline reflects the restriction to a more homogeneous set of firms rather than the

[9] Of the 4,017 transitions, 164 (4.1%) involve height changes of magnitude two or more. These larger jumps plausibly reflect more substantial organizational events than single-level boundary adjustments; characterising them is a natural extension we leave for follow-up work.

removal of noise from a regressor. We report both because the attenuation argument applies to the within-firm coefficients only (see Appendix B.4 for full details, including the height transition matrix, seasonal patterns, and persistence duration statistics).

4 Results

We prioritize magnitudes and cross-sectional patterns throughout. Each empirical statement is linked to a specific figure or table. Two weighting conventions appear in the paper and the distinction is load-bearing. Descriptive statistics and figures are *firm-equal* weighted: a worker-month enters with weight $1/(n_{ft}T_f)$, where n_{ft} is firm f's headcount in month t and T_f the number of firm-months it contributes, so every firm carries total weight one regardless of its size or how long it is observed. This is the convention that matches the descriptive question — what the typical firm's hierarchy looks like. Regression tables instead weight each *firm-month* equally, with weight $1/n_{ft}$, so a firm contributes in proportion to the months it is observed; this is the convention under which the estimation sample and the standard errors are defined. The choice is not innocuous. Firms observed for longer are larger, and larger firms are less likely to have a woman at the top: the female share of level-one positions is 23.3–11.6% across heights under firm-equal weights and materially lower when each employee-month counts equally. Appendix C reproduces the summary statistics under the firm-month convention so that the sensitivity of the size, tenure and wage rows is visible rather than asserted; the female-share figures are reported under both conventions in Section 4.5. Wage measures are restricted to full-time equivalent workers to ensure comparability. We report conditional correlations rather than causal effects: hierarchical position, team assignment, and managerial characteristics are endogenous outcomes of matching and sorting processes we do not model, so phrasings such as "is associated with" should be read as descriptive of the conditional pattern rather than as causal claims.

To organize the descriptive exercise, Table 4 maps specific predictions from the hierarchy literature to our empirical findings.

Table 4: **Theoretical predictions and empirical findings.**

Theory	Prediction	Finding
Garicano (2000); Garicano and Rossi-Hansberg (2006); Garicano and Rossi-Hansberg (2015); Garicano and Hubbard (2007); Garicano and Hubbard (2016)	Firms add layers as they grow; richer problem distributions and higher coordination requirements support deeper hierarchies	Consistent on scaling: each doubling of size adds 0.51 layers (pooled), positive in every industry. Semi-elasticity ranges from 0.90 (agriculture) to 0.65 (administrative services) among precisely estimated sectors, or 0.62 to 0.45 layers per doubling; the cross-industry ordering does not follow knowledge intensity; coordination demands are one reading, untested here
Rosen (1982); Lazear and Rosen (1981); Calvo and Wellisz (1978); Garicano and Rossi-Hansberg (2006); Garicano and Hubbard (2016)	Convex wage profiles arise via assignment-equilibrium with multiplicative spillovers (Rosen 1982), tournament prize structures (L-R 1981), or hierarchical supervision economics for ex-ante identical workers (Calvo-Wellisz 1978, eq. 29); knowledge leverage at the top (GRH 2006; G-H 2016) amplifies pay differentials over larger teams	Consistent: CEO-to-median ratios rise from 3× (height-3) to 6× (height-6); wage profiles are strongly convex

continued on the next page

Table 4, continued

Theory	Prediction	Finding
Lazear and Rosen (1981); Prendergast (1999)	Tournament prize structures truncate the binomial winner-loser variance at the top; piece-rate-style contracts at lower levels carry larger common-noise variance	Not testable on these data: the CV is computed on *monthly payroll*, which excludes the annual and deferred components through which executive incentive pay is typically delivered (Section 4.2). What we document is month-to-month take-home risk, and it is lowest at the top (CV ≈ 0.03–0.04) and highest at middle-to-lower levels (CV ≈ 0.10–0.14). Whether *total* compensation risk rises with rank cannot be determined without the off-payroll components
Garicano (2000); Garicano and Rossi-Hansberg (2006); Garicano and Rossi-Hansberg (2015)	As firms deepen, hierarchies feature management by exception, with lower layers handling routine decisions and higher layers handling rarer problems	Partly: spans at the deepest managerial rung widen systematically in taller firms, but CEO spans do not narrow — they are flat across heights (Figure 4)
Bolton and Dewatripont (1994)	Span of control decreases with hierarchical level as higher-layer agents absorb longer aggregated reports (Proposition 4); delegation arises when the apex is overloaded (Proposition 1)	Only in short firms. The prediction is about span against *level*; within height-3 firms the level gradient in log(span) is -0.518, as predicted, but by height 6 it is 0.055 (SE 0.012), positive and precisely estimated, so the profile inverts rather than flattening (Section 4.3). The apex-span version of the prediction is a different comparison, and it does not hold either: CEO spans are flat across heights

continued on the next page

Table 4, continued

Theory	Prediction	Finding
Radner (1993)	Trees (hierarchies) are efficient architectures for decentralized information processing under processor-capacity constraints (Theorem); decision-delay grows as $1 + \lceil \log_2 N \rceil$, slower than linearly in firm size (Corollary 1)	Partly. The pyramidal headcount distribution matches (Figure 5), and depth grows far more slowly than size. But the corollary is specific: $1 + \lceil \log_2 N \rceil$ implies exactly one additional layer per doubling of employment. We observe 0.51, about half of that. Firms are flatter than a balanced binary tree would be — which is the content of Fact 1, not a confirmation of the benchmark
Calvo and Wellisz (1978)	Monitoring costs create a trade-off between layers and spans	Partly: the CEO tier thins and managerial intensity drops at lower levels in taller firms, but the base tier thins too — the shape becomes a diamond, not a sharper pyramid (Figure 5)

Notes: Each row maps a theoretical prediction to the corresponding empirical pattern documented in this paper. "Consistent" indicates the data exhibit patterns in the direction predicted by the theory in a single-country convenience sample. These are conditional patterns whose generalizability remains to be established.

4.1 Hierarchical Structure: Height, Spans, and Pyramids

We begin by documenting heterogeneity in hierarchical depth, the distribution of supervisory responsibility, and the pyramidal shape of organizations.

4.1.1 Height Heterogeneity and Scaling

Firms in our sample exhibit substantial variation in hierarchical depth. Figure 2 (Panel A) shows that while height-4 firms are most common, the distribution spans from flat organizations with just 2–3 levels to deep hierarchies with 7 or more layers.[10] This dispersion is notable given that firm size explains only part of the variation: firms of similar size can differ substantially in the number of layers they employ.

Hierarchy height scales systematically with firm size, but the relationship is far from mechanical. Figure 3 plots firm height against log firm size. In levels, the slope means a doubling of employment buys about 0.51 of a layer, so a firm must quadruple in size to add one. The fitted relationship accounts for meaningful variation, but substantial residual heterogeneity remains: firms of similar size organize into hierarchies of very different depths. This suggests that hierarchy height reflects strategic design choices beyond what firm size alone would dictate.

4.1.2 Spans of Control

A second fundamental dimension is the span of control: how many direct reports managers supervise. Figure 2 (Panel B) shows that spans are strongly right-skewed. A substantial share of managers supervise very small teams: 21.4% of managers have exactly one direct report and 63% supervise four or fewer, weighting each firm equally; weighting each manager equally instead gives 15.5% and 52%.[11] Yet the mean span is materially larger than the median because a small number of managers oversee very large teams, with the 95th percentile at 15 direct reports and 74 manager-months above 500. This long tail is an important descriptive regularity in its own right. It suggests that any theory of hierarchy that relies on a "typical" span must confront substantial heterogeneity in supervisory scope.

The distribution alone, however, does not reveal how spans are organized across levels. Figure 4 shows average spans by employee level, separately for firms of different heights.

[10] Panel A plots firm-months, not firms, and the axis runs from 2 to 10. Assigning each firm its modal height instead gives an observed range of 2 to 9 layers. Measured instead at the firm-month level, twelve of 70,753 firm-months reach ten or eleven layers. We report the firm basis throughout because the descriptive question concerns firms rather than months of observation.

[11] The description of the bulk of the distribution does not depend on the weighting choice; the two bases differ more in the tail, where managers of very large teams are concentrated in the largest firms.

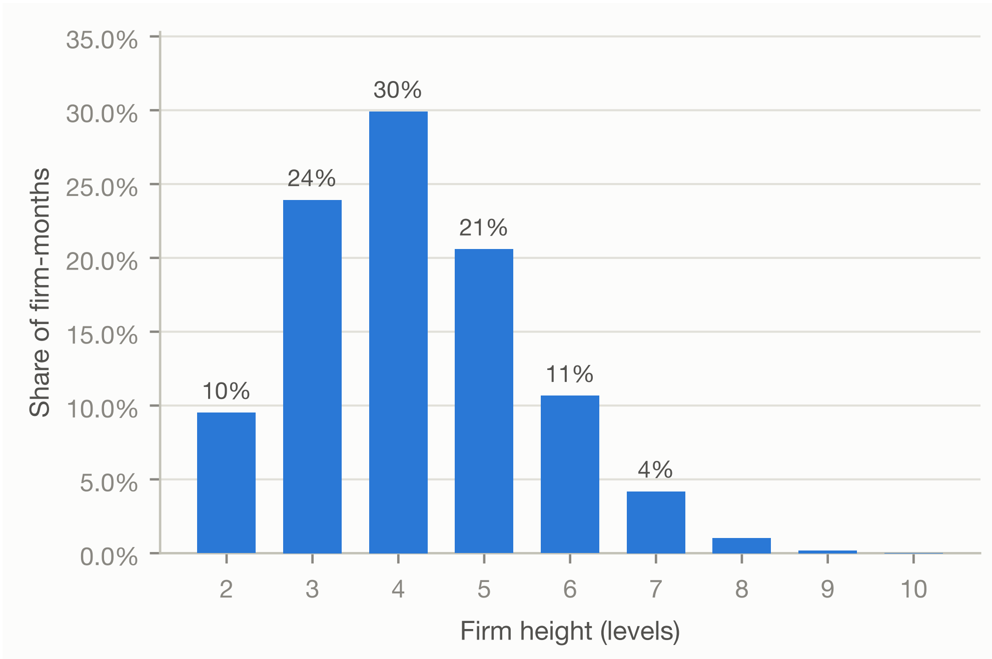

(a) Distribution of Firm Heights

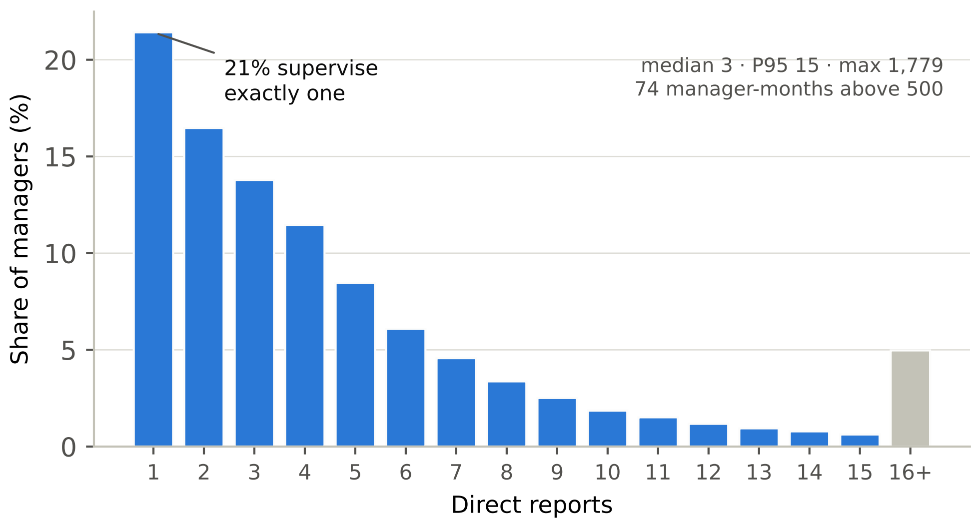

(b) Distribution of Manager Spans of Control

Figure 2: **Distributions of hierarchy height and spans of control.** Panel (a) plots the distribution of firm heights in the sample. Panel (b) plots the distribution of manager spans of control (number of direct reports), weighting each firm equally. Bars give the share of managers at each span from one to fifteen; the final bar collects all larger spans. The distribution is single-peaked at one direct report and declines monotonically, with a median of 3 and a long right tail that the bars cannot show at a useful scale: the 95th percentile is 15, the maximum is 1,779, and 74 manager-months exceed 500 direct reports.

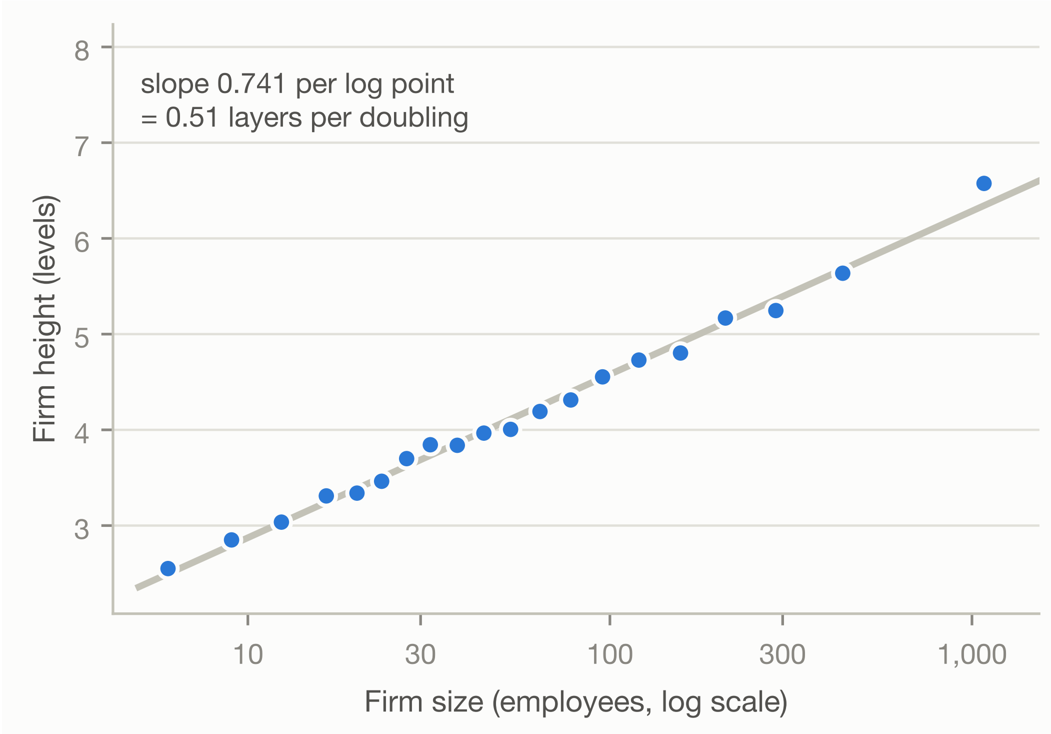

Figure 3: **Hierarchy height scales with firm size.** Firm height against log firm size. Each point is the mean height within one of twenty equal-sized bins of log employment, over firm-months with at least five employees — below that, height is close to mechanical. The unit of observation is the firm-month and no weights are applied, so each firm-month counts once; the line is the fitted regression on the underlying firm-months, not on the bins. The pooled height–size semi-elasticity is $\hat{\beta} = 0.741$ (SE 0.013, clustered by firm), implying that each doubling of firm size adds approximately 0.51 layers. This bivariate OLS on all firm-months is the headline estimate used throughout; Table 11 re-estimates it separately within each industry.

Two patterns stand out.

First, CEO spans are remarkably stable across firm heights. The chief executive supervises between 5.2 and 6.2 direct reports at every height we examine: 6.2 on average in height-3 firms and 6.0 in height-6 firms, dipping to 5.2 in between. A simple delegation logic would predict that adding layers lets the top executive push supervision downward and narrow their own span; we do not find that. What changes with height is where in the hierarchy spans are widest, not how many people report to the person at the top.

Second, the allocation of spans across levels becomes more *layered* as height increases. In taller firms, upper-middle managers tend to have relatively narrow spans, while lower-level managers supervise larger teams. For example, in height-6 firms, mean spans run 4.7 at level 2, 6.6 at level 4 and 6.0 at level 5. In other words, growth in firm size is not managed only by adding layers; it is accompanied by a systematic reallocation of supervisory scope down

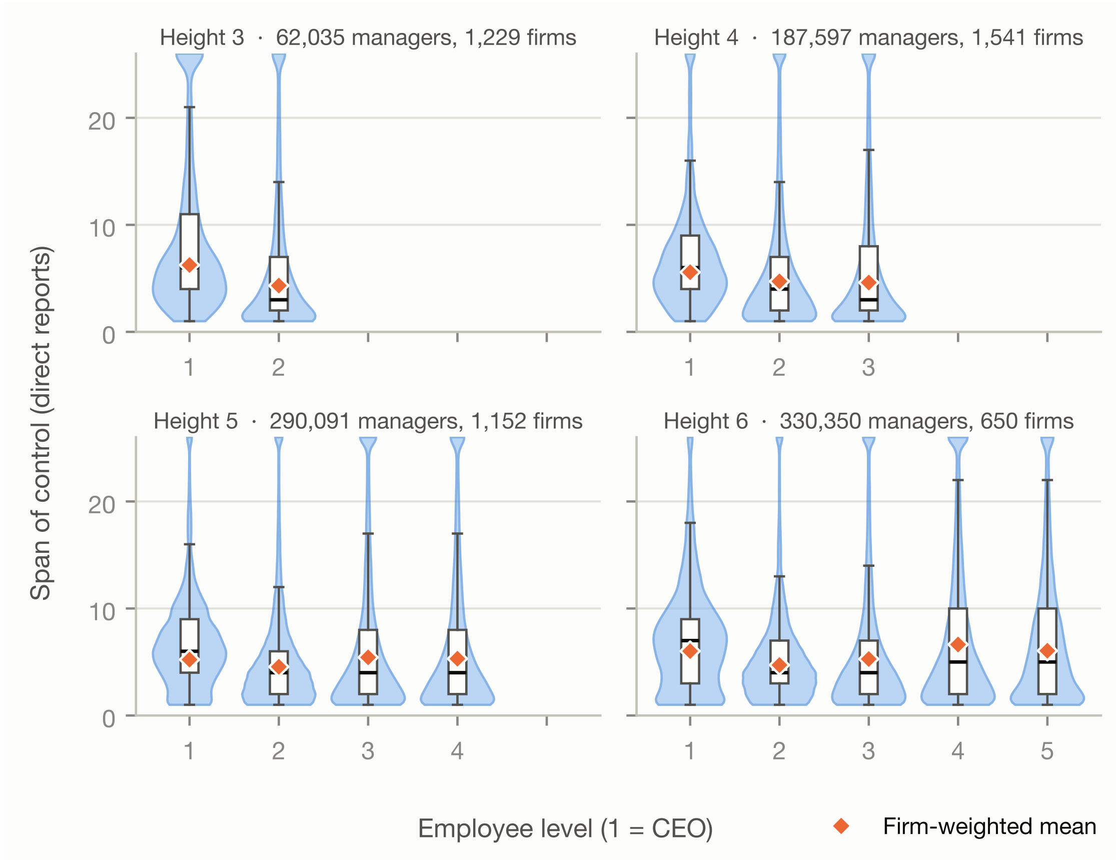

Figure 4: **Span of control by employee level and firm height (heights 3–6).** Mean and median spans (direct reports) by employee level, plotted separately for firms of different heights. The sample is restricted to employees who supervise at least one worker, so each box describes managers rather than all employees at that level. The violin and the box are unweighted and winsorised at 26 for display; the diamond is the firm-equal weighted mean, computed on unwinsorised spans, and is the figure quoted in the text. CEO spans are close to flat across heights (6.2 to 6.0 direct reports on average, dipping to 5.2 in between). The box medians here are unweighted; Table 34 reports firm-equal weighted medians for the same cells and they run lower, so read the level rather than the exact median. What moves with height is the shape below the apex: managerial spans at lower levels expand as firms deepen.

the hierarchy.

This pattern parallels themes in both Ewens and Giroud (2025) and Caliendo et al. (2015). The former document that, in large U.S. public firms, the share of employees at each layer yields pyramidal organizations. The latter show how French manufacturing firms organize production hierarchies, with layers expanding and knowledge becoming more specialized as firms grow. Both emphasize that layers and spans jointly determine the "shape" of corporate hierarchies. Our data reinforces these points using direct reporting links across a broader firm population and shows explicitly how span reallocation differs by firm height.

4.1.3 Pyramidal Structure

A classic depiction of hierarchy is the pyramid: few at the top, many at the bottom. Figure 5 makes this precise in two complementary ways — and qualifies it. Only height-3 firms are pyramids in the strict sense, with the base as the widest tier; from height 4 upward the widest tier sits one rung above the base, and the bottom tier's share of headcount falls monotonically with height. The shape is a diamond that becomes more pronounced as firms deepen.

Panel A reports the share of employees at each level by firm height. The top is sharp; the base is not. In height-6 firms, the CEO level accounts for less than 1% of employees, while the bottom two levels account for a majority of headcount. In height-5 firms, the bottom two levels similarly contain the bulk of employees, while levels 1–2 account for a small minority. Shorter firms are less extreme: in height-3 firms, the top level accounts for a nontrivial share of employment, reflecting that a "CEO level" in small organizations includes more functional heads and fewer layers of delegation.

Panel B reports *managerial intensity*: the percentage of employees at a level who manage at least one direct report. This measure adds a second dimension of structure. It shows that managerial status is not synonymous with being above the bottom; rather, the manager/non-manager boundary shifts with firm height. In taller firms, a larger fraction of employees in intermediate levels are managers, and the transition to predominantly non-managerial roles occurs at a lower rank. For example, in height-6 firms, managerial intensity remains high through level 3, drops substantially at level 4, and is close to zero at the bottom level. In height-4 firms, managerial intensity declines more quickly, consistent with fewer layers of supervision.

These pyramids are descriptive constraints on theory. Information-based and knowledge-based models often imply that decision-making is concentrated at the top with increasing specialization and routinization down the hierarchy (Garicano, 2000; Garicano & Rossi-Hansberg, 2006). Incentive- and assignment-based models emphasize, respectively, promotion tournaments and rank-based pay (Lazear & Rosen, 1981) and competitive market assignment of heterogeneous talent to ranks with multiplicative-productivity spillovers (Rosen, 1982). Both classes of models are consistent with pyramids, but they differ in what they predict about the distribution of managers vs. non-managers and the evolution of these structures with firm scale. The joint depiction in Figure 5 provides a factual target: in practice, firms adjust the *composition* of levels (managerial intensity) in addition to the *size* of levels (headcount).

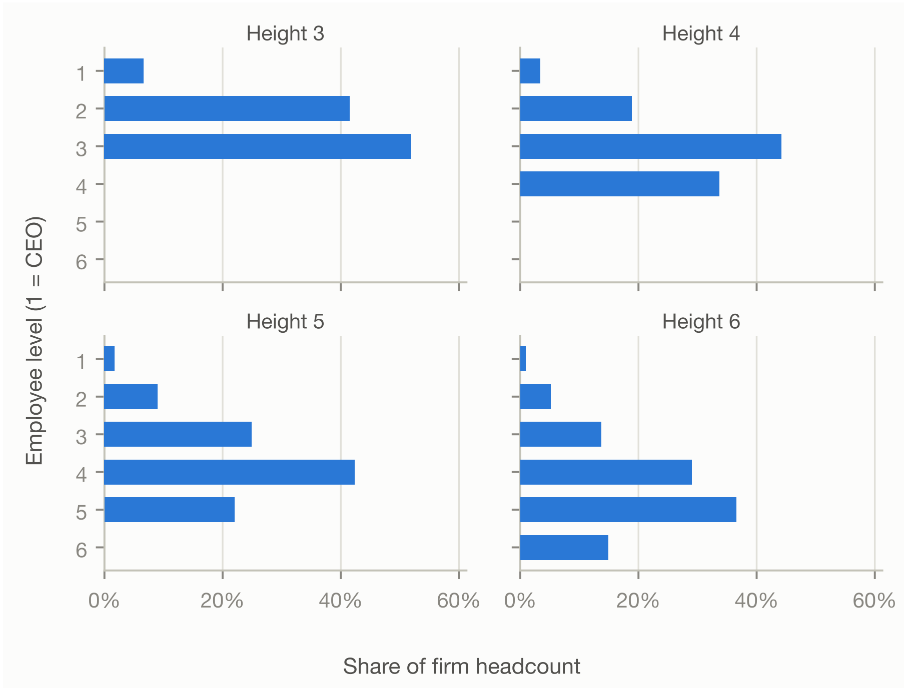

(a) Employment Shares by Level

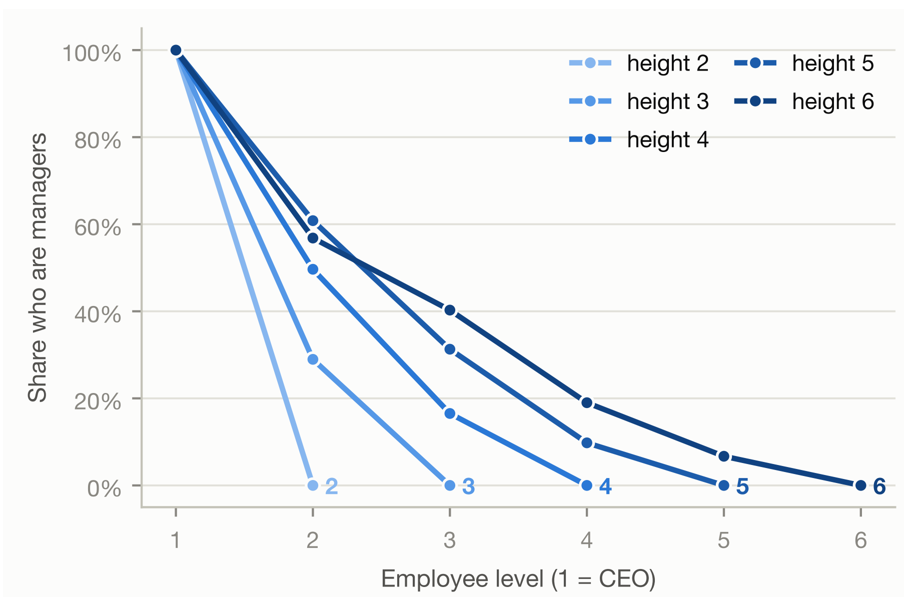

(b) Managerial Intensity by Level

Figure 5: **Pyramids in headcount and managerial intensity.** Panel (a) shows employment shares by employee level for firms of heights three to six. Panel (b) shows managerial intensity (percent of employees who are managers) by level, and runs from height two because the extensive margin is defined there even though the span gradient is not. Both panels weight each firm equally. Together, these panels describe the "shape" of organizations in terms of both headcount distribution and the allocation of managerial roles.

4.2 Compensation: Wage Levels, Convexity, and Variable Pay

How does pay vary with hierarchical position? We document three key patterns: sharp wage stratification by level, convexity in the wage profile, and systematic variation in pay volatility across layers.

4.2.1 Wage Stratification by Level

Compensation is among the most visible outputs of hierarchical design – central to incentive systems and a key source of within-firm inequality. Figure 6 documents the first two dimensions of wage structure: (i) the distribution of median wages by hierarchical level and (ii) CEO-to-median pay ratios by firm height. Figure 7 complements these with the allocation of the wage bill across levels.

Formally, for employee i in firm f at month t, we estimate

$$\log w_{ift} = \alpha_f + \beta_1 \cdot \text{level}_{ift} + \beta_2 \cdot \text{span}_{ift} + \beta_3 \cdot \text{reach}_{ift} + \mathbf{X}'_{ift}\gamma + \delta_t + \varepsilon_{ift},$$

where level_{ift} denotes the hierarchical level of employee i (with 1 = CEO), span_{ift} is the number of direct reports, reach_{ift} is the size of the subordinate subtree, α_f are firm fixed effects, δ_t are calendar-month effects, and $\mathbf{X}_{ift}$ collects demographic controls.

Panel A shows that median monthly wages fall sharply as one moves down the hierarchy. In height-5 firms, median CEO pay is approximately CLP 7.9 million against CLP 1.1 million at the bottom level; in height-6 firms the corresponding medians are CLP 10.1 million and CLP 1.1 million. The steps narrow monotonically toward the bottom: in height-6 firms the ratio falls from 2.65 at the apex to 1.09 between the two lowest levels, so the bottom rung is worth almost nothing in proportional terms.

Panel B aggregates these differences using CEO-to-median worker pay ratios. The median ratio in the sample is 4.2, but it increases systematically with firm height: from 3.15 in height-3 firms to 6.34 in height-6 firms. The relationship is not strictly mechanical, but the overall pattern is clear: deeper hierarchies are associated with greater top-to-middle wage dispersion. This parallels findings from Ewens and Giroud (2025), who document similar upward-sloping relationships in U.S. public firms. Their sample covers roughly 3,100 publicly traded U.S. firms using résumé data on about seven million employees, with a mean firm height of 9.5 layers – substantially taller than our mean of 4.1, a difference largely explained by the much larger average size of U.S. public firms. Despite this gap in absolute depth, both studies find that CEO-to-median pay ratios rise with hierarchy height and that taller firms display wider top-to-bottom wage dispersion. Pay levels and ratios are considerably higher in their setting, reflecting differences in country, firm type, and measurement approach (in-

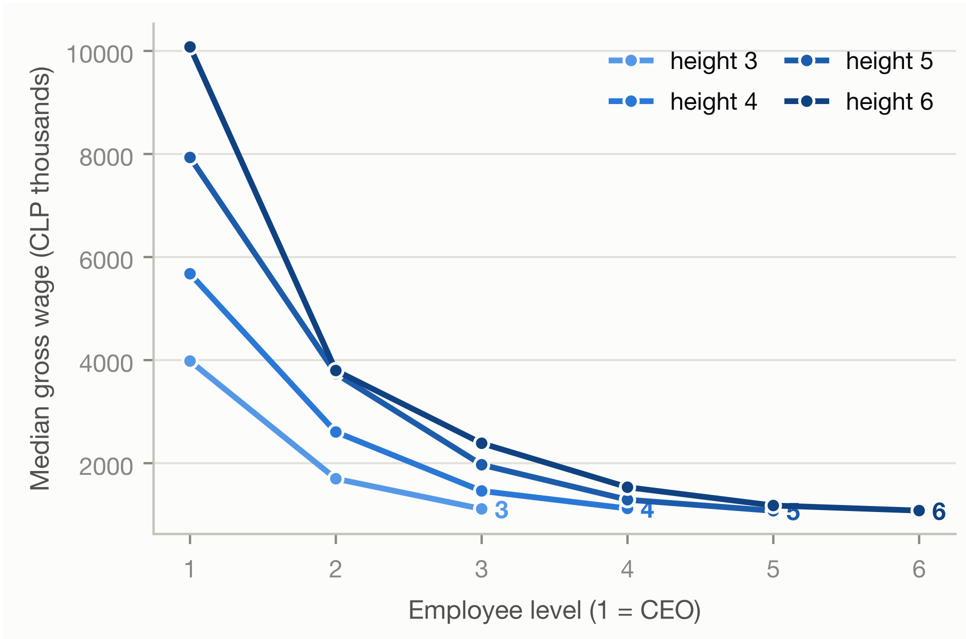

(a) Median wages by level and height

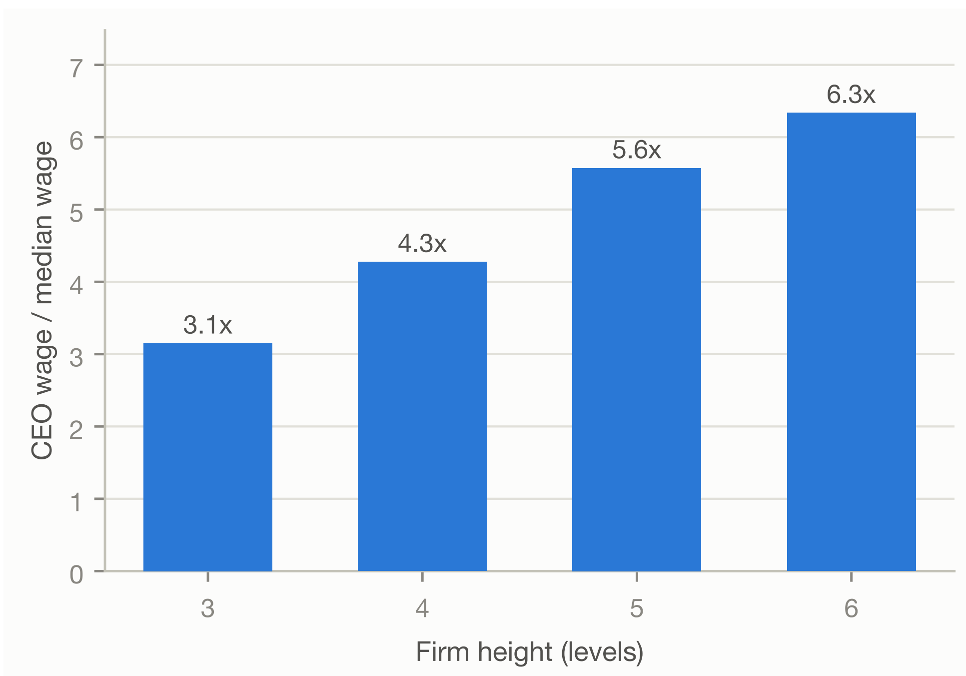

(b) CEO-to-median worker pay ratios

Figure 6: **Wage stratification and CEO pay ratios.** Panel (a) reports median wages by employee level and firm height (FTE sample). Panel (b) reports CEO-to-median worker pay ratios by firm height: the median across firm-months of the within-firm-month ratio. Together, these panels describe level-based wage stratification and how top-to-middle pay dispersion varies with hierarchy height.

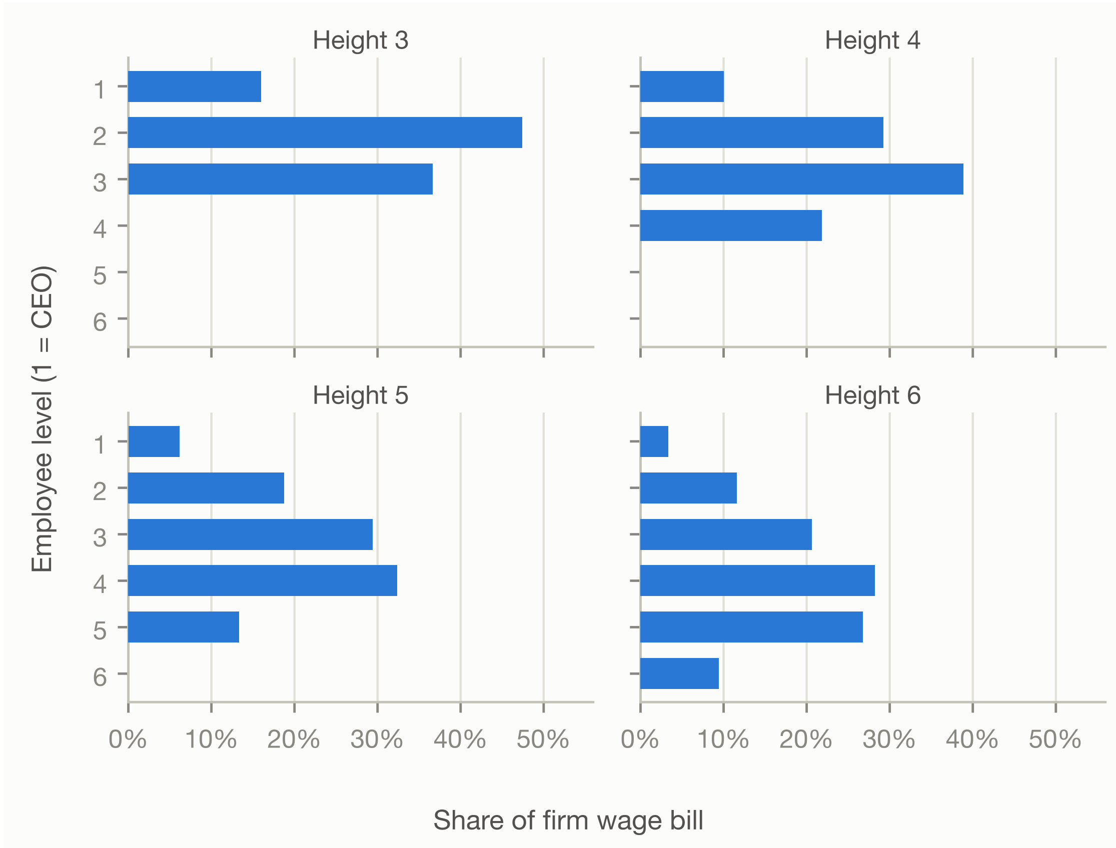

Figure 7: **Distribution of the firm wage bill across hierarchical levels.** The figure reports the share of total firm wages allocated to each hierarchical level, separately for firms of different heights. In short firms, the CEO level accounts for a substantial share of total wages; in taller firms, the middle layers absorb the bulk of compensation.

ferred from job titles vs. directly observed reporting links). The directional associations are nonetheless consistent, suggesting that the structural relationship between hierarchical depth and compensation convexity generalizes across institutional contexts. CEO-level compensation may also be affected by multi-RUT structures (Section 3): if the true organizational apex sits above the data boundary, the observed "CEO" is a subsidiary manager whose compensation understates the group-level pay gradient.

Figure 7 examines how compensation is allocated across levels. In short firms, top-level pay accounts for a substantial share of total wages – for example, the CEO level represents about 16% of the wage bill in height-3 firms. In taller firms, this share falls sharply: CEOs account for about 3% of the wage bill in height-6 firms. This reflects not only the presence of more levels, but also the growing weight of the middle of the organization – layers that combine headcount scale with meaningful compensation.

These patterns illustrate how hierarchical design structures inequality. A high CEO-to-median ratio does not necessarily imply a high CEO share of total pay, especially in large firms where the CEO is one among thousands. Conversely, in smaller organizations, the CEO may represent a disproportionate share of the wage bill even with a moderate ratio. Figures 6 and 7 make this tradeoff visible, highlighting how organizational depth is associated with both relative and absolute pay distributions within firms. Table 5 quantifies wage convexity by reporting ratios between adjacent levels. The steepest proportional wage increase occurs at the CEO transition (Level 1 to Level 2) at every firm height – from 2.35× in height-3 firms to 2.65× in height-6 – and ratios decline monotonically toward the bottom of the hierarchy, reaching 1.09–1.53× at the deepest transitions. The apex premium does not moderate as hierarchies deepen. It is not monotone either: the ratio dips to 2.12× at height 5 before returning to 2.65× at height 6, so what the table supports is that the apex step stays steep at every height, not that it grows with depth. The transition into upper management remains steep alongside it (1.59–1.90× at Level 2 to Level 3 in heights 5–6), so deep firms concentrate large proportional rewards across the top two transitions rather than at the apex alone. This steepening of rewards toward the top is consistent with Rosen (1982)'s assignment logic, in which control amplification concentrates rewards in the highest ranks. The pattern is also jointly consistent with Fox (2009), who shows that firm-size wage premiums increase with hierarchical position in Swedish data, and with our rising CEO-to-median ratios in deeper (and typically larger) firms.

An institutional caveat. Chile's minimum wage binds for a substantial share of the bottom of the distribution over the sample period, which would compress wages at the base of the hierarchy and so widen the measured convexity of the profile. We do not measure the size of that channel, and it is a reason to read the level of convexity with care; it does not bear on the comparison across firm heights, which is what the compression results turn on.

Table 6 shows that within firms, workers with broader scope – measured by reach – earn significant wage premia, even conditional on level. This supports incentive-theoretic interpretations of compensation as a function not just of vertical rank, but also of the complexity and responsibility embedded in a role. The explanatory power of hierarchy variables is substantial: adding level, height interactions, span, and reach to a baseline model with demographic and firm-level controls and firm and period fixed effects increases the within-firm R^2 from 0.107 to 0.413 – an incremental gain of 30.6 percentage points. By contrast, the incremental contribution of hierarchy to behavioral outcomes is smaller (within-R^2 rises by less than 1 percentage point for turnover and absence), consistent with the fact that rare binary events are inherently harder to predict.[12]

[12] Within-R^2 values for binary outcomes (voluntary turnover, absence, leave) should be interpreted with

Table 5: **Wage convexity: Median wage ratios between adjacent hierarchical levels.**

Height	Level Transition	Wage (Level N)	Wage (Level N+1)	Ratio
Height 3				
	$1 \rightarrow 2$	3,982,268	1,698,125	2.35
	$2 \rightarrow 3$	1,698,125	1,112,929	1.53
Height 4				
	$1 \rightarrow 2$	5,676,219	2,605,076	2.18
	$2 \rightarrow 3$	2,605,076	1,461,218	1.78
	$3 \rightarrow 4$	1,461,218	1,121,838	1.30
Height 5				
	$1 \rightarrow 2$	7,932,839	3,737,980	2.12
	$2 \rightarrow 3$	3,737,980	1,969,543	1.90
	$3 \rightarrow 4$	1,969,543	1,291,201	1.53
	$4 \rightarrow 5$	1,291,201	1,076,208	1.20
Height 6				
	$1 \rightarrow 2$	10,077,634	3,798,571	2.65
	$2 \rightarrow 3$	3,798,571	2,386,518	1.59
	$3 \rightarrow 4$	2,386,518	1,533,539	1.56
	$4 \rightarrow 5$	1,533,539	1,178,869	1.30
	$5 \rightarrow 6$	1,178,869	1,080,183	1.09

Notes: Firm-weighted median wages (CLP), heights 3–6. Ratio = Wage(Level N) / Wage(Level N+1). The CEO-to-Level 2 transition is the steepest at every height (2.12–2.65×), exceeding the Level 2-to-Level 3 ratio (1.53–1.90×). Lower-level transitions are generally smaller (1.09–1.56×).

A natural question is how much of total wage variation is attributable to differences *between* firms versus differences *within* firms. Table 7 reports a variance decomposition of log wages using firm fixed effects. Firm identity explains 31.4% of wage variance; the remaining 68.6% occurs within firms. Within firms, hierarchy variables explain a substantial share of this variation: adding level, height interactions, span, and reach to a baseline model with demographic and firm-level controls and firm fixed effects increases the within-firm R^2 by 30.6 percentage points (from 0.107 to 0.413; Table 6). Both statistics point to internal organizational structure as an economically important dimension of pay dispersion, but they capture different objects: the between-firm share reflects sorting across employers, while the within-firm hierarchy contribution reflects how positional variables structure pay within a given firm.

A related finding appears in Spanos (2022), who uses French linked employer-employee data to show that firm organization accounts for 31–50% of the density-related component

caution. These are estimated via linear probability models, where R^2 is not bounded between 0 and 1 in the usual sense, and low base rates – e.g., mean voluntary turnover of 6.4% per month – mechanically compress fit statistics relative to continuous dependent variables.

Table 6: **Regressions adding hierarchy variables: level, span, and reach.**

Dependent Variables:	log(Wage)	Voluntary Turnover	Involuntary Turnover	Absence	Absence Days†	Leave	Leave Days†
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
Firm Height (levels)	0.0188***	0.0097***	0.0033*	-0.0044***	0.0088	-0.0047***	-0.6796***
	(0.0064)	(0.0021)	(0.0017)	(0.0014)	(0.0584)	(0.0011)	(0.1828)
Employee Level	-0.2655***	0.0112***	0.0104***	0.0107***	-0.1680**	0.0063***	-1.092***
	(0.0106)	(0.0027)	(0.0022)	(0.0021)	(0.0672)	(0.0015)	(0.2553)
Level × Height	0.0141***	-0.0025***	-0.0014***	-2.57×10^{-5}	0.0109	0.0009***	0.1813***
	(0.0018)	(0.0005)	(0.0004)	(0.0004)	(0.0116)	(0.0003)	(0.0409)
Span of Control (direct reports)	-0.0074***	0.0001	0.0002**	-0.0006***	0.0099	-0.0002***	0.0051
	(0.0006)	(8.38×10^{-5})	(6.90×10^{-5})	(6.30×10^{-5})	(0.0179)	(5.05×10^{-5})	(0.0120)
log(Reach + 1)	0.2665***	-0.0068***	-0.0054***	-0.0027***	0.0050	-0.0068***	-0.7900***
	(0.0037)	(0.0008)	(0.0006)	(0.0004)	(0.0878)	(0.0004)	(0.1127)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Dep. Var. Statistics							
DV Median	14.210	0.000	0.000	0.000	1.000	0.000	11.000
DV Mean	14.333	0.064	0.037	0.036	1.936	0.052	14.248
DV SD	0.727	0.244	0.189	0.185	2.584	0.222	11.137
Fixed-effects							
Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics							
Observations	9,646,190	9,646,190	9,646,190	9,646,190	426,642	9,646,190	701,043
Within R^2	0.41259	0.00422	0.00314	0.00977	0.00141	0.00488	0.02110

Notes: Controls + Hierarchy Block. Controls (age, tenure, female, log firm size) collapsed. Firm-month weighted: each firm-month receives equal weight. FTE wage restriction applied to all columns. Standard errors clustered at the firm level in parentheses. †Conditional on the dependent variable being > 0. These columns keep the unconditional firm-month weights rather than renormalising within the conditional subsample, so a firm-month enters in proportion to the share of its workers with a positive outcome. Dependent-variable statistics are firm-month weighted, matching the estimation. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 7: **Variance decomposition of log wages.**

Component	Variance	Share
Between-firm	0.14	31.4%
Within-firm	0.31	68.6%
Total	0.46	100%

Variance of log winsorised monthly wages, split into a between-firm component (the employment-weighted dispersion of firm mean wages around the grand mean) and a within-firm residual. *Weighted by worker-month*, which is a deliberate exception to the firm-equal convention used for the other descriptive exhibits: this statistic describes the wage dispersion a worker faces, not the structure of the typical firm, so the population of interest is workers rather than firms. Weighting firms equally instead would raise the between-firm share to roughly 42%, because smaller firms are more dispersed and would carry the same weight as large ones. The between-firm share bounds what any within-firm variable can explain.

of within-firm wage disparities, suggesting that organizational structure is a quantitatively important source of within-firm pay variation across different empirical contexts. Spanos (2022) reports that firm-organization variables explain 31-50% of within-firm wage variance in cross-sectional specifications; panel firm-fixed-effects estimates yield 12-16%, suggesting that the organizational signal is partially absorbed by firm-FE and these magnitudes should not be interpreted as directly comparable.

A direct comparison with the AKM tradition is limited by a feature of our data: employee identifiers in the HR platform are firm-specific rather than portable across firms. This precludes tracking workers across employers, which is necessary for decomposing wages into worker and firm components in the AKM framework (Abowd et al., 1999; Card et al., 2013). Our between-firm share (31.4%) therefore captures both true firm wage premiums and compositional differences in worker quality across firms – it is an upper bound on the true firm-premium share, analogous to the between-establishment variance in a raw decomposition before worker fixed effects are absorbed. For comparison, Song et al. (2019) document that the between-firm component accounts for approximately two-thirds of the *rise* in U.S. earnings inequality over 1978–2013 (a change over time, not a level statistic); their full-population linked data permit a richer decomposition than our firm-specific identifiers allow.

The 30.6 percentage-point within-firm R^2 gain from hierarchy variables is measured conditional on firm fixed effects and thus absorbs all between-firm variation, but it cannot separate the contribution of organizational position from the sorting of higher-ability workers into higher positions within firms. This distinction is fundamental for interpretation: if the hierarchy-wage gradient is entirely explained by sorting – the most productive workers occupying the highest positions – then the "organizational structure explains wages" interpretation overstates the causal role of hierarchy. Separating these channels would require within-worker variation in hierarchical position (e.g., worker-mover designs exploiting transitions across positions), which the current data can support in limited form and which we identify as a priority for future work.

To assess the incremental contribution of organizational variables beyond hierarchical position, we decompose the 30.6 percentage-point within-R^2 gain sequentially (Table 8). Adding level alone to the baseline controls increases the within-R^2 by 21.6 percentage points (from 0.107 to 0.323), accounting for 70% of the total gain. This is expected: hierarchical level is a powerful positional proxy, analogous to occupation codes in other datasets. The Level $\times$ Height interaction contributes a modest additional 0.7 percentage points. However, span of control and organizational reach – variables with no occupational analog that capture the breadth and depth of a worker's position in the reporting tree – add a further 8.3

percentage points (from 0.330 to 0.413), representing 27% of the total gain. Decomposing further, reach alone accounts for 8.1 of these 8.3 percentage points ($\hat{\beta} = 0.241$, SE 0.003). Span of control, added alone, contributes 2.0 percentage points ($\hat{\beta} = 0.021$, SE 0.001), but this contribution is almost entirely shared with reach: once reach is included, span's marginal contribution is only 0.2 percentage points (the difference between the 8.3-point joint gain and reach's 8.1-point solo gain). Reach – the size of a worker's organizational subtree – is thus the primary channel through which reporting-link data add explanatory power beyond positional level. The dominance of reach over span is consistent with Rosen (1982), who predicts that the return to authority derives from the multiplicative effect of a manager's decisions on downstream workers: reach captures the economic rents from authority over a larger subordinate tree, whereas span measures only first-order supervisory load. These organizational architecture variables thus provide explanatory power for wages that cannot be replicated by positional measures alone, establishing that the reporting-link data contain information beyond what occupation-based hierarchy measures capture. Because our specification omits education and occupation, both correlated with hierarchical position, the incremental R^2 figures should be read as upper bounds on the within-firm wage variance attributable to hierarchical position; we do not interpret the magnitudes as causal contributions of hierarchy. Baker et al. (1994), in the canonical empirical-personnel-economics study of a single U.S. firm, find that level dummies alone explain 68% of cross-sectional log-wage variance, with a full battery of human-capital controls (sex, race, education, tenure) raising R^2 only to 0.71 – a 3-percentage-point increment. Read against this benchmark, the 21.6-percentage-point level contribution we report is unlikely to be primarily attributable to omitted education and tenure; the 8.3-percentage-point span-and-reach contribution is less susceptible to this concern in any case, since these variables capture organizational architecture rather than worker characteristics.

Table 8: **Sequential decomposition of within-firm wage variance.**

Specification	Within R^2	Gain (pp)
Controls only	0.107	–
+ Level	0.323	21.6
+ Level x Height	0.330	0.7
+ Span + Reach	0.413	8.3
Total		**30.6**

Within-R^2 from nested specifications, each adding a block to the previous one, with firm and period fixed effects throughout. Gains are in percentage points of within-R^2. The order is not innocuous: because level, span and reach are correlated, the gain attributed to each block depends on where it enters. Read the totals rather than the individual attributions.

Table 31 re-estimates the height–size slope, the level–wage gradient and the level–turnover gradient under all three conventions. The first two are robust to the choice. The height–size slope runs from 0.726 under firm-equal weights to 0.741 under firm-month weights and 0.780 under employee-month weights; the level–wage gradient from −0.254 to −0.265 to −0.237. All six estimates are precise and none changes sign.

The level–turnover gradient is different, and the difference is worth stating plainly: it is distinguishable from zero under *only one* of the three conventions. Under the firm-month weights used in the regression tables it is 0.011 and precisely estimated; weighting firms equally gives 0.005 with a standard error of 0.0035, and weighting worker-months equally gives 0.001 with a standard error of 0.0015. Neither alternative can reject zero. This is not a monotone pattern in firm size, so we do not read it as large firms drowning out a real effect. The honest reading is that the level–turnover association is small enough that the convention chosen for averaging determines whether it can be detected at all, and we treat it as a robustness object rather than a result, and none of the paper's claims rest on it.

The Level × Height interaction in the wage regression is positive and precisely estimated ($\hat{\beta} = 0.0141$, $p < 0.001$): the wage penalty associated with each additional level of depth is steeper in shorter firms. Computing marginal effects via the delta method, the implied wage gradient ranges from −0.223 log points per level in height-3 firms (SE 0.006, $p < 0.001$) to −0.181 log points in height-6 firms (SE 0.004, $p < 0.001$). All marginal effects are precisely estimated and economically large, and the attenuation in taller firms indicates that deeper hierarchies compress the per-layer wage differential. Full marginal effects by firm height are reported in Appendix B.3.[13]

4.2.2 Salary Volatility

Convex wages raise the question of how pay variability – the exposure to bonuses, commissions, and performance-linked components – differs across hierarchical levels. Table 9 reports intra-individual wage variability by level and firm height. We measure variability using the coefficient of variation (CV) of monthly salary within employee-years, which proxies for exposure to variable pay components.

The pattern is notable: CV is lowest at the top (Level 1: median CV 0.03–0.04) and increases toward middle-to-lower levels (Levels 4–6: median CV 0.10–0.14). The CEO-level CV estimate in the tallest firms rests on a few hundred unique individuals (334 unique workers in the sparsest cell) and should be read with corresponding care. This indicates greater

[13]Wild cluster bootstrap inference (999 Rademacher replications, 3,059 firm-level clusters) confirms the asymptotic results for the height–size regression ($p = 0.001$, the smallest value 999 replications can resolve; bootstrap 95% CI $[0.715, 0.769]$). See Appendix B.2.

Table 9: **Within-person salary variability by firm height and level.**

Firm height	Level	Mean CV	Median CV	SD CV	P25	P75	Mean salary	Person-years
3	1	0.112	0.042	0.154	0.012	0.152	6,695	1,244
	2	0.134	0.087	0.151	0.036	0.178	2,594	14,101
	3	0.139	0.105	0.137	0.057	0.177	1,549	17,349
4	1	0.111	0.040	0.145	0.012	0.164	9,260	1,476
	2	0.150	0.090	0.170	0.029	0.220	4,305	9,671
	3	0.148	0.100	0.168	0.048	0.193	2,104	31,758
	4	0.141	0.103	0.147	0.057	0.179	1,518	29,253
5	1	0.098	0.036	0.136	0.012	0.147	11,557	1,075
	2	0.165	0.100	0.179	0.027	0.256	5,885	6,764
	3	0.154	0.105	0.169	0.046	0.206	2,724	23,777
	4	0.159	0.119	0.166	0.067	0.195	1,710	53,514
	5	0.158	0.117	0.174	0.071	0.185	1,506	32,209
6	1	0.084	0.029	0.121	0.011	0.104	13,894	527
	2	0.179	0.107	0.192	0.025	0.286	7,631	3,583
	3	0.182	0.111	0.201	0.042	0.268	3,772	14,020
	4	0.171	0.121	0.180	0.060	0.223	2,177	38,012
	5	0.181	0.132	0.178	0.080	0.225	1,589	51,522
	6	0.188	0.144	0.167	0.088	0.237	1,528	19,483

Coefficient of variation of *monthly* salary within an employee-calendar-year, requiring at least 9 months of observed pay in the year; a worker's level and firm height are the modal values within that year. Cells average over person-years, each weighted equally — not the firm-month weighting used in the regression tables. Mean salary is in thousands of 2025 CLP. Because the underlying series is monthly payroll, it excludes the annual bonuses and deferred compensation through which executive incentive pay is typically delivered, so this table measures month-to-month take-home risk and not total compensation risk (Section 4.2).

salary variability further from the executive suite – a pattern that may seem counterintuitive given standard intuitions about executive bonuses. However, it is consistent with lower-level workers facing more variable hours, piece-rate components, or performance fluctuations, while senior executives receive more stable base salaries with bonuses that may be smoothed or deferred across periods. We also note that the low CV at the top could partly reflect the timing of discretionary compensation: if senior managers receive annual bonuses or deferred compensation outside the monthly payroll cycle, their monthly salary would appear more stable than their total annual compensation.

These patterns suggest that compensation variability reflects the distinct nature of work at different levels: senior roles feature high base salaries with potentially deferred incentives, while lower-level roles expose workers to more immediate pay fluctuations tied to hours, commissions, or short-cycle performance metrics. An additional institutional mechanism reinforces this gradient: under Chile's Article 22 of the *Código del Trabajo*, most managerial workers are exempt from the standard working-hours regime and are not entitled to overtime pay, whereas non-exempt workers at lower levels receive legally mandated overtime compensation. Article 22 exempt status is not directly observed in our data; we infer its relevance from the institutional fact that Chilean labor law exempts workers in senior management and supervisory positions from standard working-hour regulations. The observed volatility gradient may therefore partly reflect hours variation interacting with the overtime regime rather than incentive design alone. This gradient is not an artifact of contract-type heterogeneity. Permanent contracts (*contrato indefinido*) account for 84% of the wage sample, and fixed-term arrangements are more common at the bottom of the hierarchy, so contract mix is a natural candidate explanation for a gradient that rises with depth. Recomputing the coefficient of variation on permanent contracts only, within the same height × level cells as Table 9, leaves the pattern intact: median CV runs from 0.03–0.04 at level one to 0.10–0.14 at levels four and below, and no cell moves by more than 0.002 against the full sample. Because the comparison is cell by cell rather than pooled across heights, it is directly readable against the table.

4.3 Compression: Pay, Authority, and Span Across Firm Heights

Three margins of what a rank confers – how much it pays, whether it carries supervisory authority at all, and how much authority it carries when it does – compress as firms grow taller. The question is one that knowledge-based and assignment theories both raise once firms differ in depth (Garicano, 2000; Rosen, 1982): when a firm adds layers, does each layer remain as distinct from its neighbors as in a shallow firm, or do layers become more similar to one another as they multiply? Answering it on a single sample requires observing the reporting link, not only an assignment of workers to layers.

We measure each margin as a gradient in employee level, allowed to vary with firm height. The first is the wage step: the change in log wages per additional level of depth, recovered by the delta method from the Level × Height specification of Section 4.2 and reported in full in Appendix B.3. The second is the extensive margin of authority: the per-level change in the probability of supervising at least one direct report. The third is the intensive margin: the per-level change in log(span) among workers who supervise anyone at all. The two authority regressions include *firm-month* fixed effects, comparing workers at different levels of the same firm in the same month and absorbing firm-level composition and any common time-varying shock; the wage margin comes from the firm- and period-fixed-effects specification already estimated above, so identification is not identical across the three panels. All three use the firm-month-equal weighting convention applied throughout and cluster standard errors by firm. The three gradients are collected in Table 10.

The wage step attenuates modestly. It moves from -0.223 log points per level in height-3

Table 10: **Three margins of hierarchical compression.**

	Wage step		P(supervises)		log(span) | supervises	
Firm height	Est.	(SE)	Est.	(SE)	Est.	(SE)
3	-0.223	(0.006)	-0.476	(0.004)	-0.518	(0.055)
4	-0.209	(0.005)	-0.320	(0.003)	-0.161	(0.019)
5	-0.195	(0.004)	-0.240	(0.003)	-0.024	(0.013)
6	-0.181	(0.004)	-0.184	(0.003)	0.055	(0.012)

Per-level gradients by firm height. The wage step is a delta-method marginal effect from the Level × Height interaction on the full wage specification; the two authority margins are estimated within height, with firm-month fixed effects. All three are firm-month weighted with standard errors clustered on firm. The intensive margin conditions on supervising at least one worker.

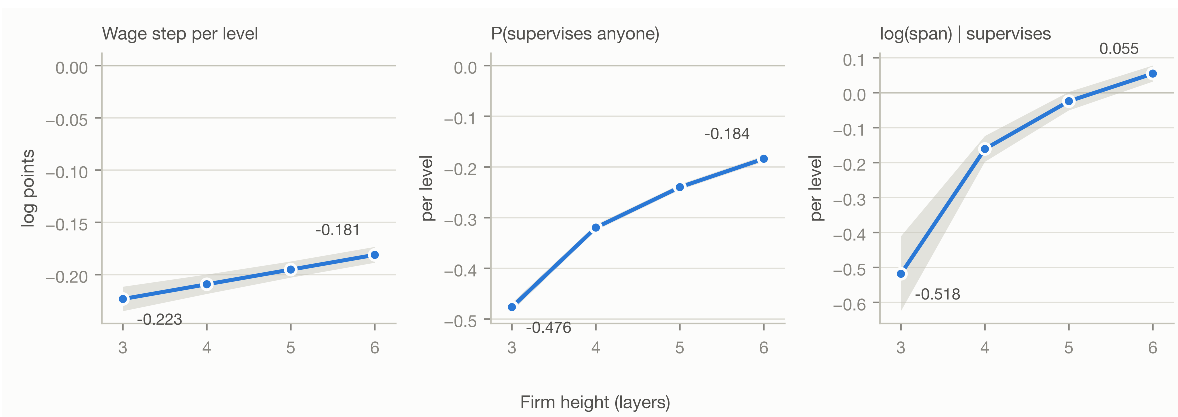

Figure 8: **Three margins of hierarchical compression.** Each panel plots a per-level gradient against firm height with 95% confidence bands: the wage step (delta-method marginal effects from the Level × Height wage specification), the gradient in the probability of having at least one direct report, and the gradient in log(span) conditional on supervising. The three are not estimated the same way. The wage panel comes from a single pooled model with a linear Level × Height interaction, so its path is linear in height by construction; the two authority panels are estimated separately within each height, so their monotonicity is an estimate rather than a functional form. The two authority regressions include firm-month fixed effects. Standard errors clustered by firm; firm-month-equal weighting throughout. Heights 3–6 are the range on which all three margins are estimated; the same specification implies -0.237 at height 2 and -0.167 at height 7, heights held back from the figure because both authority margins are degenerate at height 2 — supervisors exist only at level one, so there is no within-firm-month variation in rank among them — and height 7 has too few firms.

firms to -0.181 in height-6 firms, passing through -0.209 and -0.195 at heights 4 and 5 – a smooth monotone path and a proportional attenuation of about 19%. A rung of a tall ladder pays less than a rung of a short one, but it remains a large wage step at every height.

The extensive margin of authority attenuates much more, and part of that movement is arithmetic. The gradient runs from -0.476 in height-3 firms to -0.184 in height-6 firms, through -0.320 and -0.240; these are linear-probability coefficients, so the height-3 estimate is a decline of 47.6 percentage points per level. Because the probability of supervising falls from near one at the apex to near zero at the base, the average per-level slope in a firm of height H is bounded by $-1/(H-1)$ – from -0.500 at height 3 to -0.200 at height 6 – and our estimates lie between 92% and 96% of that bound at all four heights. The extensive margin therefore compresses at close to the rate the range restriction alone requires. We report it for completeness and place the economic content of the subsection on the other two margins, which face no analogous bound.

The intensive margin compresses furthest and crosses zero. Conditioning on supervising at least one worker, the level gradient in log(span) runs from -0.518 at height 3 to 0.055 at height 6, with -0.161 and -0.024 in between. It does not merely reach zero: the height-6 estimate is positive and precisely estimated (SE 0.012), so among supervisors in the tallest firms a manager one rung deeper supervises *more* people, not fewer. The reading is not that team size stops mattering but that its relation to rank inverts: Figure 4 shows spans in height-6 firms widening into the middle of the hierarchy and easing only below it, so the linear slope over the whole ladder comes out positive. This complements Figure 5b, where the shift to predominantly non-managerial roles occurs at progressively lower ranks as firms deepen.

These are three distinct objects – a price, a probability, and a quantity – and two of them are estimated within firm-month cells. Two limits should be kept in view. Variation in height is across firms: firm-month fixed effects absorb the level of firm characteristics but not their interaction with height, and because height rises with size (Fact 1), compression with depth cannot be separated from compression with scale or its correlates. And within a firm-month spans must sum to $N-1$, so the level–span profile is disciplined by the headcount pyramid of Figure 5a. What the estimates do establish is that "one level" is not a fixed unit of career progression: what a step on the ladder delivers, in pay and in authority, depends systematically on the shape of the organization in which the ladder is embedded (Baker et al., 1994).

Thinner rungs, crossed more often. The three margins say that a rung buys less in a tall firm. Internal mobility says that workers cross rungs more often there. The rate of individual team switching – a change of direct manager that is not explained by the worker's former colleagues moving with them (Section D.2) – rises monotonically with height, from 0.39% per month in height-3 firms to 0.82%, 1.08%, and 1.28% at heights 4, 5, and 6: a factor of 3.3 across the range over which the wage step falls by 19% and the span gradient crosses zero. Depth therefore does not simply stack more copies of the same rung. It produces a finer-grained ladder: more steps, each worth less in pay and authority, and traversed more frequently. That combination is what an internal labor market with low-stakes, high-frequency reallocation looks like, and it is the opposite of the picture in which added layers are added barriers.

We put this forward as an interpretation, not as a test. The two facts are measured on different objects – the gradients on the cross-section of levels within firm-months, the switching rate on within-firm transitions between consecutive months – and a team switch is not a promotion: our measure records that the supervisory relationship changed, not that the worker moved up. Separating lateral reallocation from advancement requires tracking level transitions for the same worker, which the panel supports and we leave to companion work. What the two facts jointly rule out is the reading in which tall firms are simply short firms with more layers bolted on. On both margins we can measure, the rungs of a tall firm behave differently: they are cheaper, and they are crossed more often.

What the average attenuation hides. The per-level averages above are taken over each firm's full ladder, and the ladders being compared are not the same length. Re-estimating the identical specification on the top three levels — two rungs, the only stretch all four heights share — the step runs from -0.298 in height-three firms to -0.205 in height-six firms, an attenuation of 31%. The same specification on the same firms over their full ladders gives 22%, so the level window alone moves the estimate by nine points and in the direction opposite to the one the framing above would suggest. Compression is not carried by the rungs that only tall firms have: it is present, and stronger, on the stretch every firm shares.

The linear index imposes one step per firm, which is the assumption worth relaxing. Estimating levels as a factor instead, and controlling for span and reach throughout, both ends of the ladder compress but they do so differently. The step at the *base* falls from -0.271 to -0.096 log points, losing 65% of its magnitude while staying negative. The level-one-to-level-two step at the apex goes from -0.223 to $+0.123$: it does not merely shrink, it changes sign, so once reach is held fixed the level-two worker in a tall firm is paid *more* than the person above them. That is what we mean by saying the apex premium of tall firms is a premium for reach rather than for rank; without the reach controls the same step is -0.756 and -0.511, steeply negative at both heights. The total apex-to-base gap meanwhile widens with height, from -1.207 to -1.867 in the uncontrolled specification: tall firms are not flatter overall, they are flatter *between neighbours*. Deeper firms therefore compress their ladder at both ends while stretching it end to end: neighbouring rungs grow more alike, and the distance from top to bottom grows.

Are the levels real? Collapsing unitary chains. Measured depth might be inflated by titling: 20.8% of managers have exactly one direct report, and such pass-through nodes add a rung without adding employment.[14] We rebuild each firm-month tree from the recorded supervisor identifier, contract every maximal chain of single-report nodes into one edge, and recompute all primitives on the collapsed tree. Mean height falls only from 4.17 to 3.92 layers, so pass-through nodes are common but rarely stacked.[15] The gradients survive: the extensive margin runs from -0.438 to -0.172 and the intensive margin from -0.373 to 0.069, again crossing zero at height 6. The height–size semi-elasticity *rises* from 0.741 to 0.765 under collapsing.[16] Titling therefore attenuates the scaling relationship rather than generating it: removing the layers most likely to be cosmetic makes depth scale more steeply with employment, not less.

What we do not claim. A balanced-tree reading of Fact 1, in which the height–size slope recovers Garicano's $1/\ln(s)$ for a constant span s, is tempting, and we evaluated it. It does not survive. The benchmark $h = 1 + \ln N/\ln s$ restricts the intercept as well as the slope, and the data reject the restriction: the fitted intercept, which the balanced tree pins at 1.000 because a one-worker firm has height one, is 1.166 (SE 0.052). The implied s also ranges from about 2.5 to 3.9 depending on how height is defined. We record the mapping as an interpretation considered and rejected, not as a structural estimate, and none of the gradients above depends on it. They are descriptive targets that a model of layer formation should reproduce, and they leave open whether the pattern reflects knowledge-based assignment (Caliendo et al., 2015; Garicano, 2000), the assignment of talent to positions of amplified control (Rosen, 1982), or job-ladder design of the kind documented by Baker et al. (1994).

[14] Firm-month-equal share on the subsample where the tree can be rebuilt from the recorded supervisor identifier. Section 4.1.2 reports it over the full sample: 21.4% weighting firms equally, 15.5% weighting managers equally.

[15] Firm-month-equal means on the reconstruction subsample; the firm-based mean used elsewhere in the paper is 4.1 (Table 1).

[16] Both re-estimated on the reconstruction subsample; the headline full-sample estimate is 0.741 (SE 0.013).

4.4 Industry Heterogeneity

How does the relationship between firm size and hierarchical complexity vary across industries? Knowledge-based theories of hierarchy predict that industries requiring greater problem-solving, coordination, and information processing should exhibit deeper hierarchies for a given firm size (Garicano, 2000; Garicano & Rossi-Hansberg, 2006). We examine this by estimating the height–size semi-elasticity – the coefficient from regressing firm height on log firm size – separately for each industry with sufficient sample size.

Table 11 reports these semi-elasticities for the seventeen industries with at least 20 firms and 300 firm-months. The pooled estimate across all industries is 0.741, implying that each doubling of firm size adds approximately 0.51 layers. Sectors do differ, but the differences are estimated with very unequal precision, and the ordering is informative only over the sectors that are well populated: of the 136 sector pairs, only 11 of the 55 pairs among them separate at the 5% level. We therefore read this evidence as heterogeneity with limits rather than as a ranking.

Table 11: **The height–size semi-elasticity across industries.**

Sector	Semi-elast.	95% CI	Firms
Agriculture	0.90	[0.79, 1.01]	103
Manufacturing	0.84	[0.76, 0.92]	275
Transport	0.81	[0.70, 0.91]	165
Finance	0.80	[0.73, 0.86]	159
Commerce	0.78	[0.73, 0.84]	571
Professional services	0.77	[0.68, 0.87]	313
Real estate	0.77	[0.51, 1.03]	58
Construction	0.76	[0.66, 0.86]	198
Health	0.68	[0.58, 0.78]	103
Hospitality	0.68	[0.57, 0.78]	141
Information & comms.	0.67	[0.55, 0.79]	298
Admin. services	0.65	[0.55, 0.75]	242
Arts & recreation	0.64	[0.44, 0.83]	36
Water & waste	0.63	[0.38, 0.87]	30
Other services	0.61	[0.42, 0.80]	112
Utilities	0.61	[0.30, 0.91]	23
Mining	0.59	[0.06, 1.12]	25

OLS of firm height on log employment, estimated separately by sector on firm-months with at least five employees; standard errors clustered on firm. Sectors in grey are estimated too imprecisely to be ranked: their confidence intervals overlap most of the table. Only 15 of 136 sector pairs separate at the 5% level, so the extremes of the ordering should not be read as distinguishable.

Among the 11 sectors estimated with precision, agriculture, livestock, forestry and fishing shows the steepest scaling (0.90) – each doubling of firm size adds 0.62 layers, about a fifth more than the pooled average – and administrative and support services the flattest (0.65, or 0.45 layers). That difference is itself precisely estimated ($t = 3.3$). Manufacturing (0.84), transportation and storage (0.81) and financial services (0.80) also scale above the pooled average. At the flatter end sit administrative and support services (0.65) and, more surprisingly, information and communications (0.67) – a knowledge-intensive sector estimated on close to three hundred firms, so its position is not a small-sample artefact. Professional and technical services (0.77), also knowledge-intensive, sits above the pooled estimate, though its confidence interval contains it. We quote only sectors estimated precisely enough to place; the remaining cells appear in grey in Table 11 and are not ranked here. The nominal minimum of the table, mining and quarrying, illustrates why: it rests on 25 firms and its confidence interval spans $[0.06, 1.12]$, wide enough to contain the maximum. The ordering therefore does not reduce to the knowledge-intensity gradient that motivated the exercise: the steepest scaling appears in primary and industrial sectors, and several knowledge-intensive service industries add depth no faster than the pooled average. One reading is that hierarchical depth responds to coordination demands that are not specific to knowledge work – geographically dispersed operations, multi-site production, and supervision-intensive manual processes may require added layers just as surely as layers of expert review do.

Knowledge-based theories link industry coordination requirements to organizational structure (Garicano, 2000; Garicano & Rossi-Hansberg, 2006), and a complementary interpretation comes from the management practices literature. Bloom and Van Reenen (2007) and the subsequent World Management Survey research consistently find that management quality varies systematically across industries, with manufacturing among the higher-scoring sectors on structured management practices. The cross-industry variation in management quality overlaps only partially with the cross-industry variation in hierarchical depth documented here, raising the question of whether depth differences reflect information-processing requirements, management quality, or industry-specific production technologies – and, by implication, whether some of the variation reflects better versus worse organizational design rather than purely efficient adaptation to technology.

Additionally, Guadalupe and Wulf (2010) (building on Rajan & Wulf, 2006) provide causal evidence that product market competition flattens hierarchies: using the Canada-U.S. Free Trade Agreement as a quasi-natural experiment, they show that increased competition reduced hierarchy depth by 11% and increased CEO span by 6%. The industry heterogeneity we document could partly reflect differences in competitive intensity: service industries in Chile face varying competitive pressures, and the Guadalupe-Wulf mechanism predicts flatter hierarchies in more competitive product markets. Without firm-level measures of competitive conditions, we cannot distinguish the knowledge-intensity channel from the competition channel, but both deserve consideration when interpreting the industry variation. We also note that compositional controls – for firm size distributions, workforce composition, and competitive conditions within industries – would strengthen the industry results and help distinguish genuine structural differences from confounds.

Several industries have limited cluster counts (e.g., utilities with 23 firms, mining with 25), and the corresponding semi-elasticity estimates carry wider confidence intervals. These industry-level results should be treated as suggestive rather than inferentially precise for the smallest-cluster industries. They highlight that the "typical" height–size relationship documented earlier masks economically meaningful variation across industries.

Worker outcomes: salary volatility and mobility. The salary volatility gradient documented in Table 9 is among the paper's most behaviorally relevant findings. The pattern – lowest CV at the CEO level (0.03–0.04), peaking at middle-to-lower levels (0.10–0.14) – has direct implications for compensation design and worker welfare. It suggests that compensation risk in the firm is borne disproportionately by lower-level workers, who face more variable hours, piece-rate components, and performance fluctuations, while senior executives receive more stable base salaries. We stress what this does and does not establish. Monthly payroll cannot adjudicate rank-order tournament models: the annual bonuses, deferred pay, and equity that carry executive risk sit outside the monthly cycle we observe, so a flat executive CV is as consistent with smoothing as with low risk. The defensible statement is narrower and, for worker welfare, more direct: *month-to-month* pay risk – the variation that must be absorbed by household budgets without savings buffers – falls disproportionately on workers at the bottom of the hierarchy, through variable hours, commissions, and short-cycle performance components. Chile's Article 22, which exempts most managerial workers from the overtime regime, is one institutional channel that would produce exactly this gradient.

Internal labor market mobility – measured as team switches excluding mass reassignments – shows a similarly clear pattern: monthly true mobility rates range from 0.39% in height-3 firms to 1.28% in height-6 firms, with moves concentrated at lower levels, suggesting that deeper hierarchies feature more active reallocation at lower layers (Appendix D, Tables 32–33). The combination of higher mobility and higher salary volatility at lower levels paints a consistent picture of the organizational periphery as a more dynamic, less stable employment environment.

Additional exploratory evidence. Turnover gradients – the tendency for exit rates to rise with distance from the top in shorter firms, and to fall with it in the tallest – are robust in the full sample: the level–turnover association is positive and statistically significant across nearly all specification variants (Table 12), with two exceptions: the restriction to firms of at least 50 employees, where the coefficient is 0.001 and indistinguishable from zero, and the anchor-rule-only variant, where it is -0.002 and changes sign. Gender representation exhibits a steep vertical gradient: women comprise 39.6% of the workforce but only 16.8% of top-level positions[17], with glass ceiling ratios ranging from 0.32 (height-6 firms) to 0.59 (height-3 firms); the within-firm gender wage gap is -0.149 log points (SE 0.005), compared to a raw gap of approximately -0.188 log points, representing approximately 79% of the raw gap (Appendix D, Figure 10, Tables 13–14). Section 4.5 decomposes the raw gap into the rung women occupy, the firm they work in, and pay within rung and firm. What that decomposition does *not* identify is why the within-rung component exists: separating discrimination from unobserved differences in tasks, experience or hours would require controls this data does not carry. Finally, manager demographics and team composition show modest conditional correlations with subordinate outcomes once the worker's own hierarchical position is controlled (Appendix D, Tables 35–36).

The main coefficients are also robust to the sample-selection rule itself. Appendix B.1 reports the three coefficients across hierarchy-resolution thresholds from 80% to 100%, within-sample dose–response bands, a panel-length grid, and a forest of estimates across industry, firm-size, region, sales-bracket, and time subgroups; the height–size and level–wage gradients keep their sign and rough magnitude throughout, and the level–turnover coefficient, while small in magnitude, is statistically significant at every resolution threshold.

4.5 Gender: Representation, Pay, and Where the Gap Lives

Women make up 39.6% of worker-months in the overall sample, but representation at the CEO level (Level 1) is far lower – 16.8% of CEO-months pooled across firms – and declines with organizational depth, from 23.3% in height-3 firms to 11.6% in height-6 firms (Fig-

[17] The two shares use different denominators: the 39.6% workforce share is computed over all employee-months in our analysis sample, while the 16.8% top-level share is computed over the subset of CEO-level (level 1) positions only. Both are reported as employee-month population shares (each worker-month contributes equally), which is the natural reading of the descriptive question and the convention in the glass-ceiling literature (Bertrand & Hallock, 2001; Cotter et al., 2001). Under firm-equal weighting (each firm contributing equally regardless of size or of how many months it is observed), the corresponding shares are 38.1% and 24.6%; the qualitative pattern of steep vertical under-representation is unchanged, though the CEO-level share is markedly higher when each firm counts once. Panel length accounts for about half of that difference: restricting the firm-equal calculation to the 2,221 firms observed for at least twelve months lowers the CEO-level share to 20.4%, a movement of 4.1 percentage points against the 7.8-point gap between the two conventions, while the overall workforce share barely moves (38.1% to 38.4%). Short panels are therefore part of the story and firm size is the rest: small firms are both more likely to have a woman at the top and given more weight when each firm counts once. Briefly observed firms are thus disproportionately woman-led, and firm-equal weighting gives each of them the same weight as a firm observed for the full window. Together these statistics document the steepness of the vertical gradient; the workforce and CEO-level shares are not directly comparable as raw shares.

Table 12: **Robustness of the main coefficients across specification variants.**

Variant	Height–Size β (log size)	Level–Wage β_{level} (log wage)	Level–Turnover β_{level} (vol. turnover)
(1) Baseline	0.741***	−0.265***	0.011***
	(0.013)	(0.011)	(0.003)
(2) Incl. anomalous teams	0.724***	−0.265***	0.011***
	(0.014)	(0.010)	(0.003)
(3) Excl. height 7+	0.628***	−0.271***	0.017***
	(0.012)	(0.014)	(0.004)
(4) Anchor rule only	0.653***	−0.192***	−0.002
	(0.036)	(0.013)	(0.003)
(5a) 100% resolution	0.691***	−0.277***	0.016***
	(0.017)	(0.014)	(0.004)
(5b) 90% resolution	0.747***	−0.256***	0.011***
	(0.013)	(0.010)	(0.002)
(6) Firms ≥24 months	0.748***	−0.279***	0.014***
	(0.018)	(0.014)	(0.003)
(7) Firms ≥50 employees	0.792***	−0.279***	0.001
	(0.028)	(0.014)	(0.003)
(8) Base salary only	0.741***	−0.310***	0.011***
	(0.013)	(0.012)	(0.003)

Notes: Each cell reports the key coefficient and standard error for one coefficient under one specification variant. Height–Size: OLS of firm height on log firm size (firm-month level, unweighted). Level–Wage: coefficient on employee level from the hierarchy wage regression (firm + period fixed effects, firm-month weighted, controls for age, tenure, gender and log firm size, plus the hierarchy block of level × height, span of control and log subordinate reach). Level–Turnover: coefficient on employee level from the same specification with voluntary turnover as the dependent variable. Standard errors for all three coefficients are clustered at the firm level. The level–turnover coefficient is reported here as a robustness object rather than as one of the three facts: it is distinguishable from zero under only one of the three weighting conventions, as discussed in the text. Every variant is rebuilt from the raw panel under its own sample rule rather than filtered out of the production frame, so that rules looser than the production one are actually looser; winsorisation cut-offs are held fixed at the production values throughout, so the variants move the sample and nothing else. Variant (2) includes firm-months containing anomalous teams. Variant (4) identifies the apex by the external-manager anchoring rule alone, dropping the subordinate-reach criterion. Variants (5a) and (5b) re-estimate at the 100% and 90% hierarchy-resolution thresholds in place of the 95% production threshold. Variants (6) and (7) are subsamples of the baseline restricted to firms observed for at least 24 firm-months and to firms whose median headcount is at least 50. Variant (8) replaces gross pay with base pay as the wage measure; Facts 1 and 3 are re-estimated on the rows that survive, not carried over from the baseline. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ure 10, Table 13). Glass ceiling ratios – the share of women at the top relative to the share at the bottom – are low throughout and fall with height, from 0.59 in height-3 firms to 0.32 in height-6 firms (Table 13). The pattern resonates with the "glass ceiling" literature: Bertrand and Hallock (2001) show that women held fewer than 2.5% of top executive positions in large U.S. corporations in the 1990s and that much of the raw gender pay gap among executives was attributable to women managing smaller firms and being severely under-represented in the top occupational titles (CEO, Chair, Vice-Chair, President). More broadly, the organizational sociology literature documents how invisible barriers disproportionately impede women's advancement into executive positions even in organizations with ostensibly equal-opportunity policies (Cotter et al., 2001). Strictly, the static ratios reported here satisfy only Cotter et al. (2001)'s hierarchy criterion (greater inequality at higher levels); the full four-criterion test of a glass ceiling effect would also require longitudinal evidence on advancement chances and career trajectories, which we leave to companion work focused on gender-specific dynamics.

Figure 9 shows how gender pay gaps vary along the ladder, comparing mean log wages by gender at each level. Gender wage gaps are generally larger at the top of hierarchies than in mid-levels. Controlling for firm fixed effects and hierarchical level, the within-firm gender gap is -0.149 log points (SE 0.005), representing approximately 79% of the raw gap (-0.188). The remaining 21% is sorting, and the reporting tree resolves it into its two parts: 12% of the raw gap (-0.188 log points) is the rung a woman occupies and 9% is the firm she works in, leaving 79% as pay within the same rung of the same firm. Sorting across rungs is therefore the larger of the two sorting channels, though both together are dominated by what women are paid on the rung they already hold. How the level control enters is not a detail. Entering level as a linear index gives a within-firm gap of -0.167 (SE 0.005); entering it as a full set of level dummies gives -0.149 (SE 0.005). The linear version overstates the gap by 0.017 log points, about 12%, because a single slope cannot absorb a step profile and leaves part of the level composition in the female coefficient. We report the dummy specification throughout. Both estimates weight each firm-month equally, the convention used in every regression table; weighting employee-months equally instead gives a within-firm gap of -0.133 (SE 0.005), the gap faced by the average worker rather than the gap inside the average firm. Firm-clustered bootstrap 95% confidence intervals (200 replications; Appendix A.2.5) place the CEO-level gaps above zero in every height, but not precisely: because firms are the resampled unit and few women hold CEO-level positions in any one firm, these cells rest on the smallest effective samples in the table, and their intervals are correspondingly wide – spanning $[27.0, 42.0]$ in height-4 firms and $[5.9, 51.4]$ in height-6. The *ordering* of gaps across levels is far better identified than any single cell's magnitude.

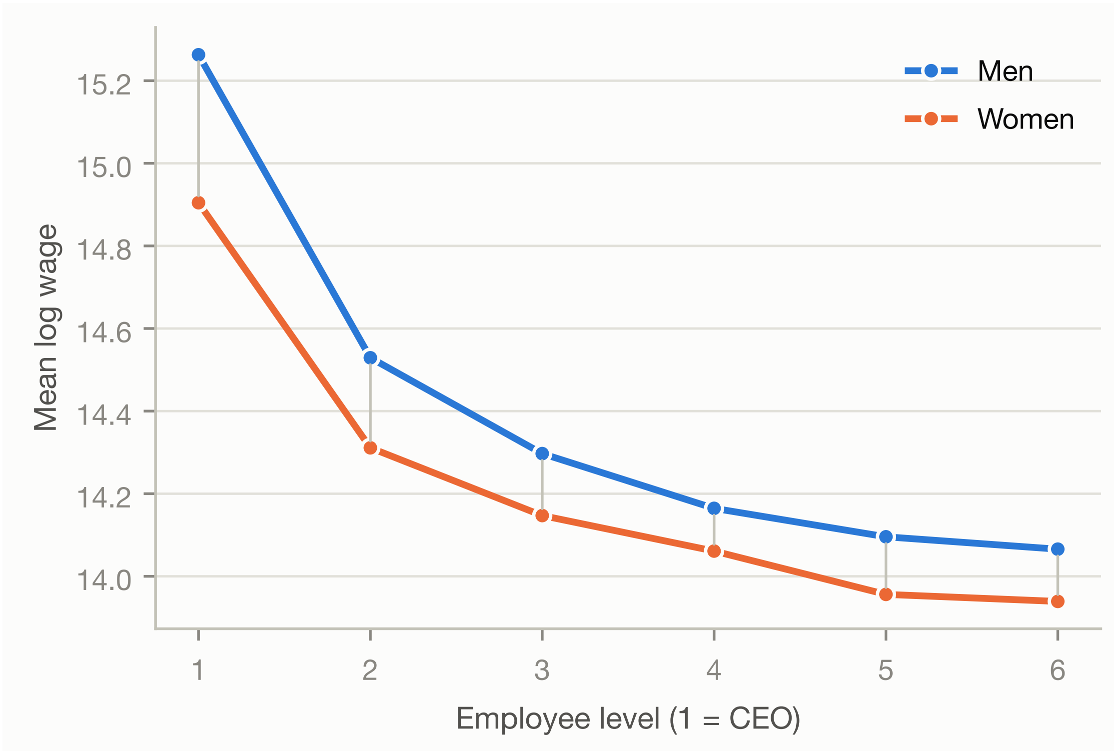

Figure 9: **Gender wage gaps by employee level.** Mean log wages by gender at each employee level, pooled across firm heights and weighted so that each firm counts once. The vertical distance between the two series is the raw gap at that level. The gap is widest at the apex (−0.358 log points at Level 1) and is narrower at the base (−0.126 at Level 6), though not monotonically in between; it is negative at every level, so no rung reverses. Table 14 reports the corresponding gaps cell by cell within firm height, with bootstrap confidence intervals.

Table 13: **Female representation at the top and bottom of the hierarchy, by firm height.**

Firm height	N (top)	N (bottom)	% female (top)	% female (bottom)	Ratio
Height 3	18,249	340,567	23.3%	39.8%	0.59
Height 4	22,705	597,987	16.2%	38.0%	0.43
Height 5	16,333	607,449	12.5%	36.8%	0.34
Height 6	8,549	384,309	11.6%	36.6%	0.32

The ratio is the female share of level-one worker-months divided by the female share at the firm's lowest level. Shares are firm-equal weighted, the convention used for descriptive exhibits throughout: each firm carries the same total weight regardless of its size or how many months it is observed. N columns are raw worker-month counts.

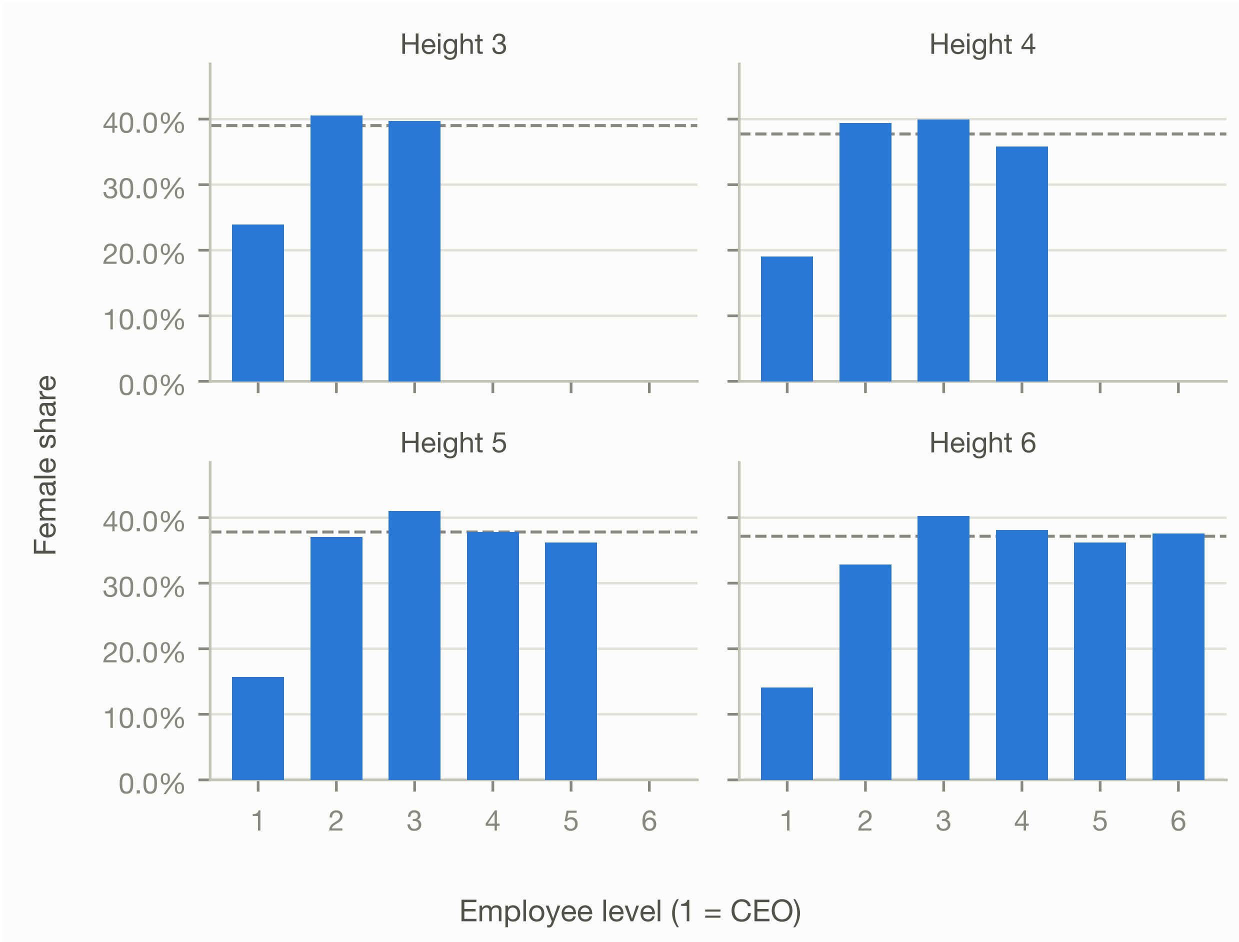

Figure 10: **Female representation by firm height and employee level.** The figure plots the percentage of female employees at each level (bars) within firms of different heights. Dashed lines indicate the overall female share for that specific firm height group. While the aggregate female share in the sample is 39.6%, representation at the CEO level (Level 1) is far lower and declines with the total number of organizational layers: from 23.3% in height-3 firms to 11.6% in height-6 firms — Table 13 reports the same quantity on the same firm-equal basis. Pooled across all firms, women hold 16.8% of CEO-months. Counting each employee-month instead of each firm gives a steeper decline; the firms that firm-equal weighting brings forward are disproportionately briefly observed, and those are more often woman-led (see the footnote cited below). The glass-ceiling literature conventionally reports employee-month shares (Bertrand & Hallock, 2001; Cotter et al., 2001); firm-equal-weighted overall and CEO counterparts (38.1% and 24.6%) are reported in the footnote to the workforce-share statement in Section 4 (paragraph on additional exploratory evidence). Heights 3–6 shown; see Section 3 for height range rationale.

5 Conclusion

This paper uses administrative HR data from 3,059 Chilean firms and 10.1 million employee-months to document new facts about organizational hierarchies. Because the data directly observe reporting relationships, we are able to measure hierarchy itself rather than infer it from occupations or titles.

The first set of facts concerns hierarchical structure. Firm height ranges from 2 to 9 layers, with substantial variation even among firms of similar size: each doubling of employment is associated with only 0.51 additional layers on average. Spans of control are highly skewed – roughly one in five managers supervises a single direct report, while some oversee hundreds – and the allocation of supervisory responsibility shifts systematically toward lower levels as firms deepen. These patterns produce a diamond shape, in which the CEO level accounts for less than 1% of headcount in tall firms.

The second set of facts concerns compensation. Pay is sharply convex in hierarchical position, with the steepest proportional increases concentrated at the top transitions of the hierarchy. CEO-to-median wage ratios rise from approximately $3\times$ in height-3 firms to $6\times$ in height-6 firms. Hierarchy variables substantially increase explained within-firm wage variance (Section 4.2), underscoring the importance of internal organizational structure for pay. Salary volatility is lowest at the top and highest in middle-to-lower levels, locating month-to-month pay risk at the bottom of the hierarchy. Whether *total* compensation risk follows the same gradient cannot be settled with monthly payroll, which does not observe annual bonuses or deferred pay.

The third set of facts concerns compression: as firms grow taller, adjacent levels become more alike in what they pay and in the authority they carry. The wage step between neighbouring rungs falls from -0.223 to -0.181 log points between height-3 and height-6 firms, and the level gradient in span among those who supervise falls from -0.518 to 0.055, crossing zero – not because team size stops depending on rank, but because in the tallest firms the direction reverses and deeper supervisors run larger teams. The averaging hides where the compression sits. Estimating levels as a factor, the step at the base of the ladder loses 65% of its magnitude between heights three and six, while the step at the apex, on the same specification, changes sign (-0.223 to $+0.123$). Deep firms compress the distance between neighbours at both ends of the ladder while widening the distance from top to bottom.

The reporting tree also changes what can be said about the gender wage gap. Splitting a raw gap of -0.188 log points along the ladder leaves 79% of it inside rungs — pay differences between men and women on the same rung of the same firm — with 12% attributable to which rung women occupy and 9% to which firm they work in. Under-representation at the top is real and steepest where the ladder is longest: the glass-ceiling ratio falls from 0.59 in height-3 firms to 0.32 in height-6. But the shortfall is numerically small, because the levels where it is most extreme hold few workers, and the weight of the gap sits inside rungs rather than across them. What the decomposition does not identify is why the within-rung component exists; separating discrimination from unobserved differences in tasks, experience or hours would require controls this data does not carry.

Industry heterogeneity sits alongside these facts rather than as a fourth one. The height–size semi-elasticity varies across industries, from 0.90 in agriculture to 0.65 in administrative services among precisely estimated sectors (0.62 vs. 0.45 layers per doubling); the noisiest sectors cannot be ranked. The ordering does not reduce to a knowledge-intensity gradient: the steepest scaling appears in primary and industrial sectors, which is not what a knowledge-intensity ordering would predict. One reading is that depth responds to coordination demands – dispersed operations, multi-site production, supervision-intensive processes – that are not specific to knowledge work, but we observe no measure of those demands and only 11 of the 55 pairs among the precisely estimated sectors separate at the 5% level, so we offer it as a reading rather than a finding.

One dimension we address only at an exploratory level is the connection between hierarchy characteristics and firm-level outcomes. The relationships we document here are descriptive; establishing causal links between hierarchical structure and firm-level performance is a task for future work.

Firms with taller hierarchies tend to fall in higher revenue categories. Using the twelve SII sales-tier categories, ordered by their revenue bands and available for 93.2% of firm-months, an OLS regression of tier rank on hierarchy depth and log employment[18] puts the coefficient on firm height at $\hat{\beta} = +0.221$ (SE 0.032, $p < 0.001$): one additional layer is associated with roughly a fifth of a tier higher, conditional on employment. Adding industry fixed effects attenuates this by about a third but does not remove it ($\hat{\beta} = +0.141$, SE 0.030, $p < 0.001$), so the association operates both across industries and within them. Standard errors are clustered on the firm in both specifications.

A rigorous treatment of hierarchy-performance linkages requires identification strategies beyond what our descriptive framework provides, but even these exploratory correlations help connect the structural facts to the firm-level outcomes that motivate much of the theoretical literature.

Height dynamics in our panel display a clean two-regime pattern: decreases reflect routine

[18] OLS on an ordinal outcome is a linear approximation; an ordered probit would be more appropriate for formal inference. We use OLS for transparency, and because this exercise is exploratory. Firm-months whose tier is missing or recorded as unavailable are dropped rather than assigned a rank.

peripheral churn, while increases combine mechanical relabeling with concurrent reporting-chain restructuring at middle management. The 2,180 persistent height transitions we observe offer a foundation for future event-study analysis of how wages, spans, and turnover evolve around such events – a natural extension we leave for follow-up work.

These descriptive regularities suggest directions for behavioral and experimental investigation. The salary volatility gradient and span distributions documented here provide empirical targets for models of organizational fairness and authority (cf. Adams, 1965); testing whether these patterns shape worker motivation and cooperation is a natural next step for the experimental economics of organizations.

Several limitations should be noted. Our sample is drawn from clients of a single HR platform, accessed through a research collaboration between the ESE Business School at Universidad de los Andes and the Buk Research Department. The sample likely selects toward more formalized, technologically oriented firms. As documented in Table 20, platform clients are disproportionately medium-to-large formal-sector enterprises concentrated in IT and professional services. The patterns we document may therefore not generalize to informal-sector firms, micro-enterprises, or sectors underrepresented in the sample such as agriculture and mining. The analysis is descriptive throughout: the conditional correlations we report do not identify causal effects of hierarchical structure on outcomes, as sorting, selection, and unobserved firm heterogeneity remain important confounds. Union status is not observable in our data, and Chile's unionization rate of approximately 20% (with substantial variation by sector and firm size) may affect hierarchical structure, wages, and turnover through negotiated work rules and bargaining arrangements. And while the Chilean labor market offers a useful setting – combining OECD membership with developing-country features – cross-country generalizability remains an open question.

Regarding external validity, Chile combines OECD membership and high-income classification with developing-country features; the formalized, knowledge-intensive firms over-represented in our sample sit closer to developed-economy benchmarks than to the smaller, less structured firms that predominate elsewhere in the Chilean economy. Stronger severance mandates than in the United States may compress turnover at upper levels, and the prevalence of family-owned firms and holding-company (multi-RUT) structures may generate hierarchy patterns distinct from Anglo-Saxon economies. These features depart from the American managerial-capitalism configuration that Chandler (1977) historicized, in which structural separation of ownership from management was the institutional default; Chilean firms may instead exhibit a hybrid template combining the formal managerial hierarchy of the modern enterprise with family-controlled ownership. The Bloom et al. (2012) trust-and-decentralization framework is relevant here: Chile's concentrated ownership structures may push toward deeper hierarchies and more centralized control than would be predicted by the country's institutional trust alone. The directional cross-country agreement on the larger-firms-are-taller pattern – documented in our cross-country comparison (Table 1) – is reassuring for the generalizability of the structural pattern, while compensation patterns may be more context-specific given Chile's minimum-wage compression and use of variable pay.

These limitations point toward a productive research agenda. First, causal identification could be pursued through organizational shocks – such as managerial departures or restructurings – or through worker-mover designs that exploit transitions across hierarchical positions within and between firms. Second, the gender disparities we document invite deeper investigation into the determinants and consequences of female leadership, including whether managerial effects on subordinate outcomes differ by gender. Third, the industry heterogeneity in height–size semi-elasticities motivates further study of how specific technologies and knowledge requirements shape hierarchical design. Fourth, cross-country comparisons using similar administrative data would clarify whether the patterns we identify reflect universal features of hierarchy or are shaped by local institutions and labor market structures.

More broadly, our results show the value of administrative HR data for studying internal firm organization. By directly observing reporting structures alongside wages and worker outcomes, these data provide new empirical benchmarks for research on organizational design, incentives, and internal labor markets.

Beyond their academic contribution, these findings carry practical implications. The structural and compensation benchmarks we report – height, spans of control, and pay ratios disaggregated by firm size and industry – provide reference distributions against which firms can evaluate their own organizational design. The finding that salary volatility is highest at lower hierarchical levels has direct relevance for compensation design and retention strategies targeting frontline and mid-level workers.

Declaration of competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The employee-level records are administrative payroll and HR data provided under a confidentiality agreement with the platform operator. They cannot be shared, and the agreement does not permit release in de-identified form either.

The material following the references — data construction, the robustness and sensitivity battery, supplementary results and additional worker-outcome results — is available from the authors on request.

Funding

Felipe Aldunate gratefully acknowledges financial support from ANID (Agencia Nacional de Investigación y Desarrollo), Chile, through FONDECYT Regular Grant No. 1262214.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Claude (Anthropic) in order to improve language clarity and to assist with literature organization. After using this tool, the authors reviewed and edited all content and take full responsibility for the content of the published article.

References

Abowd, J. M., Kramarz, F., & Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, *67*(2), 251–333. https://doi.org/10.1111/1468-0262.00020

Adams, J. S. (1965, January). Inequity in social exchange. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (pp. 267–299, Vol. 2). Academic Press. https://doi.org/10.1016/S0065-2601(08)60108-2

Alchian, A. A., & Demsetz, H. (1972). Production, information costs, and economic organization. *The American Economic Review*, *62*(5), 777–795. Retrieved January 8, 2026, from http://www.jstor.org/stable/1815199

Baker, G., Gibbs, M., & Holmstrom, B. (1994). The internal economics of the firm: Evidence from personnel data. *Quarterly Journal of Economics*, *109*(4), 881–919. https://doi.org/10.2307/2118351

Bandiera, O., Barankay, I., & Rasul, I. (2010). Social incentives in the workplace. *The Review of Economic Studies*, *77*(2), 417–458.

Bandiera, O., Barankay, I., & Rasul, I. (2013). Team incentives: Evidence from a firm level experiment. *Journal of the European Economic Association*, *11*(5), 1079–1114. https://doi.org/10.1111/jeea.12028

Bayer, C., & Kuhn, M. (2019). *Which ladder to climb? decomposing life cycle wage dynamics* (IZA Discussion Paper No. 12473). Institute of Labor Economics (IZA).

Bayer, C., & Kuhn, M. (2023, May). *Job levels and wages* (IZA Discussion Paper No. 16177). Institute of Labor Economics (IZA).

Belloc, F., Burdin, G., & Landini, F. (2023). Corporate hierarchies and workplace voice. *Journal of Institutional Economics*, *19*(6), 729–746. https://doi.org/10.1017/S1744137423000206

Bertrand, M., & Hallock, K. F. (2001). The gender gap in top corporate jobs. *Industrial and Labor Relations Review*, *55*(1), 3–21. https://doi.org/10.1177/001979390105500101

Bloom, N., Sadun, R., & Van Reenen, J. (2012). The organization of firms across countries*. *The Quarterly Journal of Economics*, *127*(4), 1663–1705. https://doi.org/10.1093/qje/qje029

Bloom, N., & Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *Quarterly Journal of Economics*, *122*(4), 1351–1408. https://doi.org/10.1162/qjec.2007.122.4.1351

Bolton, P., & Dewatripont, M. (1994). The firm as a communication network. *The Quarterly Journal of Economics*, *109*(4), 809–839. https://doi.org/10.2307/2118349

Caliendo, L., Monte, F., & Rossi-Hansberg, E. (2015). The Anatomy of French Production Hierarchies. *Journal of Political Economy*, *123*(4), 809–852. https://doi.org/10.1086/681641

Caliendo, L., Monte, F., & Rossi-Hansberg, E. (2018). Exporting and organizational change [Originally NBER Working Paper No. 23630, 2017]. In L. Y. I. Ferretti & M. Yu (Eds.), *World trade evolution: Growth, productivity and employment* (pp. 225–260). Routledge. https://doi.org/10.4324/9781351061544-9

Calvo, G. A., & Wellisz, S. (1978). Supervision, loss of control, and the optimum size of the firm. *Journal of Political Economy*, *86*(5), 943–952. Retrieved October 15, 2025, from https://www.jstor.org/stable/1828417

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, *90*(3), 414–427. https://www.jstor.org/stable/40043157

Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, *50*(2), 317–372. https://doi.org/10.3368/jhr.50.2.317

Card, D., Heining, J., & Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, *128*(3), 967–1015. https://doi.org/10.1093/qje/qjt006

Chandler, A. D. (1977). *The visible hand: The managerial revolution in american business*. Belknap Press of Harvard University Press.

Chandler, A. D. (1962). *Strategy and structure: Chapters in the history of the industrial enterprise*. MIT Press.

Cotter, D. A., Hermsen, J. M., Ovadia, S., & Vanneman, R. (2001). The glass ceiling effect. *Social Forces*, *80*(2), 655–681. https://doi.org/10.1353/sof.2001.0091

Ewens, M., & Giroud, X. (2025, August). Corporate hierarchy. https://doi.org/10.3386/w34162

Fox, J. T. (2009). Firm-size wage gaps, job responsibility, and hierarchical matching. *Journal of Labor Economics*, *27*(1), 83–126. https://doi.org/10.1086/597428

Friedrich, B. U. (2022). Trade shocks, firm hierarchies, and wage inequality. *Review of Economics and Statistics*, *104*(4), 652–667. https://doi.org/10.1162/rest_a_00998

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, *108*(5), 874–904. https://doi.org/10.1086/317671

Garicano, L., & Hubbard, T. N. (2007). Managerial leverage is limited by the extent of the market: Hierarchies, specialization, and the utilization of lawyers' human capital. *Journal of Law and Economics*, *50*(1), 1–43. https://doi.org/10.1086/512706

Garicano, L., & Hubbard, T. N. (2016). The returns to knowledge hierarchies. *Journal of Law, Economics, and Organization*, *32*(4), 653–684. https://doi.org/10.1093/jleo/eww008

Garicano, L., & Rossi-Hansberg, E. (2006). Organization and inequality in a knowledge economy*. *The Quarterly Journal of Economics*, *121*(4), 1383–1435. https://doi.org/10.1093/qje/121.4.1383

Garicano, L., & Rossi-Hansberg, E. (2015). Knowledge-based hierarchies: Using organizations to understand the economy. *Annual Review of Economics*, *7*, 1–30. https://doi.org/10.1146/annurev-economics-080614-115748

Guadalupe, M., Li, H., & Wulf, J. (2013). Who lives in the C-suite? Organizational structure and the division of labor in top management. *Management Science*, *60*(4), 824–844. https://doi.org/10.1287/mnsc.2013.1795

Guadalupe, M., & Wulf, J. (2010). The flattening firm and product market competition: The effect of trade liberalization on corporate hierarchies. *American Economic Journal: Applied Economics*, *2*(4), 105–127. https://doi.org/10.1257/app.2.4.105

Hamilton, B. H., Nickerson, J. A., & Owan, H. (2003). Team incentives and worker heterogeneity: An empirical analysis of the impact of teams on productivity and participation. *Journal of Political Economy*, *111*(3), 465–497. https://doi.org/10.1086/374182

Jung, P., & Kuhn, M. (2019). Earnings losses and labor mobility over the life cycle. *Journal of the European Economic Association*, *17*(3), 678–724.

Kandel, E., & Lazear, E. P. (1992). Peer pressure and partnerships. *Journal of Political Economy*, *100*(4), 801–817. https://doi.org/10.1086/261840

Lawrence, P. R., & Lorsch, J. W. (1967). *Organization and environment: Managing differentiation and integration*. Division of Research, Harvard Business School.

Lazear, E. P., & Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, *89*(5), 841–864. https://doi.org/10.1086/261010

Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. *2008 Eighth IEEE International Conference on Data Mining*, 413–422. https://doi.org/10.1109/ICDM.2008.17

Prendergast, C. (1999). The provision of incentives in firms. *Journal of Economic Literature*, *37*(1), 7–63. https://doi.org/10.1257/jel.37.1.7

Radner, R. (1993). The organization of decentralized information processing. *Econometrica*, *61*(5), 1109–1146. https://doi.org/10.2307/2951495

Rajan, R. G., & Wulf, J. (2006). The flattening firm: Evidence from panel data on the changing nature of corporate hierarchies. *The Review of Economics and Statistics*, *88*(4), 759–773. https://doi.org/10.1162/rest.88.4.759

Rajan, R. G., & Zingales, L. (2001). The firm as a dedicated hierarchy: A theory of the origins and growth of firms. *The Quarterly Journal of Economics*, *116*(3), 805–851. https://doi.org/10.1162/00335530152466241

Rosen, S. (1982). Authority, control, and the distribution of earnings. *The Bell Journal of Economics*, *13*(2), 311–323. https://doi.org/10.2307/3003456

Rumelt, R. P. (1986). *Strategy, structure, and economic performance* (Rev. ed). Harvard Business School Press.

Russo, G., & van Houten, G. (2021, December). Complex Job Design and Layers of Hierarchy. In S. W. Polachek, K. Tatsiramos, G. Russo, & G. van Houten (Eds.), *Workplace Productivity and Management Practices* (p. 0, Vol. 49). Emerald Publishing Limited. https://doi.org/10.1108/S0147-912120210000049004

Song, J., Price, D., Guvenen, F., Bloom, N., & von Wachter, T. (2019). Firming up inequality. *Quarterly Journal of Economics*, *134*(1), 1–50. https://doi.org/10.1093/qje/qjy025

Spanos, G. (2022). Organization and density-related differences in within-firm wage disparities. *Regional Science and Urban Economics*, *95*, 103787. https://doi.org/10.1016/j.regsciurbeco.2022.103787

Tåg, J. (2013). Production hierarchies in Sweden. *Economics Letters*, *121*(2), 210–213. https://doi.org/10.1016/j.econlet.2013.08.001

Williamson, O. E. (1967). Hierarchical control and optimum firm size. *Journal of Political Economy*, *75*(2), 123–138. Retrieved January 8, 2026, from http://www.jstor.org/stable/1829822