NBER WORKING PAPER SERIES

CORPORATE HIERARCHY

Michael Ewens
Xavier Giroud

Working Paper 34162
http://www.nber.org/papers/w34162

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2025, Revised September 2026

We are grateful to Daniel Bias, Lukas Freund, Luis Garicano, Maria Guadalupe, Andrea Prat, and Esteban Rossi-Hansberg, as well as seminar participants at the NBER Organizational Economics Meetings, LSE, the University of Houston, Rice University, SKEMA, INSEAD, HEC, PolyU, University of Oxford (Said), University of Bristol, Nova, University of Georgia, Purdue University, University of Michigan, the University of Warwick, and the University of Porto for helpful comments and suggestions. We thank Shiyang Chu, Joseph Choi, and Sonny Song for their excellent research assistance. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w34162

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Michael Ewens and Xavier Giroud. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Corporate Hierarchy
Michael Ewens and Xavier Giroud
NBER Working Paper No. 34162
August 2025, Revised September 2026
JEL No. D21, D23, G3, J24, L22, M12, M51

ABSTRACT

We introduce a novel measure of corporate hierarchy for 3,454 U.S. public firms. We construct this measure using online resumes of 18.1 million employees and a network estimation technique that allows us to identify hierarchical layers. Equipped with this measure, we document several facts about corporate hierarchies. Firms have on average six hierarchical layers and a pyramidal organizational structure. Firms with more layers have higher internal promotion rates and longer employee tenure. They exhibit higher operating performance and produce more patents, but not higher-quality patents. They also exhibit lower stock return volatility and operating asset volatility. At the same time, flatter firms fared better in response to the COVID-19 shock, consistent with greater organizational agility. We also examine potential determinants of hierarchical structures. Firms facing greater product market threats adopt flatter hierarchies. Moreover, firms increase their number of layers after going public, while they flatten their hierarchies following the adoption of artificial intelligence (AI). These findings are consistent with the theoretical predictions of existing models of corporate hierarchies and suggest that firms adjust their organizational structures to balance the benefits and costs of hierarchical layers.

Michael Ewens
Columbia University
Columbia Business School
and NBER
me2731@columbia.edu

Xavier Giroud
Columbia University
Columbia Business School
and NBER
xavier.giroud@gsb.columbia.edu

1 Introduction

Corporate hierarchies come in many shapes. While some companies adopt vertical, multi-layered hierarchies (e.g., IBM, GE), others prefer flatter organizational structures (e.g., Netflix). Anecdotal evidence abounds with potential pros and cons of different hierarchical structures. For example, it is often argued that vertical hierarchies benefit from clearly defined roles, responsibilities, and accountability, which might enable them to achieve the organization's strategic goals more effectively. On the other hand, these advantages might come at the cost of reduced flexibility and slower decision-making (e.g., McKinsey, 2020).

The theoretical literature provides several frameworks that help conceptualize hierarchies. The early literature sees hierarchies as a form of governance that allows managers to monitor their subordinates better and hence align their incentives with the firm's objective (e.g., Williamson, 1967; Alchian and Demsetz, 1972; Calvo and Wellisz, 1978, 1979; Lazear and Rosen, 1981; Rosen, 1982; Qian, 1994). Other models highlight the role of hierarchies in processing and communicating information within the boundaries of the firm (e.g., Radner, 1993; Bolton and Dewatripont, 1994; Stein, 2002). In these models, a manager's span of control is limited by how much information she can process, which induces a need for hierarchical layers. Layers are costly, however, since they distort the transmission of information within the firm. In a similar spirit, the literature on "knowledge hierarchies" (e.g., Garicano, 2000; Garicano and Rossi-Hansberg, 2006; Caliendo, Monte, and Rossi-Hansberg, 2015) argues that hierarchies allow workers and managers to specialize and communicate, allowing managers to solve the more complex problems faced by the production workers. In these models, by adding a hierarchical layer, the company invests in a manager who acts as a "problem solver" for the workers under her supervision. The manager is costly, however, since she does not engage in production and has a fixed training cost. Ultimately, the decision to add a layer is driven by the trade-off between the benefits and costs of adding more problem solvers to the organization.

While the theoretical literature provides helpful guidance in thinking about corporate hierarchies, there is surprisingly little empirical work that aims to understand the determinants and implications of different organizational structures. The existing evidence comes primarily from middle-sized firms in Europe with relatively straightforward organizational forms (e.g., Caliendo, Monte, and Rossi-Hansberg, 2015; Caliendo, Mion, Opromolla, and Rossi-Hansberg,

2020; Gumpert, Steimer, and Antoni, 2022; Bias, Lochner, Obernberger, and Sevilir, 2026).[1] In contrast, much less is known about the hierarchies of more complex organizations such as U.S. public companies. In this regard, the main evidence to date has been gathered by Rajan and Wulf (2006), Guadalupe and Wulf (2010), and Guadalupe, Li, and Wulf (2014), who rely on survey data from the consulting firm Hewitt Associates for a sample of about 300 U.S. firms over the years 1986-1999. These data provide information on the number of managers under the supervision of the CEO and other C-suite executives. While these data allow them to study the top layers of the organizational chart, they only capture a subset of the hierarchy.

This scarcity of empirical evidence is due to the lack of comprehensive data on corporate hierarchies. Even the most granular datasets available to date—such as the Longitudinal Employer-Household Dynamics (LEHD) of the U.S. Census Bureau that contains employer-employee matched data for U.S. firms—do not allow researchers to characterize hierarchies. In the LEHD, information is available on each employee's salary, along with employee characteristics (such as gender, race, and highest educational degree). Still, no information is collected on the employee's role and rank within the organization. This makes it challenging to construct informative metrics of corporate hierarchies.

In this paper, we propose a novel measure of corporate hierarchy that can be used to study publicly traded U.S. companies. Equipped with this measure, we then explore the role of hierarchies for corporate decision-making and a wide range of outcomes, guided by the various theories of corporate hierarchies.

To construct our hierarchy measure, we combine data on worker job histories with a network estimation technique that allows us to infer corporate structures. A large sample of online resumes matched to public firms provides a view of the firm's internal labor market, allowing us to track job transitions and title changes. We use these primary inputs to determine the number of hierarchical layers using the network estimation methodology of Huitfeldt, Kostøl, Nimczik, and Weber (2023). Intuitively, job transitions within a firm help estimate job ranks. The estimated job ranks come with statistical uncertainty because transitions can be promotions, demotions, and lateral moves. Identifying the most likely set of lateral moves provides information to infer the layers of the corporate hierarchy.

The variable construction leverages online resume data from LinkedIn for 18.1 million

[1] Moreover, most of this work characterizes organizational forms based on occupation codes (rather than job titles or ranks), which need not capture hierarchies per se.

workers at U.S. public firms from 2015 to 2024. The measure we obtain, which captures the number of hierarchical layers, indicates that the average firm has six layers and that the number of layers has been flat to weakly increasing over the sample period.

We validate our hierarchy measure in a number of ways. First, we show that the layers group workers into reasonably homogeneous groups based on their title and work experience. Second, while the Huitfeldt, Kostøl, Nimczik, and Weber (2023) methodology does not impose any restriction on the shape of the hierarchy, we nevertheless find that the average firm's hierarchy is pyramidal (i.e., more workers are in the bottom layers). This is reassuring, as most models of firm hierarchy predict this shape. Third, the fraction of workers with "manager" or an executive role in their title increases significantly as we move up to higher layers. This is again consistent with the models of firm hierarchy in which employees at higher layers are more likely to engage in managerial activities. Fourth, we show that our measure of hierarchical layers performs well when benchmarked against a small set of firms with known hierarchical layers. Fifth, the roles assigned to each layer line up with external measures of the work they involve. BLS occupational wages rise as one moves up the hierarchy; the O*NET work styles that gain importance up the hierarchy are leadership, innovation, and initiative, while attention to detail and independence are flat; and the O*NET skill requirements of the occupations increase at each successive layer. Overall, these patterns indicate that our measure captures several inherent features of corporate hierarchies.

We also examine potential limitations of using online resume data. Because online resumes are voluntarily disclosed, they may provide more complete coverage of some occupations than others. We therefore compare the occupations and employment shares in the resume data with those reported by the Bureau of Labor Statistics (BLS) for the same industries, identify where underrepresented occupations are located within firms' hierarchies, and use simulations of the estimator to assess whether these gaps may systematically bias our estimates of the number of layers. Our simulations indicate that, at the observed levels of coverage in the LinkedIn data, our estimate is unlikely to be biased in a material way.

After validating our measure, we begin the analysis by exploring the relationship between the number of layers and different dimensions of firms' human capital. We find that advanced degrees are more prevalent at higher hierarchical layers. This pattern is consistent with models of knowledge hierarchies, in which higher layers add "problem solvers" who require greater training costs. We further find that employees in firms with more layers have longer tenure, with promotions playing a greater role in transitions than outside hires. This finding is

consistent with the notion that hierarchies provide employees with clearer career trajectories and promotion paths.

Next, we examine the relationship between hierarchies and a variety of outcomes. First we show that high-layers firms are more profitable—as captured by the return on assets (ROA) and net profit margin (NPM)—but generate lower revenues per employee. We further show that high-layer firms generate more patents. However, there is no significant difference in patent quality (as measured by patent citations).

Second, we show that high-layer firms exhibit lower stock volatility. This lower volatility is paired with lower operating asset volatility rather than lower financial risk (e.g., leverage). These differences align with the models of knowledge hierarchies, in which high-layer firms have more specialist managers—that is, "problem solvers"—who reduce the firm's operating volatility through their problem-solving activities.

Third, we examine whether firms with flatter hierarchies are more "agile" in response to shocks affecting their operating environment. Specifically, we investigate the relationship between the number of organizational layers and the stock price reaction to the COVID-19 shock. We find that firms with flatter hierarchies experienced a smaller stock price decline in response to the COVID-19 shock. This complements the findings of Fahlenbrach, Rageth, and Stulz (2021) that firms with greater financing flexibility fared better during COVID-19; our results suggest that *organizational flexibility* also contributes to firms' resilience to shocks. This evidence echoes the models that emphasize the role of hierarchies in communicating information within the firm (e.g., Radner, 1993; Bolton and Dewatripont, 1994). In these models, flatter hierarchies facilitate more direct communication with higher layers of the organization, potentially accelerating decision-making and enhancing corporate agility in response to external disruptions.

In the last part of the paper, we flip the regression specification and examine the determinants of corporate hierarchies. First, we examine the relationship between product market competition and the number of layers. To capture product market competition, we use the fluidity measure of Hoberg, Phillips, and Prabhala (2014), which captures the rate at which a firm's product market is disrupted by competitive threats. We find that firms with more fluid product markets—that is, firms that face a higher rate of product market disruption—tend to adopt flatter hierarchies. This finding complements our COVID-19 results, suggesting that flatter hierarchies may help firms maintain corporate agility and respond more quickly to disruptions in their operating environment.

Second, we examine how firms adjust their hierarchies over their life cycle, focusing on the transition from private to public ownership. Going public increases the complexity of a firm's operations (e.g., by imposing additional reporting requirements). One prediction of the models of knowledge hierarchy is that firms respond to higher complexity in their environment by increasing the number of layers, as more layers add "problem-solvers" who are better qualified to handle the higher complexity faced by the organization (e.g., Garicano, 2000; Caliendo, Monte, and Rossi-Hansberg, 2015). Consistent with this argument, we find that firms substantially increase their number of hierarchical layers following their IPO. Interestingly, the increase begins before the IPO itself, suggesting that firms may expand their hierarchies in preparation for going public.

Third, we examine how technological change affects corporate hierarchies. A common prediction of models of corporate hierarchy is that a reduction in the cost of acquiring, processing, and communicating knowledge reduces the need for hierarchical layers (e.g., Radner, 1993; Bolton and Dewatripont, 1994; Garicano, 2000). In particular, in models of knowledge hierarchies, managers create value by acting as specialized problem solvers: rather than requiring each worker to acquire the knowledge needed to handle every possible problem, the organization can concentrate specialized knowledge in managers and route difficult or unusual problems up the hierarchy. Hierarchical layers are therefore a way of economizing on the knowledge that workers need to possess individually. Technologies that make knowledge easier to acquire or problems easier to solve reduce the benefits of this specialization. Artificial intelligence (AI) represents a particularly important example because it can perform tasks that were previously assigned to specialized managers, including information retrieval and processing, decision support, and the resolution of problems that require specialized knowledge. In this sense, AI can substitute for some of the problem-solving and information-processing functions performed by managers, allowing employees at lower levels of the organization to handle a broader set of tasks without relying on hierarchical escalation. The resulting reduction in the value of managerial specialization might lead firms to operate with fewer layers. Consistent with this prediction, we find that firms flatten their hierarchies following the adoption of AI technologies. Taken together, the IPO and AI results suggest that firms use hierarchical layers to expand their problem-solving capacity when organizational complexity increases, and reduce them when technology lowers the cost of acquiring and processing knowledge.

The remainder of this paper is organized as follows. Section 2 discusses the theoretical underpinnings and our conceptualization of hierarchical layers. Section 3 reviews existing

measures of hierarchy. Section 4 introduces our novel measure. Section 5 describes the data. Section 6 provides a validation of our metric. Section 7 presents the analysis of firm outcomes. Section 8 examines the determinants of corporate hierarchies. Finally, Section 9 concludes.

2 Corporate hierarchy: theory

This section summarizes the theoretical literature on corporate hierarchy that guides our variable construction, variable validations, and empirical tests. For comprehensive surveys of corporate hierarchy theory, see Garicano and Van Zandt (2013) and Malenko (2024).

2.1 Models of corporate hierarchy

It is useful to start with a summary of how economists view the organization. Caliendo, Mion, Opromolla, and Rossi-Hansberg (2020) write that the organization must "determine how much each employee knows, how many employees to hire, and how many layers of management to use in production" (p. 4213). Garicano (2000) argues that organizations exist "to solve coordination problems in the presence of specialization" (p. 874). Given that most firms' production depends on the skills of their workforce, the organizational problem can be viewed as a means to allow workers to combine their time and knowledge, which allows workers to gain the help of other workers in production (Garicano and Rossi-Hansberg, 2006).

One class of models views the firm as a "knowledge hierarchy" (e.g., Garicano, 2000; Garicano and Rossi-Hansberg, 2006; Caliendo, Monte, and Rossi-Hansberg, 2015). In these models, workers must contribute time and knowledge to produce, and workers depend on the output of other workers to complete production. A knowledge hierarchy divides tasks—or the knowledge space—among workers. Layers emerge with managers who act as problem solvers and handle exceptions from the layers below. Adding layers thus adds specialists to the workforce, allowing the firm to leverage worker knowledge at the lower levels better. The hierarchy ensures that managers deal with the least common problems, leaving the more common ones to the lower layers of the firm. As Caliendo and Rossi-Hansberg (2012) summarize, when we observe an organizational structure, we can infer the agents' knowledge, how information flows within the firm, and the communication paths.

These benefits of hierarchy come with costs. The first is the fixed cost of training managers, as they must learn how to solve problems regardless of how the firm uses the knowledge. The second is the communication required between managers and workers that does not generate

output. Accordingly, when a firm adds a layer to its hierarchy (and hence another set of problem solvers), it weighs the gains from specialization against the cost of coordination and lost production time. Ultimately, firms choose a structure—the number of layers and the share of workers in each layer—to best match problems to knowledge. The organization then tries to economize on the knowledge of managers to solve problems.

The main ideas from the models of knowledge hierarchy echo well with a second class of models, which we refer to as "information hierarchy" (e.g., Radner, 1993; Bolton and Dewatripont, 1994; Stein, 2002). In these models, the span of control of the CEO is restricted by the CEO's ability to process information, which induces a need for delegation in the form of hierarchical layers. As in the models of knowledge hierarchy, adding layers gives rise to a trade-off between the gains from specialization and the costs of communication. In these models, specialized workers collect information (within their layer), which they then communicate to managers (across layers). If the costs of communication are high, multilayered hierarchies become less attractive.

Finally, a third class of models sees hierarchies as a form of governance that allows managers to better align their subordinates' incentives with the firm's objective (e.g., Williamson, 1967; Alchian and Demsetz, 1972; Calvo and Wellisz, 1978, 1979; Rosen, 1982; Qian, 1994). These models, which we refer to as "incentive hierarchy," argue that, because much of employees' incentives are provided through promotions, their effectiveness depends on the career path within the organization's hierarchy. In these models, multilayered hierarchies provide a career path that is conducive to internal promotions and hence contribute to employees' incentives. Relatedly, tournament models view promotions as prizes in contests among employees, with the prospect of advancement providing incentives for effort (e.g., Lazear and Rosen, 1981; Malcomson, 1984). Accordingly, the number and structure of hierarchical levels shape both the opportunities for promotion and the incentives generated by these internal contests.

Our analysis studies whether empirical analogs of these models' assumptions can be found in our sample, and then explores the empirical relevance of these models for firm decisions and outcomes.

2.2 What is a layer?

Before introducing our hierarchy measure, we describe how we conceptualize hierarchical layers throughout the paper. Caliendo and Rossi-Hansberg (2012) and Caliendo, Monte,

and Rossi-Hansberg (2015) define a layer as a group of employees with similar knowledge levels performing tasks at a comparable level of authority. In keeping with this definition, our measure of hierarchical layers aims to capture jobs with similar complexity, tasks, and responsibilities. Another way to characterize layers is in terms of information flows between roles. For example, in the model of Bolton and Dewatripont (1994), workers in a given layer send messages to those above but not below. In the knowledge hierarchy model of Garicano (2000), a layer is a set of workers with the same scope of knowledge and who communicate problems to the layer above them. Since we do not observe communication flows across roles, we cannot measure layers in this fashion. Nevertheless, we believe that our measure captures similar duties and close rankings within the firm, in keeping with the above conceptualization of hierarchical layers.

3 Existing measures of hierarchy

Our attempt to measure corporate hierarchy is not the first. Appendix Table A1 summarizes sixteen papers that capture the broad scope of empirical work in this area. The table allows us to discuss different methodologies and resulting counts of hierarchical layers.

There are three methods that researchers use to identify job rankings and levels. The first approach is to use direct reports or surveys from firms or consultants (e.g., Baker, Gibbs, and Holmstrom, 1994; Liberti and Mian, 2009; Guadalupe and Wulf, 2010; Guadalupe, Li, and Wulf, 2014; Garicano and Hubbard, 2016). This approach minimizes measurement error but lacks external validity due to the small number of firms for which such data are available. The second and most common approach is to connect job titles to occupation codes (e.g., Caliendo, Monte, and Rossi-Hansberg, 2015; Gumpert, Steimer, and Antoni, 2022; Bias, Lochner, Obernberger, and Sevilir, 2026). The key assumption underlying this method is that the occupation codes reflect job or task complexity, and such complexity correlates with rank. The unique number of occupational codes limits this approach, however. Appendix Table A1 shows that the number of levels based on occupation codes does not exceed 4. Finally, the third approach, which is the one we implement in this paper, uses observed transitions between jobs to infer rankings. One can then infer levels by using an algorithm that captures the "clustering" of jobs into promotions. This approach leverages the dynamics of a firm's internal labor market. Baker, Gibbs, and Holmstrom (1994) were the first to implement this approach and Huitfeldt, Kostøl, Nimczik, and Weber (2023) introduce a formal algorithm to

create hierarchical layers.

While each method and approach has strengths and weaknesses, we believe the internal labor markets method is ideal for public firms. First, alternative methods have too few layers to capture the complex structure of large public firms. Second, it is unclear whether a general mapping between occupation codes and job titles would provide a "one size fits all" that is suitable for a wide variety of public firms across different industries. Using revealed transitions in an internal labor market *at the firm level* circumvents this limitation, as this approach does not rely on a general mapping and classification system. Finally, the data that are obtained as a by-product of the internal labor markets approach are helpful as they capture promotion rates, work experience, and other human capital measures.

4 A new measure of hierarchy

4.1 Internal labor markets: an algorithm

We aim to create a firm-year measure of the number of hierarchical layers. Suppose we have data on worker resumes over a 5- to 8-year horizon, where resumes reveal titles, start dates by title, and transitions between jobs. Further, assume that titles are standardized within the firm. This data provides a view of a firm's internal labor market and the inputs we need to implement the method of Huitfeldt, Kostøl, Nimczik, and Weber (2023).[2] There are two critical pieces of intuition for this method. First, job transitions ("worker flows") reveal the ranking of jobs (Clauset, Arbesman, and Larremore, 2015). Second, a layer in a firm is a collection of roles that are more likely to have lateral moves between them rather than promotion or demotion (Bonhomme, Lamadon, and Manresa, 2022).[3]

The first step is to identify the connected jobs in the firm, or internal labor markets (ILMs). The method requires tracking connections between roles over the sample period, effectively finding the largest connected roles in the firm. Firms can have multiple ILMs if they have separate divisions where job roles never transition to one another (e.g., a diversified firm that produces jet engines and financial services). Ideally, a firm has few ILMs, or the largest ILM comprises a large fraction of worker-years. Figure A1 shows example ILMs from the data for four sample firms. Our analysis will focus on the largest ILM in the firm. Fortunately,

[2] An earlier formulation of a similar methodology is found in Baker, Gibbs, and Holmstrom (1994).

[3] What follows is a high-level summary of their method. Interested readers should consult Huitfeldt, Kostøl, Nimczik, and Weber (2023) for more details.

this ILM comprises over 85% of the worker-years for the median firm in our sample (see Section 6.1). One cost of this focus is related to multi-divisional firms, as the method can only summarize the layers of the largest division. Nevertheless, this is unlikely to be a concern for our analysis as long as firms structure their divisions—controlling for size—in a similar fashion.

The second step leverages the worker transitions in the largest ILM to create a role ranking.[4] Note that all these steps are performed at the firm level. We use all years available in the resume data (see Section 5). Intuitively, if role A is followed by role B (50%) and role C (50%), while role C only transitions to A 5% of the time, the ranking would put B and C above A. Huitfeldt, Kostøl, Nimczik, and Weber (2023) introduces a minimum violation ranking (MVR) to account for the noise in resume data and job transitions. The algorithm starts by ranking roles based on the promotions they send, with the role sending the most promotions at the bottom. The method then swaps ranks of connected (or random) role pairs and checks whether the implied ranking has the same or a higher number of promotions (accepting it if it does). This procedure likely has many ties, so the final ranking is the average of optimal ranks. The uncertainty of a given rank—the standard deviation—is used in the final step. Figure A2 plots the estimated ranks, their uncertainty, and the resulting layers for four sample firms.

Once roles are ranked, we put them into hierarchical levels or layers. This is facilitated by the estimated mean rank and the rank standard deviation obtained from the estimation's replications. Figures A3 to A5 in the Appendix provide examples of how ranks become layers, and Appendix A1 provides a walkthrough of the estimator.

A layer in the hierarchy is a group of adjacent ranks. For a given K, roles are assigned to K layers using one-dimensional k-means on their mean ranks to minimize within-layer rank dispersion, following Huitfeldt, Kostøl, Nimczik, and Weber (2023). Intuitively, a role whose rank floats equally across good orderings has lateral peers, and that is a layer. Rank uncertainty enters only through the choice of K. Conditional on K, roles are grouped by their mean ranks alone, so roles with overlapping rank uncertainty share a layer when their mean ranks are also close.

We then choose the number of levels in the hierarchy, K, by continuing to add layers until the dispersion of the mean ranks left inside those K layers is no larger than the noise in the

[4] We can alternatively rank roles using their O*NET Job Zone instead of the ranking observed from promotions. That ranking produces no variability in the number of layers across firms (5), providing more evidence that our new measure is more informative about firm hierarchies.

ranks themselves. Here, the noise floor is the average squared rank standard deviation over the roles the data can take. Roles whose rank is indistinguishable from random placement are left out of the noise floor and the choice of K, but they are still assigned to a layer. We scale this floor by the firm's LinkedIn coverage, capped at 100%, so incomplete coverage can only make the standard stricter, weakly adding layers.

For some of the largest promotion-heavy firms with many incomparable roles, the noise-floor rule above breaks down. Simply, the search for the minimum violation ranking does not settle, no matter how much computation is done. The staged search flags these firms. If too little search is done, the sampled noise is too high, and the rule collapses these firms into too few layers. If we instead give the search more compute, the replications agree on an arbitrary ordering of incomparable roles, and the number of layers increases toward one role per layer. We implement an alternative rule, motivated by another estimation approach in Huitfeldt, Kostøl, Nimczik, and Weber (2023), that does not depend on the noise floor whose sensitivity we just described.

Here, we sort roles by mean rank and repeatedly merge the neighboring pair of roles whose ranks are least separated by the data. We stop when every neighboring pair differs by more than $z(N)$ standard errors, where $z(N)$ is the Bonferroni family-wise threshold at 5%. Huitfeldt, Kostøl, Nimczik, and Weber (2023) stop when a gap exceeds one standard deviation; the family-wise threshold is ours. Because $z(N)$ rises with the number of roles, the threshold gets stricter as roles multiply, so the rule does not manufacture layers in large firms the way a shrinking noise floor does.[5] We describe our final measure of hierarchical layers in Section 6 after discussing the data inputs.

5 Data

5.1 LinkedIn data

We leverage data on worker resumes provided by two vendors, Revelio and CoreSignal.[6] Both extract resume data—including work histories, job titles, employers, and educational background—from publicly available profiles of U.S.-based workers. The Revelio panels cover

[5] Given K, the assignment of roles to layers is the same k-means step for every firm, whichever rule selects K. Of the 3,454 sample firms, 1,962 (57%) use the noise-floor rule and 1,492 (43%) the alternative rule. For the latter, the two rules give similar counts at the stage where the search stopped: mean layers 7.31 under the noise-floor rule and 7.53 under the alternative rule, medians 7 and 7, identical in 29% of firms. Appendix A1 gives the full rule; Figure A5 shows a firm where the search does not settle.

[6] See `https://www.reveliolabs.com/` and `https://coresignal.com/`.

2013 to 2024; the CoreSignal panels run through 2023.[7] While our analysis focuses on workers who are active during the sample period, we use their complete work history for worker-level variables such as tenure and years of experience. The promotion network that ranks roles is built from title changes observed inside the sample window; a job that began before the window and continues into it enters at the window's first year. We combine Revelio and CoreSignal to draw on the strengths of each: CoreSignal carries longer work histories, while Revelio is updated continuously. Revelio is our default. The vendor is fixed once per firm, so a change of vendor can never register as a within-firm change in layers. For a firm covered by both providers, we use CoreSignal only when the firm's most recent Revelio year does not yield a usable ranking with at least two layers and at least twenty roles, or when the two panels' role counts differ by more than a factor of three, a sign of a mismatched firm. Of the 3,454 sample firms, 3,050 come from Revelio and 404 from CoreSignal, of which 141 have no Revelio match.

We exclude interns, summer and incoming hires, retirees, board members, and contractors, identified from their titles. These profiles are self-filled by individuals, and while the online resume site does require some structure, there are few requirements for titles or accuracy of job timing. Admittedly, it is challenging to assess the coverage of the actual set of workers at firms or the sample selection within firms. Which workers appear, and which do not, is examined against BLS occupation data in Section 6.4.

A comparison of headcounts between our data and Compustat reveals that coverage is 68% for the median firm (see Section 6.1). However, this likely understates the actual coverage since the Compustat headcount includes temporary workers and non-U.S. employees. Importantly, the algorithm to measure hierarchical layers described below is largely robust to over- or under-representation of a given title.[8] Rather, all we require is that the titles occur at least once in the data and that a promotion or demotion exists between titles.[9] In the final sample described below, we have 6,687,972 title transitions from 2013 to 2024.[10] Again, we caution that we do not have a means of directly assessing whether we have complete coverage of titles or transitions. To validate our data, we rely on an extensive set of validation tests that

[7] For Revelio we use the raw job-title string as written on the profile, not Revelio's standardized role or occupation fields, so both sources enter the title cleaning described below on the same footing.

[8] Our cleaning of the LinkedIn data folds rare title variants into their parent role rather than dropping them, and removes a job transition only when its reverse also occurs in the same year with at most two workers.

[9] In the current sample, 31 firms have too few transitions between jobs to form a promotion network; all of them also fail the coverage or employee-count filters described below.

[10] A transition is a change in a worker's title between consecutive observations at the same firm, with both observations in the sample window.

demonstrate the data capture features of hierarchies predicted by theory and consistent with industry practice (see Section 6.2).

5.1.1 Title cleaning, clustering, and roles

One caveat of using worker-submitted resumes is that they feature a large set of often idiosyncratic titles. Indeed, the average firm in the resume data has an average of 2,839 unique titles (median 843) across the sample period. Although our methodology does not directly map jobs to occupational codes, the large number of titles makes it challenging to classify transitions accurately. We thus implement a cleaning and clustering method to reduce the number of titles within each firm. The cleaning step removes commonly added suffixes and prefixes, normalizes abbreviations, and removes rare words where "rare" is defined across and within firms. Additionally, we apply stemming (i.e., reducing words to their primary part) to titles to standardize different forms of the same role; for instance, "Software Engineer" and "Software Engineering" are both stemmed to "Software Engin." This cleaning step alone reduces the average number of titles per firm to 2,671 (median 812).

The next step recognizes that individuals may use multiple descriptions for the same role, often sharing common keywords. Our clustering algorithm identifies common phrases across titles within firms based on the frequency of strings or bi-grams. Specifically, it keeps the most frequent strings and bi-grams for each firm (the exact frequency rule is stated in Appendix A1), prioritizing bi-grams within a title. For example, consider a firm with only two titles over the years: "Software Engin, UX" and "Software Engin, backend." For individual strings, the algorithm would detect "Software" and "Engin" each appearing twice, while "UX" and "backend" appear once each. For bi-grams, "Software Engin" would appear twice, and "Engin UX" and "Engin backend" would each appear once. Since the algorithm prioritizes bi-grams over individual strings, both titles would map to the single role "Software Engin." If a title does not map to a common phrase, the algorithm maps it to a common string using the most common role in the set. This process results in a meaningful reduction in the number of titles, now classified as roles, with an average of 304 roles (median 188) per firm. At the median firm, each role encompasses about 4.5 raw titles (4.3 cleaned titles).

5.2 Public firm data

We merge the resume data with data on U.S. publicly-traded firms from Standard & Poor's Compustat. This merge is facilitated by the company's website (URL), which is reported in both datasets, and by Revelio's own match of firms to Compustat gvkeys. As there are still many false positives and negatives, we supplement the merge with a fuzzy string comparison of firm names and a manual review of hundreds of potential matches. In total, we match 8,128 firms to the resume site data (across all years and industries) before filters and additional data cleaning. These firms have 30.6 million unique workers in all years (again, before cleaning). After filtering from data cleaning and implementing the algorithm, we are left with 18.1 million unique workers.

The final sample of firms includes those in Compustat as of 2015 or later that have at least one year of measured hierarchical layers. We exclude firms in the following industries based on SIC codes: 4900-4999 (regulated utilities), 5200 and 5300-5399 (building-materials and general-merchandise retailers, with many part-time workers), 7100-7299 (service firms with poor LinkedIn coverage), and 4400-4599 (airlines and shipping, with a large fraction of firms with near-zero coverage on LinkedIn). Any remaining firms after these filters, whose LinkedIn employee count is less than 15% of Compustat employees (Compustat item "emp") or with less than 25% of worker-years in the largest internal labor market are dropped. We also require a firm to report employees, total assets, and a U.S. incorporation and headquarters in Compustat, to have at least ten employees in some sample year, and to be an operating company rather than a REIT (Real Estate Investment Trust), BDC (Business Development Company), or MLP (Master Limited Partnership) shell. At the firm-year level, we drop years in which the LinkedIn headcount exceeds twice the Compustat headcount and years with a single estimated layer; Appendix A2 lists every condition with the number of firms it removes. With these criteria, we are left with 3,454 firms (though the sample size varies depending on the regression). The remaining 1,755 firms outside the excluded industries fail at least one of the conditions listed in Appendix A2, which reports the number of firms each condition removes. Table 1 compares the two samples. In-sample firms tend to be larger (based on total assets), have higher leverage, and higher profitability (based on ROA).

We use several other datasets in the paper, which we briefly introduce here. The stock price data are obtained from CRSP and merged using the CRSP-Compustat bridge maintained by WRDS. Execucomp provides information on the top executives (the C-suite) that we combine

with the resume data, along with information on their compensation. BoardEx provides information on board size. Finally, we use the patent data of DISCERN (Arora, Belenzon, and Sheer, 2021) and Kogan, Papanikolaou, Seru, and Stoffman (2017), as well as the data of Babina, Fedyk, He, and Hodson (2024) on public firms' use of artificial intelligence (AI).

6 Hierarchical layers

This section describes the measure of hierarchical layers and presents a series of validation tests motivated by the theory literature reviewed in Section 2.

6.1 Summarizing the data

What do corporate hierarchies look like over time and across firms? Panel A of Table 2 reports summary statistics on layers and other characteristics. The unit of observation is a firm observed in the most recent year of its observable hierarchy. Recall that the clustering algorithm (see Section 4) identifies the firms' internal labor markets (ILMs) from job transitions and infers hierarchical layers from the largest ILM. Panel A shows that the average firm has 104 ILMs. Most of these are single-role titles that never connect to another role, a feature of user-entered LinkedIn titles. Counting only ILMs with at least two roles connected by observed worker moves, the average firm has 8 ILMs. The hierarchical layers measure focuses on the largest ILM, which contains 191 unique roles on average (recall that titles are mapped to roles). On average, there are 304 roles in a firm.

Two statistics provide comfort that the largest ILM captures most roles and workers. Specifically, for the median firm, 59% of the roles and 85% of the worker-years are found in the largest ILM. That is, most of the workers observed on LinkedIn help identify a firm's corporate hierarchy. The table also reports the number of employees recorded in Compustat and those found in LinkedIn. The former includes part-time, seasonal, and non-US employees. The LinkedIn data includes US-based employees with public profiles. As can be seen, the LinkedIn data covers 68% of the number of Compustat employees for the median firm. While this is an inherent limitation of the LinkedIn data, this is less of a limitation for our algorithm since it only needs to capture job transitions across roles within the firm.

The average number of hierarchical layers across all firms is 6.4 (median 6). This exceeds the number of layers—typically 4—obtained in the studies reviewed in Appendix Table A1.[11]

[11]The exceptions are the three papers that study a single firm—such as Adhvaryu, Bassi, Nyshadham,

The higher number of layers at U.S. public firms is not surprising, since they are larger than most firms considered in the prior literature. Importantly, this higher number of hierarchical layers suggests that the existing methods (e.g., those based on occupational codes that have a built-in cap on the number of layers) may be too coarse to capture the rich organizational structure of publicly traded firms.

The last three variables in Panel A report different versions of the "hierarchization" measure of Bias, Lochner, Obernberger, and Sevilir (2026) that captures the relative share of employees in each layer.[12] The first two measures use their partition based on occupation codes. We construct these two measures by mapping job titles to occupation codes (using the German Classification of Occupations and the O*Net occupation codes, respectively). The third measure uses our definition of layers. The average of 0.77 is higher, which is expected since this variable increases in the number of layers. Note that all three exceed the 0.41 average in Bias, Lochner, Obernberger, and Sevilir (2026), who study the hierarchization of German firms before and after going public. Their firms are substantially smaller than the U.S. public firms in our sample, which likely explains the lower degree of hierarchization.

Panel B reports the statistics from Panel A for the complete set of firm-year observations. These statistics mirror those from Panel A.

Moving to the time series evidence, Figure 1 shows that the number of layers for U.S. public firms has been flat to weakly increasing over the sample period: the average rises from 6.39 in 2015 to 6.55 in 2024 (+0.16 layers), while the median stays at 6 throughout. Over 2015-2024, which post-dates the samples of Rajan and Wulf (2006) and Guadalupe and Wulf (2010), the flattening they document for the top of the hierarchy does not continue in this larger sample of public firms, which takes into account a more complete characterization of their hierarchy.

Finally, Figure 2 shows that hierarchical layers vary substantially across industries, as well as across firms within industries (Panel A), while there is less variation within industries over time (Panel B). The latter further shows that many industries experience a weak increase in their number of layers, in keeping with the broader evidence from Figure 1.

Tamayo, and Torres (2023) who identify 16 layers in their study of a large Argentinean automobile firm—as well as Huitfeldt, Kostøl, Nimczik, and Weber (2023), who develop the network estimation methodology that we build on.

[12] This measure is constructed as follows. Suppose a firm has K layers and β_{ij} is the share of employees in layer j for firm i. Hierarchization is then $1 - \sum_{j=1}^{K} \beta_{ij}^2$. Thus, a flat firm has a value of 0 and a firm with equal shares across K layers approaches 1 as K increases.

6.1.1 Within-firm variation

Baker, Gibbs, and Holmstrom (1994)'s analysis of a single medium-sized services firm from 1969–1988 showed stability in the organizational form. As they note (p. 893), "the hierarchy was surprisingly stable over time. With few exceptions, titles important in 1969 were similarly important in 1988." We observe similar stability among U.S. public firms in our sample. During our sample period from 2015–2024, only 19% of firms experienced a change in layers, and this occurred in 6.4% of firm-years. When layers do change, they increase 66% of the time, and the median change is one layer. This stability limits within-firm inference. The results in Section 7 rely on differences across firms, and the within-firm designs later in the paper (the IPO and AI analyses) are identified off the 654 firms whose layer count changes during the sample period.

6.2 Validation

In this section, we present a series of tests to validate our proposed measure of corporate hierarchy.

6.2.1 Ranks, titles, and promotion

Ultimately, a measure of hierarchy should correctly rank workers and jobs. The first set of validation tests examines whether layers have the expected titles, work experience, and connections via promotion.

Table 3 compares the most common titles in firms' top layers with those in their bottom layers in 2024. Reassuringly, the titles one would expect at the top, those of management and leadership such as "vice president," "director," and "managing director," are far more likely to appear in the top layer, with the exception of "financial advisor," while the bottom layer holds the expected titles, such as "sales associate," "barista," and "cashier." This split in titles suggests the algorithm is doing its work, both in ranking roles and in grouping them into layers.

The hierarchical layers also correlate with external measures that characterize the types of positions one would expect at different levels of the hierarchy, none of which enter the estimation:

- *Work styles.* The content of work changes with hierarchical position. Appendix Figure A6 maps each title to its O*NET occupation and reports the importance of 16 work

styles by within-firm layer quintile. Leadership is far more important in the top quintile than in the bottom quintile, followed by innovation and initiative, while attention to detail, self-control, and independence are essentially flat across the hierarchy.

- *Wages.* Wages rise with hierarchical position. Using the merge of LinkedIn titles to BLS occupations described in Section 6.4, occupational wages increase monotonically across quintiles of a firm's roles sorted by rank or by layer (Appendix Table A2). Occupations in the top rank quintile pay about twice as much as those in the bottom quintile, and sorting by layer yields a similar pattern.

- *Skill requirements.* As one moves up the hierarchy, skill requirements rise, as measured by the O*NET Job Zone, a five-step scale of the education and training an occupation requires. Occupations in the top rank quintile are more than one full Job Zone above those in the bottom quintile, roughly corresponding to the gap between jobs that require vocational training or an associate's degree and jobs that require a bachelor's degree. Moreover, the difference is positive at every step between quintiles.

A hierarchical structure should also group workers into experience buckets. Workers in the lower layers should have fewer years of work experience, either because they have not yet been promoted or because being at higher levels requires more know-how that takes time to accumulate. Figure 3 shows that work experience—measured since the worker graduates or starts their first job—increases with hierarchical layer.[13] As experience tends to coincide with higher pay, this pattern is consistent with most theory models of hierarchies (e.g., Lazear and Rosen, 1981; Malcomson, 1984). Overall, the characteristics of titles and worker experience across layers confirm that our measure captures meaningful features of firm hierarchies.

Finally, Appendix A3 benchmarks the measure against an industry with a standardized job ladder. Investment banks promote through a known sequence of titles, from analyst to associate, director, and managing director, so the estimated ranks and layers can be checked against an order that is public knowledge. At Jefferies and Evercore, the estimated ranks place these titles in the expected order, with the entry rungs in separate layers. Re-running the ranking and clustering on the investment-banking roles alone recovers the same ladder.

The theory literature reviewed in Section 2 guides the remainder of the validation tests

[13] In Figure 3, as well as all other figures in which we plot layer-by-layer statistics for groups of firms with the same number of layers, we consider firms whose number of layers is between 4 and 14. This range covers about 95 percent of the firms in our sample; the interquartile range of the number of layers is 5 to 8.

that are based on firm size, the shape of the hierarchy, and the distribution of managers across layers.

6.2.2 Firm size

Most models of corporate hierarchies predict that, as a firm grows in scale, the number of layers increases (e.g., Calvo and Wellisz, 1979; Caliendo, Monte, and Rossi-Hansberg, 2015). Figure 4 shows how the distribution of hierarchical layers varies across size quintiles. This figure focuses on the manufacturing sector in 2024 for ease of exposition, but similar patterns are found across all other sectors. As shown, we find evidence for a strong, positive relationship between size (captured by total assets in Panel A and the number of employees in Panel B) and the number of hierarchical layers.

Table 4 further investigates the relationship between firm size and the number of layers. We characterize firm size in eight different ways, using total assets, employment, revenues, the number of business segments, the number of U.S. subsidiaries, board size, firm age, and the measures of Ewens, Peters, and Wang (2025) that capture a firm's stock of organizational and knowledge capital. The dependent variable is the log of the number of layers in year t, while all right-hand variables are measured in year $t-1$. As shown in columns 1-8, all size proxies are positively correlated with a firm's layer count, with R^2 ranging from 16% to 22%. The fit is highest for employment, consistent with the theoretical prediction that larger workforces require additional layers of supervision and coordination. Importantly, the regression in column 9, which includes all size controls, has an R^2 of 24%, so size does not fully explain the variation in hierarchical layers. The remaining variation, which cannot be explained by size alone, is what we exploit in our empirical analysis.

In sum, the evidence from Table 4 confirms the strong association between the number of layers and firm size, as one would expect from theory. This strong association motivates the inclusion of size controls in our regressions to ensure that our results do not merely capture differences in size that are correlated with the number of layers. In our preferred specification, we include granular size controls in the form of asset decile fixed effects (see Section 7).

6.2.3 Firm shape

Next, we turn to the shape of the firm's organizational structure. Most theories predict that the corporate structure will be pyramidal, that is, higher layers will have a smaller share of

employees. For example, this follows from the models of information hierarchy where hierarchies divide workers into those specializing in tasks and those coordinating them (e.g., Bolton and Dewatripont, 1994), the models of knowledge hierarchy where hierarchies distribute problem solvers across layers (e.g., Garicano, 2000), and the models of incentive hierarchy where promotion incentives rely on having fewer employees at higher layers (e.g., Lazear and Rosen, 1981).[14]

Note that the pyramidal shape is not imposed when constructing our measure of hierarchical layers. Thus, evidence that firms are pyramidal on average would provide additional validation in light of the different theories of corporate hierarchy. We investigate firm shape using the share of a firm-year's LinkedIn employees in each layer, counting each worker in the layer of the role held in that year. As discussed in Section 5.1, LinkedIn does not capture all employees. While this incomplete coverage is unlikely to materially affect our estimate of the number of layers—as we discuss in Section 6.4—it does impact employee shares. Thus, estimating firm shapes entails some measurement error.

In Figure 5, we report the average employment share in each layer across all firms in the same industry with the same number of layers (e.g., the first drawn cell in the top row refers to all manufacturing firms that have five layers; the four-layer cell in that row is empty because fewer than ten manufacturing firms have four layers). As can be seen, the evidence is generally consistent with a pyramidal shape. In Figure 6, we repeat this analysis across all firms (regardless of their industry). The evidence is again consistent with a pyramidal structure.

6.2.4 Worker characteristics

The next set of validation tests aims to confirm that worker characteristics across layers vary as one would expect from theory. In particular, the knowledge hierarchy models posit that higher layers are management layers, while lower layers are mostly populated by production workers. In this spirit, we expect to see an increase in the share of managers as we move up the firm's hierarchy. Figure 7 reports the share of workers in each layer with either the string "manager" or signals of an executive role in their title. Each line groups firms that have the same number of layers. Conditioning on the number of layers ensures we are comparing the same layer across firms. As can be seen, the share of managers or executives

[14]Relatedly, Hart and Moore (2005) show that pyramidal hierarchies can be optimal as a result of the trade-off between coordination and specialization.

increases as we move up the hierarchy, flattening only at the top of the highest-layer firms. This higher prevalence of managers—or "problem solvers" in Garicano (2000)'s terminology—confirms that our measure captures meaningful features of the firm's organizational structure, in keeping with existing models of corporate hierarchy.

6.2.5 Summary

In sum, the validation tests provided in this section indicate that publicly available resume data can be used to construct meaningful proxies for corporate hierarchies that are consistent with a large body of theoretical work on organizations. Before moving to the empirical analysis, we provide additional validation addressing title measurement issues (Section 6.3) and the representativeness of the LinkedIn data (Section 6.4).

6.3 Title measurement issues

As mentioned in Section 4, our hierarchy measure requires the identification of unique roles and promotion paths within the firm. The LinkedIn resume data provide imperfect information on both, which could bias our measure of hierarchy, as the firms' specific titles and promotion paths may not align with the economic mechanisms we consider. In particular, issues such as title inflation documented in Cohen, Gurun, and Ozel (2026) could produce too many managers or promotions that are not connected to decision-making or job tasks. This could, in turn, bias upward the number of layers.

Nevertheless, title inflation is unlikely to significantly impact the estimated number of hierarchical layers. This is because the algorithm creates a new layer only when there are substantial outbound promotions from a specific title to higher-level roles, and "fake" or superfluous titles—often found at lower echelons (e.g., "Director of Customer Service" or "Assistant Manager")—typically lack such promotional pathways. The title-to-role algorithm, designed to reduce dimensionality, tends to aggregate these less influential titles rather than separating them into distinct layers.

The evidence presented so far also indicates that title inflation is unlikely to be a first-order concern for our analysis. In particular, title inflation would predict an increase in the number of layers over time, while Figure 1 shows a flat median and an average that rises by only +0.16 layers over 2015–2024. Moreover, title inflation would predict a surge of (alleged) managerial roles at lower levels of the hierarchy, which is contrary to what we observed in

Figure 7.

6.4 Representativeness of the LinkedIn data

A common concern with using LinkedIn data is that coverage of workers is incomplete and potentially non-random, with some roles or parts of the hierarchy potentially overrepresented relative to others. Such selection could affect our estimates and results. Because we lack a complete accounting of employees and titles within firms, we instead benchmark the distribution of occupations in our LinkedIn data against a government database that reports employment by occupation and narrow industry. Specifically, we use the BLS Occupational Employment and Wage Statistics (OEWS) database, which provides industry-level employment counts for U.S. private-sector firms by occupation.[15]

The BLS categorizes firms by industry using NAICS codes and occupations using the Standard Occupational Classification (SOC). Each SOC code has an employee count aggregated across firms, an O*NET Job Zone (i.e., skill index), and a wage estimate. We merge these industry-occupation data with the LinkedIn data by matching each raw LinkedIn title to an SOC code using the O*NET Alternate Titles lexicon.[16] The LinkedIn data are the main estimation sample and include every worker, regardless of whether they are in the firm's largest internal labor market. We match 69.8% of LinkedIn worker-years to a BLS occupation. A large fraction of those that fail to merge are generic rank or seniority titles (e.g., "Vice President") that describe the level rather than the occupation.

With these data, we ask four questions:

1. Of the occupations that the BLS reports for an industry, which are never observed in LinkedIn, and how do these occupations compare in terms of pay and skill?

2. Among occupations matched between BLS and LinkedIn, how does the employment share reported by the BLS compare with the corresponding share in LinkedIn?

[15]The benchmark is the May 2024 OEWS release (`https://www.bls.gov/oes/tables.htm`), which reports, for every industry, the employment count and employment share of each detailed occupation (a six-digit SOC code, e.g., 15-1252 Software Developers). We keep U.S. national rows only, private-sector employment only (government is excluded because our Compustat firms are all private-sector), and the detailed SOC level only (about 830 occupations), not the major or minor aggregates. Occupational wages come from the same OEWS file; the O*NET Job Zone skill index (1–5) comes from the O*NET database (`https://www.onetcenter.org/database.html`).

[16]O*NET publishes thousands of real-world title variants for each SOC code (for 15-1252 Software Developers: "Programmer," "Software Engineer," "SDE," etc.). For titles that match several codes, the firm's BLS industry-occupation mix breaks the tie. A small set of manual and model-proposed rules fills the remaining gaps, with every title resolved to a code through O*NET.

3. Within a firm, where do the missing occupations lie in the pay and skill distributions, in the hierarchy, and on or off the promotion ladder?

4. Does this imperfect coverage affect the number of layers recovered by the estimator?

For the first question, we find that LinkedIn covers 74.2% of the occupations reported by the BLS, which account for 91.1% of BLS employment in the industry's reported occupations. As shown in Table A3, coverage is even higher in the 10 largest industries. For banks, the BLS lists 157 occupations, of which only two are not observed across all banks in our public-firm LinkedIn data. For software, 120 occupations are reported by the BLS, of which only one is missing. For data processing, no occupation is missing. For restaurants, only four of the 89 reported occupations are missing. So what is missing from LinkedIn?

Total missed employment is 9,714,530 workers. Missingness is defined at the industry-occupation cell, so an occupation can be missed in one industry and covered in another. By size, the largest single missing occupations are construction laborers, first-line supervisors of transportation workers, nursing assistants, teaching assistants, and electrical and electronic assemblers. Measured against the national pay distribution, the missing occupations sit in the middle: 57% of their workers are in the second and third wage quintiles and 8% are in the top quintile. Weighted by BLS employment, missed occupations pay 14% less than covered ones ($59,488 versus $69,047), and their O*NET Job Zone is somewhat lower (2.4 versus 2.6). Thus, the occupations not represented in LinkedIn are drawn from the middle of the pay distribution, and they are modestly lower-paid and lower-skill than the occupations LinkedIn covers.

The second question is whether LinkedIn overrepresents certain parts of the firm, for example, middle management and executives. To answer it, we compare the share of workers reported by the BLS with the share reported on LinkedIn. Shares are computed over the occupations LinkedIn matches, rescaled to sum to one on both sides, so an occupation LinkedIn misses entirely enters the coverage results above but not the share comparison. Weighting by employment, the median BLS worker is in an occupation whose LinkedIn share is 61% below its BLS share. As shown in Table A4, the most underrepresented occupations in LinkedIn include home health and personal care aides, security guards, miscellaneous assemblers and fabricators, and construction laborers. These workers rarely have a LinkedIn profile. There are, of course, overrepresented titles on LinkedIn as well. The most pronounced case is chief executives. Consistent with Cohen, Gurun, and Ozel (2026), there are more managers on

LinkedIn than in the BLS. Measured in worker counts, the most overrepresented occupation is sales managers. One explanation for the over-representation of managerial occupations is that LinkedIn is primarily a recruiting platform. Individuals who seek or have received a promotion are more likely to have a profile, making workers at the top of the hierarchy, or on a promotion ladder, more likely to appear in the data. Because the shares are computed across all LinkedIn workers, under-coverage at the bottom mechanically appears as over-representation at the top.

The third question moves inside the firm. Most occupations that a firm lacks are observed in at least one other firm in its industry, allowing us to use these peer firms to infer where the missing occupations would fall in the firm's own pay and skill distributions and within its own hierarchy. Within firms, the occupations that are missing from LinkedIn are substantially lower-skill and lower-paid than those that are observed: as shown in Appendix Table A5, their O*NET Job Zone is 0.5 to 0.6 points lower on the 1-5 scale (17% to 20% of the sample mean of 3.1), and their BLS pay is 23% to 26% lower. This contrasts with their position in the national distribution: although these occupations are mid-pay nationally, they tend to lie near the bottom of the distribution within these firms because firms in our sample have relatively highly paid workforces. Consistent with this observation, when located within the peer firms that employ them, about half of the missing occupations fall in the bottom wage quintile and the bottom job-skill quintile of the peer firm's workforce.

In the hierarchy itself, the missing occupations tend to sit in the middle. We rank occupations within each industry by the fraction of firms in which they are missing from LinkedIn, restricting the ranking to occupations observed in at least one firm in the industry. The occupations that are most frequently missing are much more likely to lie outside the largest internal labor market, which is the focus of the paper. Among occupations that can be located in a peer firm's largest internal labor market, where ranks are defined, the modal missing occupation falls in the third quintile of the peer firm's ranks; 11% fall in the top quintile and 15% in the bottom quintile. The wage and skill placements above use the peer firm's whole workforce, which is why they are more bottom-heavy than the rank placement. Thus, LinkedIn primarily misses occupations that are either outside the main internal labor market or located in the dense middle of the firm's hierarchy.

The fourth question is whether these coverage gaps affect the estimated number of layers. The concern is that missing roles could correspond to missing layers, particularly when a role is the only node on a promotion path. To assess this possibility, we built a simulator of

the layer estimator that incorporates the coverage gaps documented above.[17] The simulation builds synthetic firms with a known number of layers. It removes whole roles at the rate LinkedIn misses occupations in the median industry (16% of roles), drawn from the middle of the ladder as in Appendix Table A5, thins the junior workforce to the median firm's coverage of 68% of its true headcount, and runs the paper's estimator on the remaining data. At these measured levels of coverage, the estimator recovers the true number of layers in 99.5 percent of simulated firms, and otherwise misses by one layer. Intuitively, the observed workers within a role can be substantially thinned without affecting the estimated number of layers, as long as enough workers remain to reveal the promotion links connecting adjacent roles. Similarly, removing a role should change the estimate only if it breaks the promotion paths needed to identify the layers. The estimated number of layers changes little as coverage deteriorates, with meaningful loss of layers only at extreme levels of missingness—for example, when most mid-ladder roles are deleted or when junior careers are hidden so completely that the bottom role has no observed workers—levels of deterioration that go far beyond the observed coverage gaps in LinkedIn. Note that, lacking a BLS benchmark for promotion records, the simulator treats that channel as a range to explore rather than a measured rate.

To conclude, our simulations—informed by the measured levels of coverage in the LinkedIn data—indicate that the incomplete LinkedIn coverage is unlikely to materially affect our estimate of the number of layers.

7 Corporate hierarchy and firm outcomes

In the previous section, we validated our measure of hierarchical layers in several ways. We now use this measure to examine whether hierarchical layers help explain firm outcomes. To do so, we estimate regressions of the outcome of interest Y_{it} (such as ROA or the number of patents) for firm i in year t on $\log(\#\text{ layers})_{it-1}$, that is, the logarithm of the number of hierarchical layers of firm i in year $t-1$. The regressions include year, industry (based on 4-digit SIC codes), and size decile fixed effects (based on the firm's assets at $t-1$). That is, we compare firms in the same year, same industry, and of similar size, that differ based on their number of layers.[18] When estimating these regressions, we cluster standard errors at

[17] See `https://hierarchy-coverage-sim.onrender.com/`.

[18] We obtain qualitatively similar results throughout if the size decile fixed effects are based on the number of employees. We also obtain similar results if we include age decile fixed effects, where age is measured as the number of years since the firm went public.

the company level.

Our regressions aim to characterize how corporate hierarchies correlate with a series of outcome variables. Accordingly, we caution that our results are correlational and hence need not warrant a causal interpretation. Breaking ground on the latter would require exogenous variation in corporate hierarchy, which is difficult to obtain given the inherent endogeneity of organizational structures.

7.1 Human capital

We first examine how the number of layers correlates with measures that capture different dimensions of the firm's human capital. We consider two dimensions: education and career paths.

In the theories of corporate hierarchies, additional layers require employees with managerial skills that are costly to acquire. For example, in the models of Rosen (1982) and Garicano (2000), this is captured by a fixed cost of knowledge acquisition. We thus expect education to rise with hierarchical layer, as the higher layers are where the managerial skills sit. To examine whether this is the case, we use the educational histories of employees from their online resumes.[19]

Figure 8 presents the percent of workers in each layer who have a post-graduate degree. Each line represents the average percent of workers for firms with the same total number of layers. As can be seen, post-graduate degrees tend to be more prevalent as we move up the firm's hierarchy, consistent with the above arguments.

Next, we examine the relationship between the number of layers and the extent to which workers build their careers within the firm's internal labor market. A large literature highlights the importance of intrafirm labor mobility for employees' incentives and their willingness to invest in firm-specific human capital (e.g. Lazear and Rosen, 1981; Waldman, 1984; Baker, Gibbs, and Holmstrom, 1994; Lazear and Oyer, 2004). Accordingly, the provision of a clear career path ("job ladders") in multi-layer firms is likely to affect employees in multiple ways.

In Table 5, we examine several dimensions related to the firm's human capital. The first two columns show that high-layer firms are less likely to bring in external hires (column 1) and have higher internal promotion rates (column 2). These patterns confirm that the job ladders of promotion are more prevalent in high-layer firms. The higher promotion rates, combined

[19] Since the employees' educational histories are sometimes unreported or incomplete, this analysis is based on a smaller sample.

with the longer tenure observed in columns 3 (referring to the "non-new" employees, that is, excluding employees who joined during the year) and 4 (referring to all employees), show that the average employee in high-layer firms has more experience within the firm, consistent with the notion that employees have stronger incentives to invest in firm-specific human capital in high-layer firms.

In column 5, we further explore the extent to which high-layer firms provide employee stock option plans (ESOPs) to their employees. ESOPs are commonly used as a way to enhance employee retention and long-term commitment to the firm (e.g., Jones and Kato, 1995). To the extent that high-layer firms are more reliant on internal job ladders and firm-specific human capital investments, one might expect a higher prevalence of ESOPs among them. We collect data on ESOPs from the firms' filings of Form 5500 that we match to our sample of public firms.[20] As can be seen, we find that firms with more layers are indeed significantly more likely to offer ESOPs to their employees.

An important part of the firm's internal labor market is the top executive team. In columns 6 to 8 of Table 5, we examine features of the CEO. First, we examine whether high-layer firms are more likely to appoint internal CEOs. We find that the point estimate is positive, in keeping with our finding of higher internal promotion rates and tenure. However, it is not significant at conventional levels, suggesting that the role of internal career ladders may be less pronounced at the very top of the hierarchy. Turning to CEO compensation, we find that CEOs of high-layer firms have higher compensation, although the difference is again not significant at conventional levels. Perhaps more importantly, we find that multi-layer firms have significantly higher wage inequality, as captured by the CEO pay ratio (that is, the ratio of the CEO's pay to the median worker's pay). This finding is consistent with several models of corporate hierarchy. In particular, it is consistent with the tournament-type models, in which higher rewards are provided at the top of the job ladder (e.g., Lazear and Rosen, 1981), and the models of knowledge hierarchy in which higher-level problem-solvers require skills that are increasingly more costly (e.g., Garicano and Rossi-Hansberg, 2006).

[20] Form 5500 filings are annual reports required for employee benefit plans subject to the Employee Retirement Income Security Act (ERISA)—including retirement plans such as 401(k)s, as well as certain welfare benefit plans—in order to provide the Department of Labor (DOL) and the IRS with key information about a plan's financial condition, investments, operations, and compliance with federal regulations. Using these data, we construct an indicator variable ("ESOP") equal to one if the firm has an employee stock option plan or employee stock bonuses.

7.2 Operating performance

Arguably, how firms organize production with workers in their hierarchy should impact firm productivity. Large, persistent differences in firm productivity have been documented across countries and time (see the survey of Syverson, 2011). These differences can be partially explained by market structure, technology, demand, diversification, and exposure to trade, among others. Organizational structure is likely to be another determinant of firm productivity.

Many models highlight how organizational structures can contribute to productivity. For example, in the models of knowledge hierarchy, adding a layer brings in more managers ("problem solvers") who contribute to productivity by solving complex problems faced by the production workers at lower layers. This benefit comes at a cost since managers spend their time problem-solving rather than producing (e.g., Garicano, 2000; Garicano and Van Zandt, 2013; Caliendo, Mion, Opromolla, and Rossi-Hansberg, 2020).

Columns 1-3 in Table 6 explore the relationship between the number of layers and three measures of profitability. As can be seen, a higher number of hierarchical layers is associated with higher profitability: the return on assets (ROA) and net profit margin (NPM) are significant at the 5 percent level or better, and the return on equity (ROE) is positive but not significant at conventional levels. Column 4 further shows that high-layer firms have a significantly lower revenue-to-employee ratio. This echoes the argument that high-layer firms have more managers who are not directly involved in the production process.

Column 5 examines wages (scaled by total assets). The point estimate is positive, indicating that firms with more hierarchical layers have more expensive human capital. This points to the higher cost of maintaining multilayered hierarchies, and aligns with the theories of corporate hierarchies in which additional layers require employees with managerial skills (or the ability to solve more complex problems in the models of knowledge hierarchies) whose human capital is likely more costly. However, we caution that the point estimate is not significant at conventional levels.[21]

7.3 Innovation

A common refrain about large, public firms is that they are at a disadvantage in producing high-impact innovation. The small, high-growth startups are considered critical suppliers of

[21] Note the smaller sample size due to the incomplete coverage of wages in Compustat.

such innovation. This supposed advantage of startups could reflect the startup's size, flexibility, and quick decision-making. For example, Ewens and Marx (2024) show that younger firms produce higher-impact patents. Relatedly, in his analysis of newly-public firms, Bernstein (2015) shows that the firm's innovation productivity decreases after the IPO. Our data allow us to explore how hierarchical structures can explain heterogeneity in public firms' innovation.

Table 7 provides regression estimates of innovation activity and quality. Investments in innovation often start with R&D. Column 1 shows that firms with more layers invest less in R&D (as a ratio of total assets). The Poisson estimate in column 2 shows that this lower R&D investment does not hinder the firm's ability to generate patents. In fact, firms with more layers generate significantly more patents. This finding is confirmed in column 3 where we use OLS (instead of Poisson) and code the dependent variable as the IHS (inverse hyperbolic sine) of the number of patents.[22] Do these patents differ in quality? To examine this question, we use the average number of citations per patent as dependent variable. As can be seen from columns 4 (Poisson) and 5 (OLS with the IHS transform), we find little evidence of an association between hierarchical layers and this proxy for patent quality.[23]

Overall, these innovation outcomes indicate that multi-layer hierarchies are not detrimental to innovation production—in fact, they tend to produce more innovation—and do not face a disadvantage in generating high-quality patents.[24]

[22]The IHS(y) transform is preferred over the $\log(y)$ transform as it behaves similarly to the logarithm but allows to retain zero-valued observations without having to arbitrarily rescale y in the logarithm operator (e.g., by adding one). See Bellemare and Wichman (2020) for more details.

[23]In additional analyses, we considered other proxies for patent quality such as novelty and the number of right-tail patents. We found no significant differences in these measures between firms with different hierarchical layers.

[24]In addition to the analysis of innovation, we also examine another type of investment, namely acquisitions, in Appendix Table A6, using SDC Platinum and the Compustat-SDC merge from Ewens, Peters, and Wang (2025). Theoretically, the relation between hierarchical layers and acquisition activity is ambiguous. On the one hand, to the extent that high-layer firms have more specialized employees—as in the models of knowledge and information-based hierarchies—they might be better able to identify suitable targets and potential synergies. On the other hand, the multi-layer structure may complicate the integration of other firms and their workforce. Using various metrics of acquisition activity, we find weak evidence that high-layer firms engage in more acquisitions, the estimates being positive and marginally significant for the number of acquisitions (column 3) and the total valuation of acquisitions (column 4). Finally, Appendix Table A7 examines layers following acquisitions. We find that firms with more acquisitions (column 3), higher valuation of acquisitions (column 4), or more acquisitions of public targets (column 5) in a given year have significantly more layers in the following year, consistent with acquisitions increasing organizational complexity and companies responding by adding layers.

7.4 Firm risk

While the theories of corporate hierarchy do not provide explicit predictions for firm risk, there are many reasons to expect hierarchies to matter along this dimension. In particular, in the models of knowledge hierarchy, the addition of layers increases the firms' ability to "solve problems," which may enhance their ability to manage risks. This, in turn, would reduce the volatility of the firms' cash flows.

In columns 1 and 2 of Table 8, we examine the relationship between the number of layers and the firm's stock volatility, measured as the square root of the sum of squared daily stock returns (from CRSP) over the year. Column 1 refers to the raw measure, while it is winsorized at the 5% level in column 2. As can be seen, both estimates indicate that more layers are associated with lower stock volatility.

In columns 3 to 6, we refine this analysis by asking whether the lower stock volatility reflects lower business and/or financial risk. In columns 3 and 4, we consider measures of operating asset volatility and cash flow volatility, respectively, which we compute as in Gormley and Matsa (2016). As is shown, firms with more layers have lower business risk based on these two metrics.[25] In columns 5 and 6, we examine leverage (the debt-to-asset ratio) and cash holdings (the cash-to-asset ratio) and find that the number of layers is not associated with more conservative financing policies. If at all, the evidence in column 6 suggests that high-layer firms hold less cash (corresponding to higher financial risk). Overall, these findings indicate that the lower stock volatility of high-layer firms is due to lower business risk rather than financial risk.

7.5 Resilience to shocks

A potential downside of high-layer firms is that they may have less flexibility to quickly adapt to disruptions in their operating environment. This argument aligns with models that emphasize the role of hierarchies in communicating information within the firm (e.g., Radner, 1993; Bolton and Dewatripont, 1994). In these models, flatter hierarchies facilitate more direct communication with higher layers of the organization, potentially accelerating decision-making and enhancing corporate agility in response to external disruptions.

[25] Note that the point estimate for cash flow volatility is not significant at conventional levels. However, since this metric is based on the standard deviation of only four quarterly cash flows (from Compustat's quarterly file), it is inherently noisy, which limits our ability to detect a statistically significant relationship, even if one exists.

To examine whether firms with flatter hierarchies are more resilient to external disruptions, we consider the COVID-19 crisis as an exogenous shock to firms' operating environment. Specifically, we follow Fahlenbrach, Rageth, and Stulz (2021), who measure the extent to which firms were affected by the pandemic using their stock price performance during the "collapse period" of COVID-19 from February 3 to March 23, 2020. We compute daily excess returns as the logarithm of one plus the stock's total return (from CRSP) minus the risk-free rate, measured using the 1-month daily Treasury-bill rate (from the Federal Reserve Bank of St. Louis). We then cumulate these excess returns over the collapse period.[26] In our sample, the average cumulative excess return is -36%, comparable to the overall decrease of the S&P 500 during the collapse period.

We then regress cumulative excess returns on the number of layers and, in keeping with our baseline specification, fixed effects for size decile and industry. The results are reported in column 1 of Table 9. The coefficient on the number of layers is negative, indicating that firms with more layers experienced larger stock price declines, although the estimate is not statistically significant at conventional levels.

In column 3, we distinguish between firms in "COVID industries" and those in other industries, using the classification of Fahlenbrach, Rageth, and Stulz (2021). COVID industries are those that were particularly affected by the pandemic, such as restaurants, hospitality, and transportation. We find that high-layer firms in COVID industries experienced significantly larger stock price declines. The point estimate of -0.073 implies that a 10% higher number of layers is associated with a larger stock price decline by about 0.7 percentage points. This finding is consistent with the notion that flatter hierarchies enhance firms' agility in responding to external disruptions such as the COVID shock.

In columns 2 and 4, we repeat the analysis controlling for financing flexibility, measured by leverage and cash holdings. The coefficients on these controls are consistent with the findings of Fahlenbrach, Rageth, and Stulz (2021): firms with higher debt capacity (lower leverage) experienced significantly smaller stock price declines during COVID-19, and the coefficient on cash holdings is positive, though not significant at conventional levels. Importantly, the coefficient on the number of layers remains negative and statistically significant among firms in COVID industries. This indicates that financing flexibility and organizational flexibility represent distinct dimensions of firms' resilience to external shocks.

[26] We exclude financial firms (SIC 6000-6999) and utilities (SIC 4900-4999), as well as stocks with an average price of less than \$1. In addition, we winsorize cumulative excess returns at the 5th and 95th percentiles to mitigate the influence of outliers.

8 What explains firms' hierarchical layers

The analysis presented so far indicates that corporate hierarchy—as measured by the number of layers—has predictive power for a series of outcomes in a way that is consistent with various theories of firm organization. That being said, the results presented so far tell us little about *how* firms choose their layers. In what follows, we turn to this question by exploring potential determinants of the firms' decision to add or remove hierarchical layers.

8.1 Product market competition

In Section 7.5, we showed that firms with flatter hierarchies were more resilient to the COVID-19 shock. To the extent that flatter hierarchies enhance firms' resilience to shocks, we would expect firms facing greater exposure to potential disruptions to adopt flatter hierarchies.

To examine this question, we study the relationship between product market threats and the number of organizational layers. To capture product market threats, we use the product market fluidity measure of Hoberg, Phillips, and Prabhala (2014), which captures the extent to which a firm's product market is subject to disruptions from competitors.[27] We then regress the number of layers on the fluidity measure, size decile fixed effects, and year fixed effects. The results are reported in column 5 of Table 9. The specification in column 6 additionally includes industry fixed effects and therefore relies on within-industry variation.

As can be seen, the coefficient on the fluidity measure is negative and statistically significant in both specifications. Thus, firms facing a higher rate of product market disruption tend to have flatter hierarchies. This finding complements our COVID-19 results in Section 7.5, suggesting that firms exposed to greater disruption may adopt flatter organizational structures to enhance their agility in responding to changes in their operating environment.[28]

[27] Hoberg, Phillips, and Prabhala (2014) measure product market fluidity as the cosine similarity between a firm's current product description vocabulary and the aggregate change in other firms' product description vocabularies from the previous year. Higher fluidity therefore indicates that rivals are changing their product offerings in ways that increasingly overlap with the firm's product space, implying greater potential competitive threats.

[28] This finding is also consistent with the earlier evidence in Guadalupe and Wulf (2010). Using confidential data from Hewitt Associates on the upper layers of about 300 U.S. firms from 1986-1999, they find that greater competition from abroad led firms to reduce the number of management levels and broaden the CEO's span of control, consistent with our finding of flatter hierarchies.

8.2 Corporate life cycle

Next, we examine how firms adjust their hierarchies over their life cycle, focusing on the transition to becoming a public firm in the IPO.

This analysis echoes the work of Bias, Lochner, Obernberger, and Sevilir (2026), who explore the shift in hierarchization among German IPOs using the Rajan (2012) framework. We examine a similar question for U.S. public firms whose IPOs occurred during our sample period and for which we can observe worker resumes in the years prior. As the transition to being public increases the complexity of the firm's operations and imposes additional reporting requirements, we expect firms to increase their hierarchical layers. Firms may also expand their layers if going public increases the size of the organization, in keeping with the strong relationship between firm size and the count of layers documented in Section 6.2.2.

Figure 9 presents regression coefficient estimates from a simple event study analysis for firms that had their IPO during our sample period. Consistent with the above arguments, we find that the number of layers increases substantially post-IPO. Note that the number of layers already increased two years before the IPO, likely because the firm was preparing for the listing. Because the decision to go public is endogenous, we caution that these results are correlational and hence need not warrant a causal interpretation.[29]

8.3 Artificial intelligence

Artificial intelligence (AI) reduces the cost of knowledge acquisition and information processing. From a theoretical perspective, we would expect hierarchies to flatten following the adoption of AI. For example, in the models of knowledge hierarchy, lowering the cost of knowledge acquisition allows workers to solve a wider range of problems, which reduces the demand for problem solvers and hence the need for hierarchical layers (e.g., Garicano, 2000; Garicano and Van Zandt, 2013). Similarly, in the information-based model of hierarchies, the better ability to process information reduces the need for delegation through additional layers (e.g., Radner, 1993; Bolton and Dewatripont, 1994).

[29] To gain additional insight into how firms' hierarchical structures vary over their life cycle, we use the data of Hoberg and Maksimovic (2022), which characterize firms according to four stages of their product life cycle: product development, process optimization, mature product, and product decline. For each firm-year, they construct a weight for each stage. In Appendix Table A8, we regress the log number of layers on these stage weights, combining mature product and product decline into a single weight and using product development as the omitted category. As shown, the number of layers is significantly lower for firms with a greater weight on product development. This pattern is consistent with flatter hierarchies being associated with the need for quicker decision-making and greater organizational agility during product development.

The recent rise in the adoption of AI provides an opportunity to evaluate these predictions. To do so, we regress the log number of hierarchical layers on the measures of AI adoption proposed by Babina, Fedyk, He, and Hodson (2024) based on the number of AI job postings. The regressions also include firm, year, and size decile fixed effects. The inclusion of firm fixed effects is important since we are interested in how firms adjust their hierarchies upon adopting AI. As discussed in Section 6.1.1, a limitation is that the firm fixed effects reduce the power of our tests due to the limited within-firm variation in hierarchical layers.

The results are presented in Table 10. We consider four different metrics of AI adoption: the share of AI job postings, broadly (column 1) and narrowly (column 2) defined, the average share of AI skills (column 3), and the average maximum share of AI skills (column 4). As can be seen, we find evidence for a significant reduction in the number of layers in all four columns. These results are consistent with the organizational mechanisms emphasized in the theory literature: by lowering the cost of knowledge acquisition and information processing, AI reduces the need for hierarchical layers, inducing firms to flatten their organizational structures.[30]

9 Conclusion

We propose a measure of corporate hierarchy for U.S. public firms that is obtained from online resumes of 18.1 million employees and a network estimation technique that allows us to identify hierarchical layers. After validating this measure, we document several facts pertaining to corporate hierarchies that we link to existing theories of firm organization. In particular, we show that more hierarchical firms have higher internal promotion rates and longer employee tenure, and that education rises with the hierarchical layer within the firm. Their operating performance is higher, and they produce more patents, but not higher-quality patents. They exhibit lower stock return volatility and lower business risk (i.e., operating asset volatility). At the same time, we find that firms with flatter hierarchies experienced a lower stock market drop during the COVID-19 crisis, consistent with greater organizational agility. Turning to potential determinants of corporate hierarchies, we find that firms facing greater product market disruptions have flatter hierarchies, which points again toward flatter hierarchies providing greater agility in response to external disruptions. We further find that firms add hierarchical

[30]The evidence in Table 10 is related to the finding of Babina, Fedyk, He, and Hodson (2025) that firms with more AI postings tend to have a higher share of employees in junior vs. senior positions. To the extent that flatter hierarchies have a higher share of junior positions, their finding is consistent with our evidence.

layers after going public, while they flatten their hierarchies following the adoption of artificial intelligence. These findings suggest that firms use hierarchical layers to expand their problem-solving capacity when organizational complexity increases, and reduce them when technology lowers the cost of acquiring and processing knowledge.

In their characterization of the literature on hierarchies and internal labor markets, Baker and Holmstrom (1995) noted that the literature has "too many theories, too few facts" (p. 255). What was true in 1995 remains true 30 years later, especially for U.S. public firms, as there is no available data on their hierarchy. This study aimed to fill this gap, by providing a measure of hierarchy that can be used to gain insights into the determinants and implications of hierarchies. In this spirit, we will release the hierarchy data to other researchers, in the hope that they will spur future work on this topic.[31]

Our paper leaves several avenues open for future research, of which we highlight two. First, our findings suggest that the optimal number of hierarchical layers depends on the nature of the problems firms face, the speed at which their operating environments change, and the technologies available to process and communicate information. The ability to observe hierarchies at the firm level opens the door to studying these forces in greater detail, including how organizational structures evolve in response to other technological innovations and shifts in firms' strategies. Our data may also enable future work to examine how hierarchy shapes the allocation of talent and the diffusion of information within firms. Second, our measure provides a starting point for future work to refine and expand the measurement of corporate hierarchies. In particular, future work could develop measures that capture other dimensions of organizational structure beyond the number of layers. Such work could help distinguish more precisely between different organizational forms that have the same number of layers and provide a richer picture of how firms are organized.

[31] The data will be posted on a public repository. For now, researchers interested in using the data should email the authors to obtain the most updated version of the data.

References

Adhvaryu, Achyuta, Vittorio Bassi, Anant Nyshadham, Jorge A. Tamayo, and Nicolas Torres, 2023, Organizational responses to product cycles, NBER Working Paper 31582.

Alchian, Armen A., and Harold Demsetz, 1972, Production, information costs, and economic organization, *American Economic Review* 62, 777–795.

Arora, Ashish, Sharon Belenzon, and Lia Sheer, 2021, Matching patents to Compustat firms, 1980–2015: Dynamic reassignment, name changes, and ownership structures, *Research Policy* 50, 104217.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151, 103745.

Babina, Tania, Anastassia Fedyk, Alex X. He, and James Hodson, 2025, Firm investments in artificial intelligence technologies and changes in workforce composition, in Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich Strassner, ed.: *NBER Volume on Technology, Productivity, and Economic Growth* (University of Chicago Press) 75–117.

Baker, George, Michael Gibbs, and Bengt Holmstrom, 1994, The internal economics of the firm: Evidence from personnel data, *Quarterly Journal of Economics* 109, 881–919.

Baker, George, and Bengt Holmstrom, 1995, Internal labor markets: Too many theories, too few facts, *American Economic Review* 85, 255–259.

Bellemare, Marc F., and Casey J. Wichman, 2020, Elasticities and the inverse hyperbolic sine transformation, *Oxford Bulletin of Economics and Statistics* 82, 50–61.

Bernstein, Shai, 2015, Does going public affect innovation?, *Journal of Finance* 70, 1365–1403.

Bias, Daniel, Benjamin Lochner, Stefan Obernberger, and Merih Sevilir, 2026, Going public and the internal organization of the firm, *Journal of Finance* 81, 459–505.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen, 2012, The organization of firms across countries, *Quarterly Journal of Economics* 127, 1663–1705.

Bolton, Patrick, and Mathias Dewatripont, 1994, The firm as a communication network, *Quarterly Journal of Economics* 109, 809–839.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa, 2022, Discretizing unobserved heterogeneity, *Econometrica* 90, 625–643.

Caliendo, Lorenzo, Giordano Mion, Luca David Opromolla, and Esteban Rossi-Hansberg, 2020, Productivity and organization in Portuguese firms, *Journal of Political Economy* 128, 4211–4257.

Caliendo, Lorenzo, Ferdinando Monte, and Esteban Rossi-Hansberg, 2015, The anatomy of French production hierarchies, *Journal of Political Economy* 123, 809–852.

Caliendo, Lorenzo, and Esteban Rossi-Hansberg, 2012, The impact of trade on organization and productivity, *Quarterly Journal of Economics* 127, 1393–1467.

Calvo, Guillermo A., and Stanislaw Wellisz, 1978, Supervision, loss of control, and the optimum size of the firm, *Journal of Political Economy* 86, 943–952.

——— , 1979, Hierarchy, ability, and income distribution, *Journal of Political Economy* 87, 991–1010.

Clauset, Aaron, Samuel Arbesman, and Daniel B. Larremore, 2015, Systematic inequality and hierarchy in faculty hiring networks, *Science Advances* 1, e1400005.

Cohen, Lauren, Umit Gurun, and N. Bugra Ozel, 2026, Too many managers: The strategic use of titles to avoid overtime payments, *Review of Financial Studies* (forthcoming).

Ewens, Michael, and Matt Marx, 2024, Firm age and invention: An open-access dataset, Working paper, Columbia Business School.

Ewens, Michael, Ryan H. Peters, and Sean Wang, 2025, Measuring intangible capital with market prices, *Management Science* 71, 407–427.

Fahlenbrach, Rüdiger, Kevin Rageth, and René M. Stulz, 2021, How valuable is financial flexibility when revenue stops? Evidence from the COVID-19 crisis, *Review of Financial Studies* 34, 5474–5521.

Friedrich, Benjamin U., 2022, Trade shocks, firm hierarchies, and wage inequality, *Review of Economics and Statistics* 104, 652–667.

Garicano, Luis, 2000, Hierarchies and the organization of knowledge in production, *Journal of Political Economy* 108, 874–904.

——— , and Thomas N. Hubbard, 2016, The returns to knowledge hierarchies, *Journal of Law, Economics, and Organization* 32, 653–684.

Garicano, Luis, and Esteban Rossi-Hansberg, 2006, Organization and inequality in a knowledge economy, *Quarterly Journal of Economics* 121, 1383–1435.

Garicano, Luis, and Timothy Van Zandt, 2013, Hierarchies and the division of labor, in Robert Gibbons, and John Roberts, ed.: *Handbook of Organizational Economics* (Princeton University Press) 604–654.

Gormley, Todd A., and David A. Matsa, 2016, Playing it safe? Managerial preferences, risk, and agency conflicts, *Journal of Financial Economics* 122, 431–455.

Guadalupe, Maria, Hongyi Li, and Julie Wulf, 2014, Who lives in the C-suite? Organizational structure and the division of labor in top management, *Management Science* 60, 824–844.

Guadalupe, Maria, and Julie Wulf, 2010, The flattening firm and product market competition: The effect of trade liberalization on corporate hierarchies, *American Economic Journal: Applied Economics* 2, 105–127.

Gumpert, Anna, Henrike Steimer, and Manfred Antoni, 2022, Firm organization with multiple establishments, *Quarterly Journal of Economics* 137, 1091–1138.

Hart, Oliver, and John Moore, 2005, On the design of hierarchies: Coordination versus specialization, *Journal of Political Economy* 113, 675–702.

Hoberg, Gerard, and Vojislav Maksimovic, 2022, Product life cycles in corporate finance, *Review of Financial Studies* 35, 4249–4299.

Hoberg, Gerard, Gordon Phillips, and Nagpurnanand Prabhala, 2014, Product market threats, payouts, and financial flexibility, *Journal of Finance* 69, 293–324.

Huitfeldt, Ingrid, Andreas R. Kostøl, Jan Nimczik, and Andrea Weber, 2023, Internal labor markets: A worker flow approach, *Journal of Econometrics* 233, 661–688.

Jones, Derek C., and Takao Kato, 1995, The productivity effects of employee stock-ownership plans and bonuses: Evidence from Japanese panel data, *American Economic Review* 85, 391–414.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665–712.

Lazear, Edward P., and Paul Oyer, 2004, Internal and external labor markets: A personnel economics approach, *Labour Economics* 11, 527–554.

Lazear, Edward P., and Sherwin Rosen, 1981, Rank-order tournaments as optimum labor contracts, *Journal of Political Economy* 89, 841–864.

Liberti, Jose, and Atif Mian, 2009, Estimating the effect of hierarchies on information use, *Review of Financial Studies* 22, 4057–4090.

Malcomson, James, 1984, Work incentives, hierarchy, and internal labor markets, *Journal of Political Economy* 92, 486–507.

Malenko, Nadya, 2024, Information flows, organizational structure, and corporate governance, in David J. Denis, ed.: *Handbook of Corporate Finance* (Edward Elgar Publishing) 511–546.

McKinsey, 2020, Fitter, flatter, faster: How unstructuring your organization can unlock massive value, McKinsey & Company (August 24, 2020).

Qian, Yingyi, 1994, Incentives and loss of control in an optimal hierarchy, *Review of Economic Studies* 61, 527–544.

Radner, Roy, 1993, The organization of decentralized information processing, *Econometrica* 61, 1109–1146.

Rajan, Raghuram G., 2012, Presidential address: The corporation in finance, *Journal of Finance* 67, 1173–1217.

——— , and Julie Wulf, 2006, The flattening firm: Evidence from panel data on the changing nature of corporate hierarchies, *Review of Economics and Statistics* 88, 759–773.

Rosen, Sherwin, 1982, Authority, control, and the distribution of earnings, *Bell Journal of Economics* 13, 311–323.

Skrastins, Janis, and Vikrant Vig, 2019, How organizational hierarchy affects information production, *Review of Financial Studies* 32, 564–604.

Stein, Jeremy C., 2002, Information production and capital allocation: Decentralized versus hierarchical firms, *Journal of Finance* 57, 1891–1921.

Syverson, Chad, 2011, What determines productivity?, *Journal of Economic literature* 49, 326–365.

Tag, Joacim, Thomas Astebro, and Peter Thompson, 2016, Hierarchies and entrepreneurship, *European Economic Review* 89, 129–147.

Waldman, Michael, 1984, Job assignments, signalling, and efficiency, *RAND Journal of Economics* 15, 255–267.

Williamson, Oliver E., 1967, Hierarchical control and optimum firm size, *Journal of Political Economy* 75, 123–138.

Figures and Tables

Figure 1: Organizational layers by year

This figure reports the average and median number of hierarchy layers over time for the firms in our sample. The y-axis is truncated (not zero-based) to make the year-to-year variation visible; the change in the average number of layers from the first to the last sample year is reported in the text.

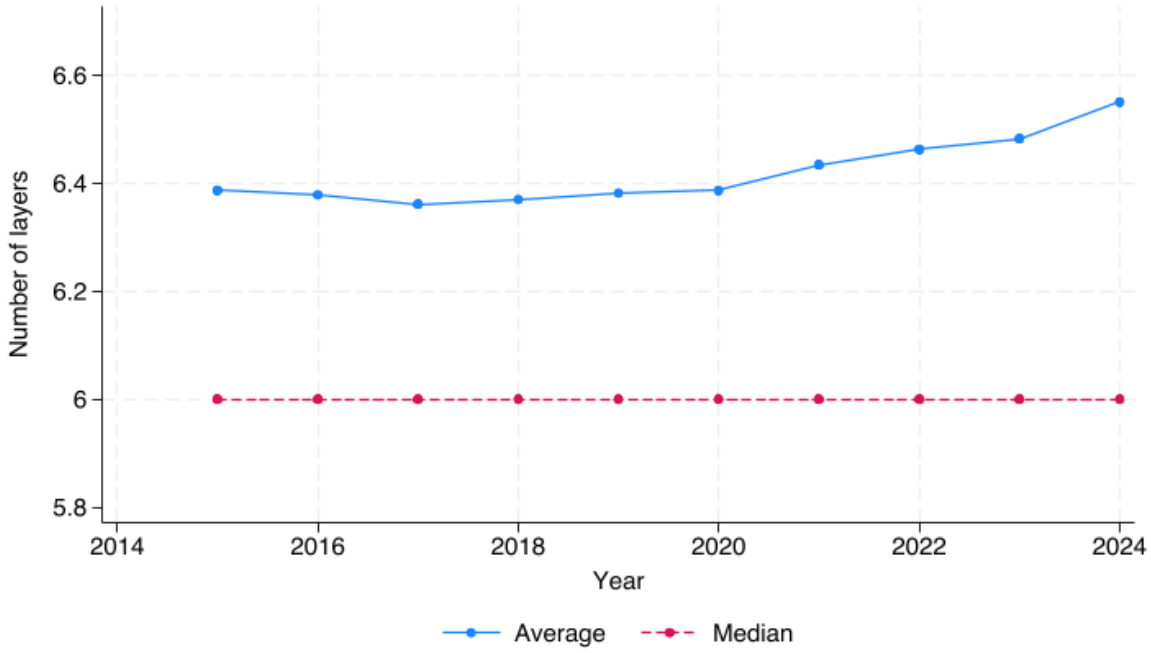
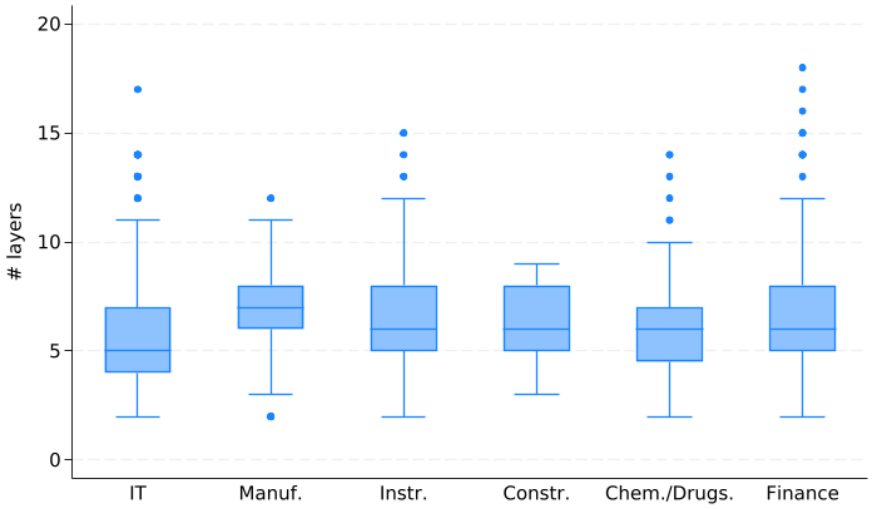


Figure 2: Organizational layers by industry

Panel A displays the box plot of the number of hierarchy layers across firms for the major industries in the sample. For each firm, the number of layers is recorded in the last year in which the firm's corporate hierarchy is observed. Panel B plots the average number of hierarchical layers by industry-year using all firm-year observations.

Panel A: Hierarchical layers by major industry



Panel B: Hierarchical layers by major industry over time

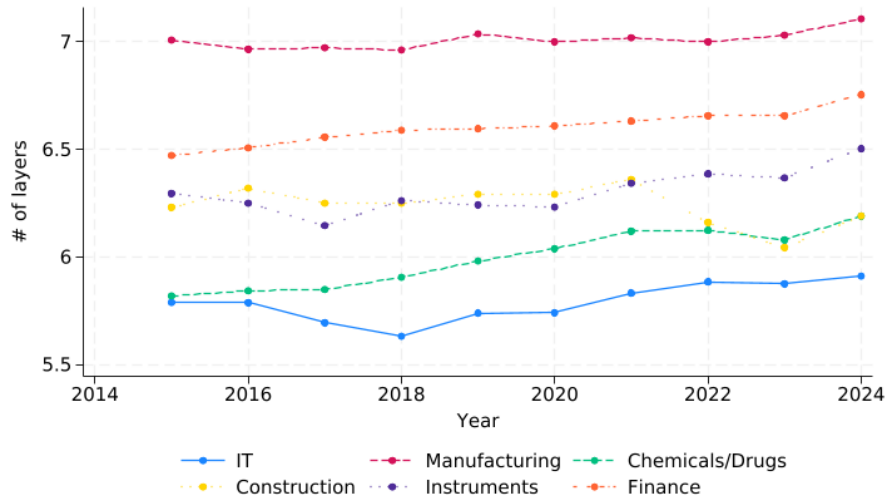


Figure 3: Work experience across the firm hierarchy

This figure reports the average years of work experience (computed as the number of years since college or since the first job) across all jobs within a given layer. Each line represents a group of firms with the same number of layers (i.e., longer lines correspond to groups of firms with more layers). The figure includes firms in our sample whose number of layers is between 4 and 14, which covers about 95 percent of sample firms.

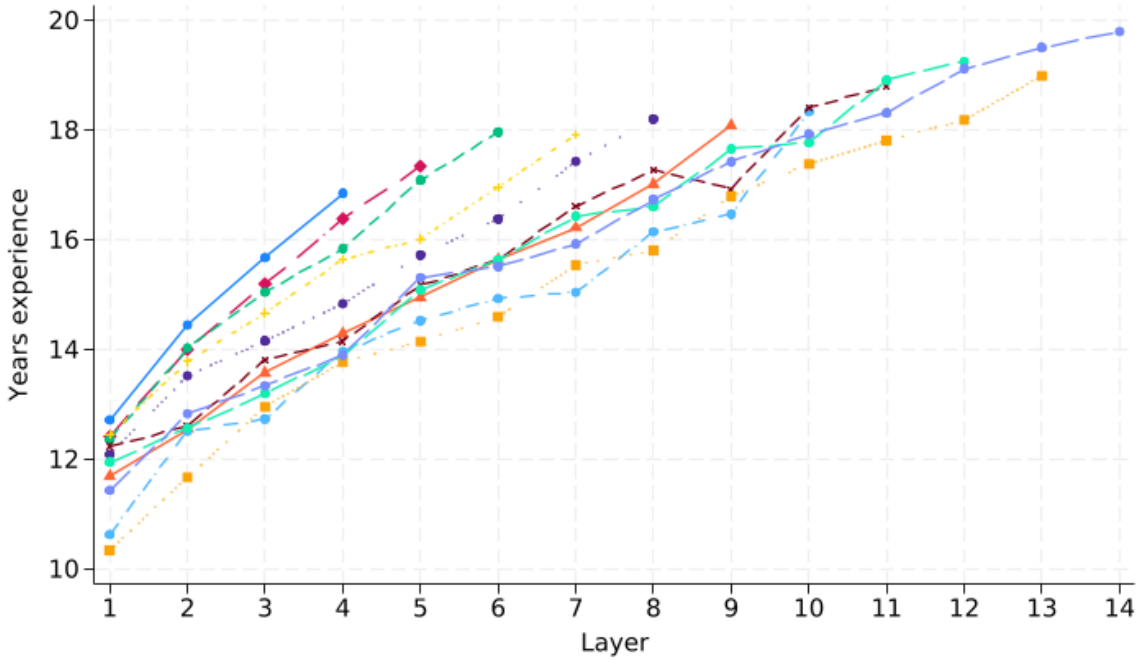
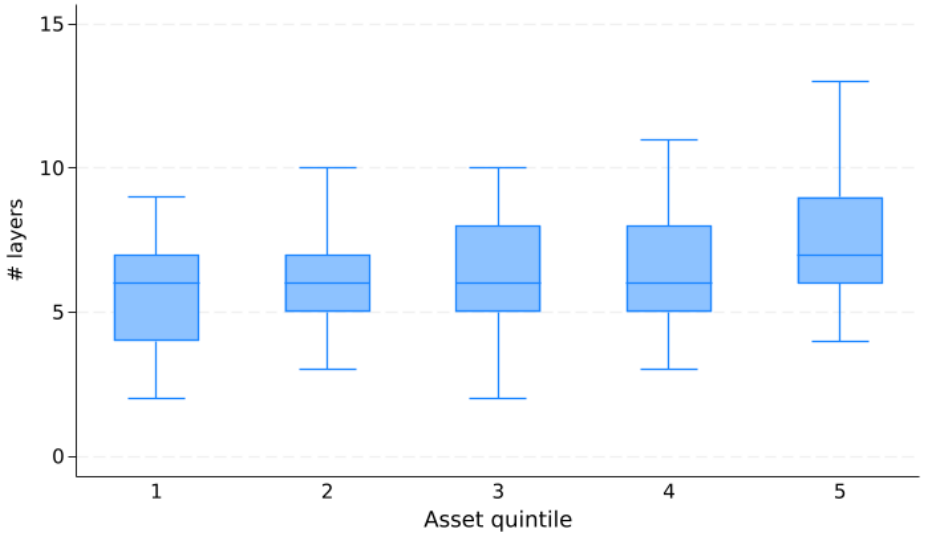


Figure 4: Variation in organizational layers by size

This figure displays the box plot of the number of hierarchical layers by size quintile for the sample of manufacturing firms in 2024. In Panel A, the size quintiles are based on total assets. In Panel B, they are based on the number of employees from Compustat.

Panel A: Assets quintiles



Panel B: Employee count quintiles

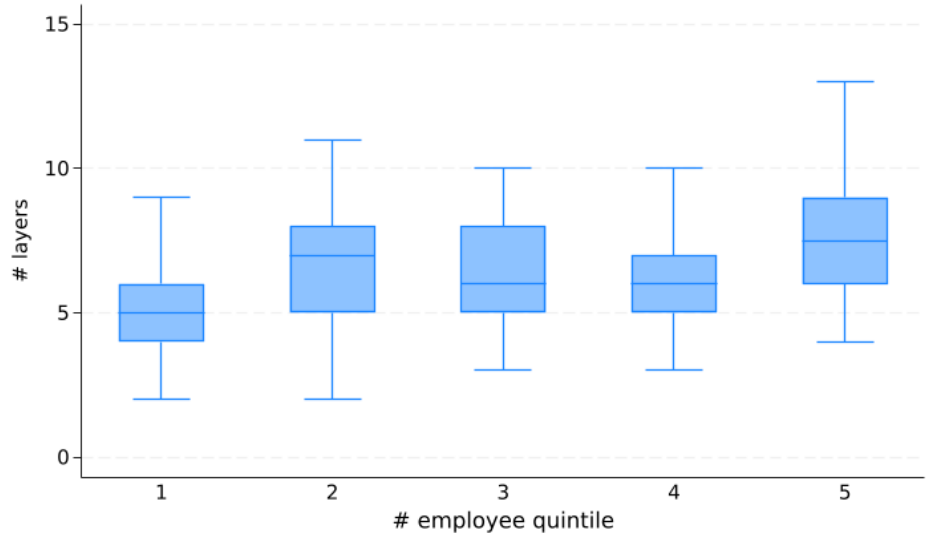


Figure 5: Share of employees by hierarchy layer, by industry and number of layers

This figure shows the shape of the corporate hierarchy as a set of employment-share pyramids, arranged by industry (rows) and by the number of hierarchical layers K (columns). Within each panel, a horizontal bar is drawn for each layer; the width of the bar represents the average share of the firm's employees in that layer, with layer 1 (the bottom layer) at the base and layer K (the top layer) at the apex. The share of employees in a layer is the number of unique employees in the LinkedIn data whose title is associated with that layer, divided by the firm's total number of employees on LinkedIn, averaged across all firm-year observations in the cell. Bar widths share a common scale across all panels. A panel is drawn only for (industry, K) cells containing at least ten distinct firms; the number of firms in the cell is printed above each pyramid (n). The Construction industry is omitted because it has too few firms at any single layer count to form a meaningful pyramid.

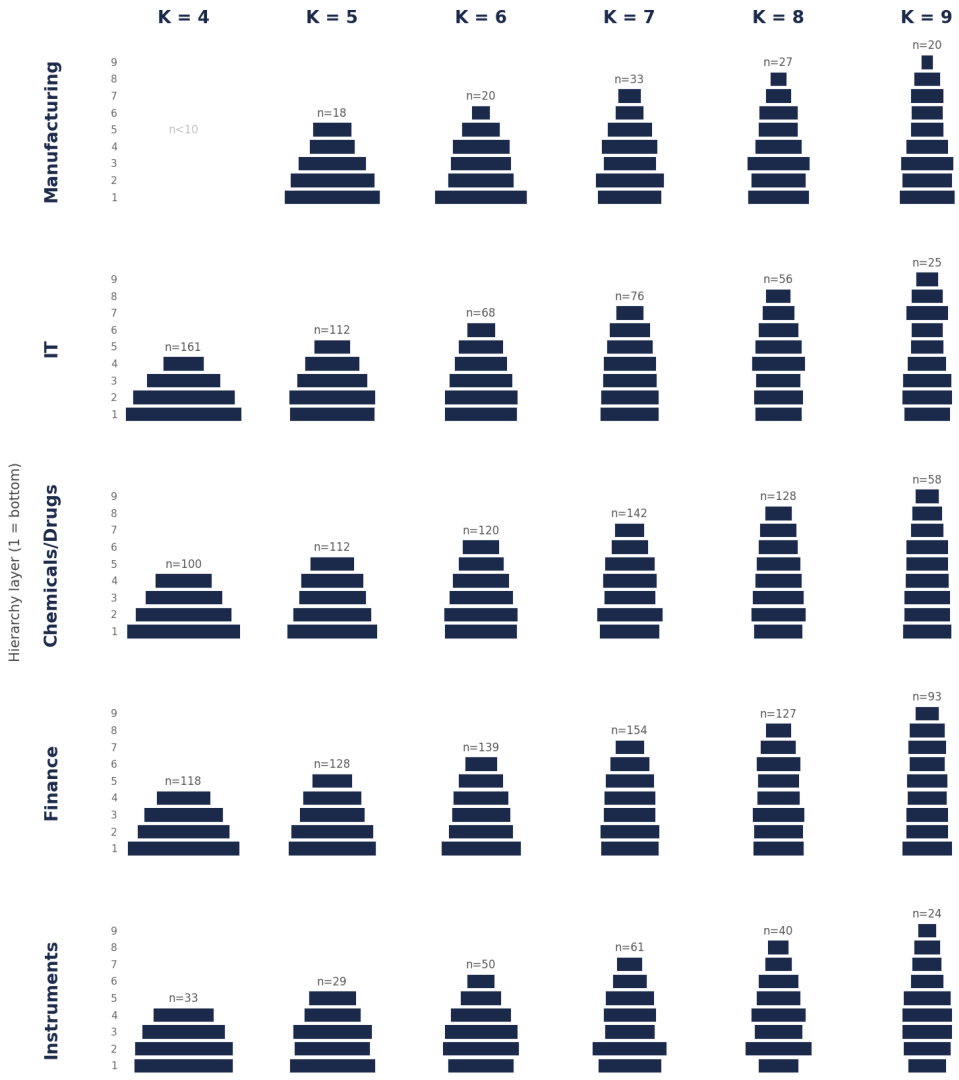
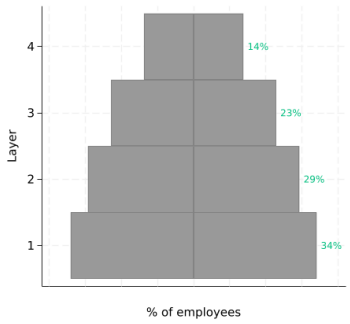


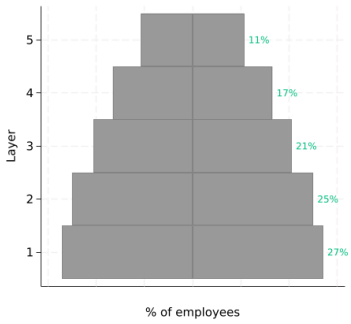
Figure 6: Share of employees by hierarchy layer, by number of layers

This figure reports the average share of employees in each layer for firms with four (Panel A), five (Panel B), six (Panel C), seven (Panel D), eight (Panel E), and nine (Panel F) layers. The share of employees in a layer is the number of unique employees in the LinkedIn data whose title is associated with that layer, divided by the firm's total number of employees on LinkedIn. The figure reports the average across all firm-year observations with the indicated number of layers.

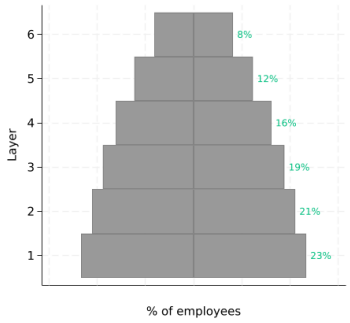
Panel A: Four layers



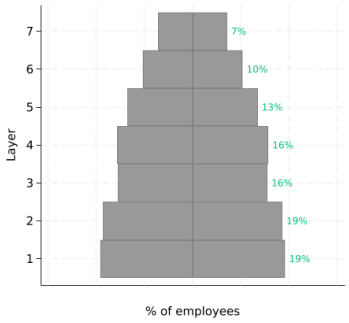
Panel B: Five layers



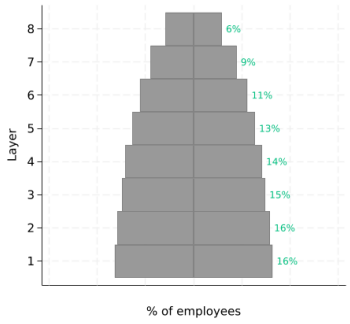
Panel C: Six layers



Panel D: Seven layers



Panel E: Eight layers



Panel F: Nine layers

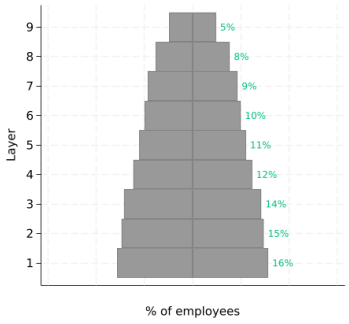


Figure 7: Managerial and executive titles by hierarchical layer

This figure reports the fraction of employees within a given layer whose title is managerial or executive. A title counts as managerial if it contains the string "manager." A title counts as executive if it contains "chief" or a C-suite acronym (*ceo*, *cfo*, *coo*, *cto*, *cio*, *cmo*, *chro*, *cpo*, *clo*, *cro*, *cso*, *cao*), "president" (which also captures vice, senior vice and executive vice president), a vice-president acronym (*vp*, *svp*, *evp*, *rvp*), or "director" (which also captures managing, senior and executive director). Titles containing "assistant," "associate," "deputy" or "staff" are excluded, the last of these so that "chief of staff" does not count; titles containing "owner," "partner" or "founder" are excluded because franchise, agency and professional-services structures are not part of the employment hierarchy. Matching is on whole words. Each line represents a group of firms with the same number of layers (i.e., longer lines correspond to groups of firms with more layers). The figure includes firms whose number of layers is between 4 and 14, which covers about 95 percent of sample firms.

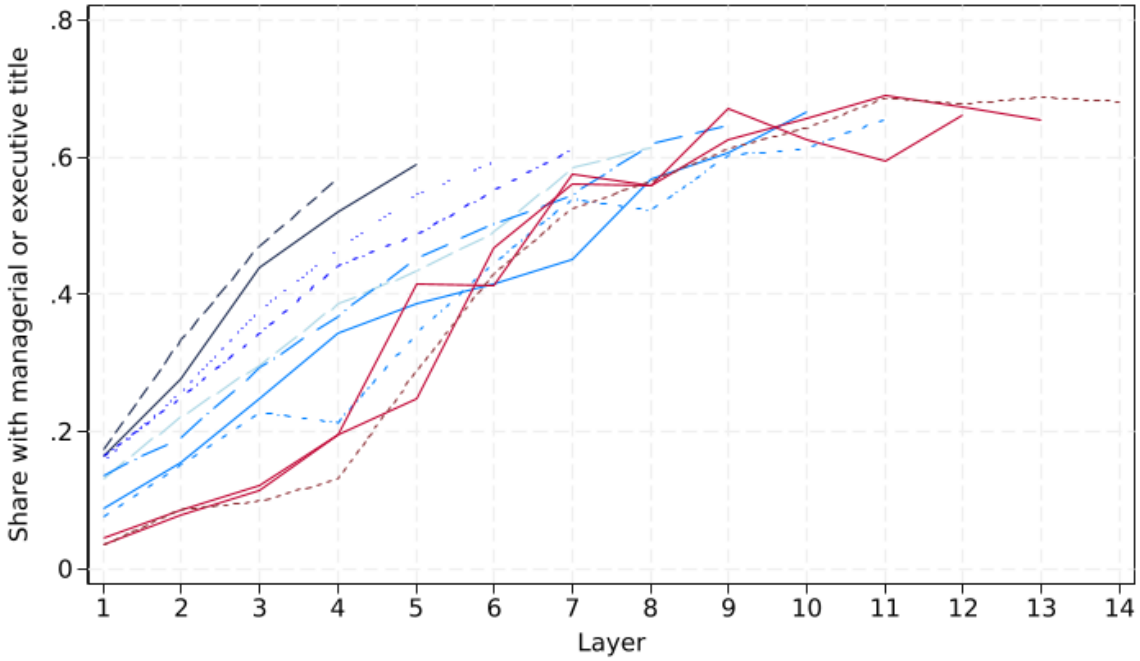


Figure 8: Advanced degrees by hierarchical layer

This figure reports the percentage of workers with an advanced degree (i.e., postgraduates) within a given layer. Each line represents a group of firms with the same number of layers (i.e., longer lines correspond to groups of firms with more layers). The figure includes firms whose number of layers is between 4 and 14, which covers about 95 percent of sample firms.

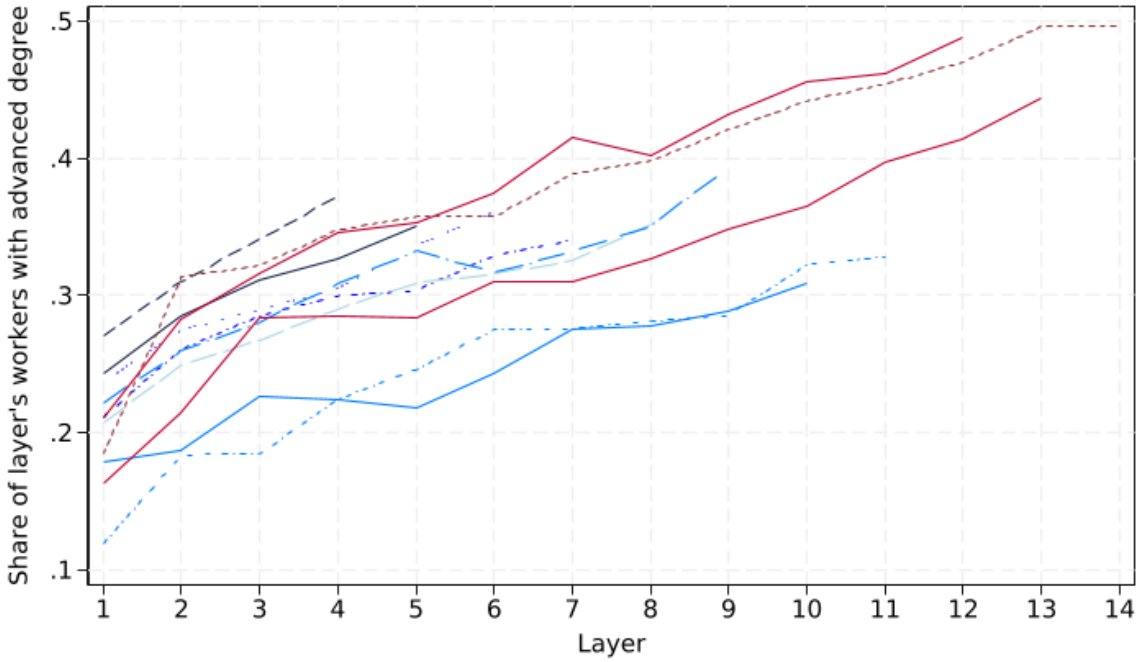


Figure 9: Change in the number of hierarchical layers around the IPO

This figure plots the coefficient estimates $\hat{\beta}_k$, along with their 95% confidence intervals, from the following regression:

$$\log(\#\text{ layers}_{it}) = \alpha + \sum_{k=-3}^{k=4} \beta_k \times \mathbf{1}(k \text{ years since IPO}) + \delta_i + \gamma_t + \epsilon_{it},$$

where i indexes firms and t indexes years for the set of firms that had an IPO during the sample period; $\log(\#\text{ layers}_{it})$ is the natural logarithm of the number of hierarchical layers; $\mathbf{1}(k \text{ years since IPO})$ is an indicator variable equal to one if the observation is recorded k years after (or before when $k < 0$) the firm's IPO; δ_i is a firm fixed effect; and γ_t is a year fixed effect. The excluded year is the year of the IPO. Standard errors are clustered at the firm level.

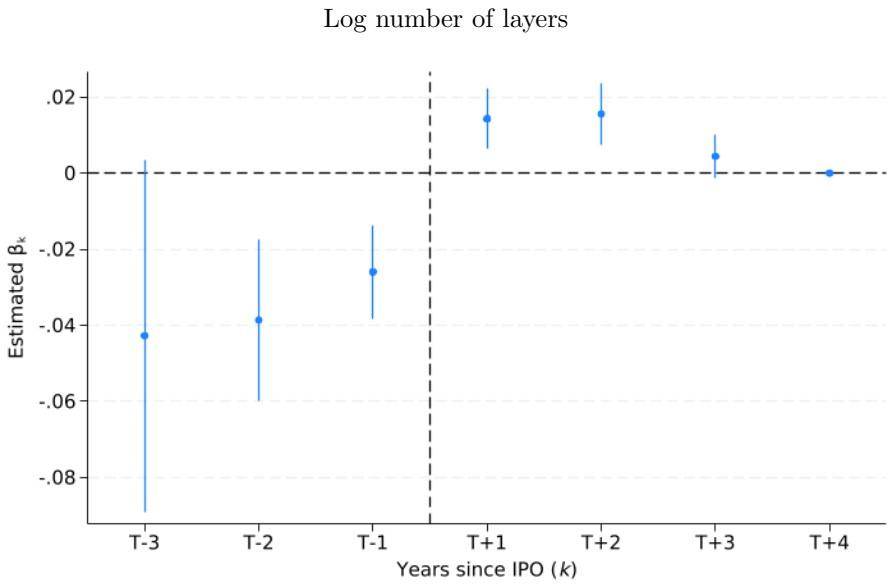


Table 1: Comparison of firms in and out of the sample

This table reports differences in observables between the firms that are in the final sample ("In sample") and those that are not ("Out of sample") within the set of Compustat firms. Each firm is recorded once, in the last fiscal year in which it is observed during our sample period (2015–2024). The table reports the sample mean in each group and the difference-in-means (with the corresponding standard error in parentheses). All variables are from Compustat. "Assets" is the book value of total assets (Compustat item "at"); "Total liabilities" is the sum of all liabilities (item "lt"); "PPE (net)" is the net property, plant, and equipment (item "ppent"), "Debt / equity" is the ratio of total liabilities to total equity (item "teq"); "Cash" is cash and marketable securities (item "che"); "Revenue" is total revenue (item "revt"); "Net income" is the firm's net income (item "ni"); "EBIT" is the firm's earnings before interest and taxes (item "ebit"); and ROA (return on assets) is the ratio of operating income before depreciation (item "oibdp") to assets. All ratios are winsorized at the 5% level. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	In sample	Out of sample	Diff.	s.e.	obs.
Assets	17769.73	3512.39	14257.34***	(2437.01)	7,340
Total liabilities	14578.92	2490.62	12088.31***	(2257.94)	7,330
PPE (net)	1745.09	609.46	1135.63***	(204.14)	7,129
Debt / equity	2.30	1.66	0.65***	(0.09)	7,319
Cash	2021.39	368.38	1653.01***	(371.81)	7,340
Cash / assets	0.22	0.25	-0.04***	(0.01)	7,300
Revenue (m)	5169.15	1556.71	3612.44***	(443.47)	7,286
Net income	112.03	42.22	69.81***	(6.13)	7,287
EBIT	837.03	207.52	629.51***	(105.03)	7,285
Net income / revenue	-0.06	-0.09	0.03***	(0.01)	5,627
ROA	-0.08	-0.38	0.30***	(0.01)	6,988
Firms (Assets obs.)	3,454	3,886			

Table 2: Summary statistics

This table reports summary statistics for the hierarchy variables. "Number of ILMs (incl. single-role titles)" is the number of internal labor markets (ILMs) in the firm, where each ILM is a connected group of roles on the promotion network, before the pruning step described in Appendix A1, including roles that never connect to another role. "Number of ILMs (excl. single-role titles)" counts only ILMs with at least two roles connected by observed worker moves. "Size of largest ILM" is the number of roles in the firm's largest ILM. "Number of unique roles" is the number of unique roles within the firm. "% role in largest ILM" is the fraction of all roles in the firm found in the largest ILM. "% worker-years in largest ILM" is the fraction of all worker-years whose role is in the firm's largest ILM. "Employees (Compustat)" is the number of employees in Compustat (in thousands). "Number employees (LinkedIn)" is the number of employees in the LinkedIn data (in thousands). "Number hierarchy layers" is the number of hierarchical layers of the firm. "Hierarchization" is the metric of hierarchization from Bias, Lochner, Obernberger, and Sevilir (2026), computed using layers based on the occupation codes of the German Classification of Occupations, the occupation codes of O*Net, and our measure of hierarchical layers, respectively. Panel A considers the cross-section of firms using the last year in which the firm's corporate hierarchy is observed. Panel B considers all firm-year observations.

	Panel A: Cross-section					
	Mean	SD	Min	Median	Max	N
Number of ILMs (incl. single-role titles)	103.76	85.10	1.00	83.00	710.00	3,454
Number of ILMs (excl. single-role titles)	8.36	4.08	1.00	8.00	34.00	3,454
Size of largest ILM (roles)	190.53	316.58	1.00	91.00	3900.00	3,454
Number of unique roles	304.14	392.41	2.00	188.00	4623.00	3,454
% role in largest ILM	55.00	20.35	3.77	58.82	100.00	3,454
% worker-years in largest ILM	77.27	20.47	25.00	85.04	100.00	3,454
Employees (Compustat, 1000s)	7.99	37.41	0.00	0.98	1556.00	3,406
Number employees (LinkedIn, '000s)	3.23	11.37	0.00	0.64	325.56	3,454
Number hierarchy layers (selected)	6.35	2.14	2.00	6.00	18.00	3,454
Hierarchization (German)	0.60	0.10	0.00	0.62	1.00	1,091
Hierarchization (ONet)	0.51	0.12	0.00	0.52	0.76	3,454
Hierarchization (this paper)	0.77	0.10	0.14	0.79	0.91	3,326
	Panel B: Firm-year					
	Mean	SD	Min	Median	Max	N
Number of ILMs (incl. single-role titles)	109.21	89.98	1.00	86.00	710.00	26,112
Number of ILMs (excl. single-role titles)	8.48	4.10	1.00	8.00	34.00	26,112
Size of largest ILM (roles)	215.09	346.57	1.00	104.00	3900.00	26,112
Number of unique roles	334.33	427.47	2.00	203.00	4623.00	26,112
% role in largest ILM	58.63	19.50	3.77	62.67	100.00	26,112
% worker-years in largest ILM	79.13	19.71	25.00	87.04	100.00	26,112
Employees (Compustat, 1000s)	9.24	37.62	0.00	1.10	1608.00	24,451
Number employees (LinkedIn, '000s)	3.55	12.14	0.00	0.65	367.66	26,112
Number hierarchy layers (selected)	6.42	2.17	2.00	6.00	18.00	26,112
Hierarchization (German)	0.60	0.11	0.00	0.62	1.00	20,162
Hierarchization (ONet)	0.52	0.12	0.00	0.53	0.77	26,112
Hierarchization (this paper)	0.77	0.09	0.14	0.79	0.91	23,447

Table 3: Common titles by hierarchy layer

This table lists the most common job titles at the bottom and the top of the corporate hierarchy. For each firm in the sample in 2024, employees are assigned to layers using the firm's estimated hierarchy; the bottom layer is layer one and the top layer pools the two highest layers. Titles are ranked by the number of employees holding them, aggregated across all firms. Executive-assistant titles are excluded from the top layer.

Bottom layer			Top layer		
Rank	Job title	Employees	Rank	Job title	Employees
1	sales associate	68,772	1	vice president	21,394
2	barista	63,610	2	director	16,716
3	customer service representative	39,674	3	managing director	13,196
4	crew member	30,195	4	senior vice president	8,067
5	software engineer	24,475	5	financial advisor	6,531
6	cashier	23,952	6	senior director	6,525
7	warehouse associate	22,703	7	associate director	5,760
8	associate	21,023	8	senior manager	5,350
9	shift supervisor	20,012	9	executive director	4,789
10	fulfillment associate	14,467	10	manager	4,420

Table 4: Regressions of the number of hierarchical layers on measures of firm size

This table reports regression estimates using observations at the firm-year level. The dependent variable is the logarithm of the number of hierarchical layers. On the right-hand side, the regressions include various measures of firm size, lagged by one year. "Assets" is the book value of total assets (Compustat item "at"); "employees" is the number of employees in thousands (item "emp"); "revenue" is total revenue (item "revt"); "firm age" is the number of years since the firm's founding; "business segments" is the count of historical business segments reported in the segment file of Compustat; "# US subsidiaries" is the count of subsidiaries from the WRDS subsidiaries data (set to one if missing); "log org. stock" and "log knowledge stock" are the logarithms of the firm's organization and knowledge capital stocks from Ewens, Peters, and Wang (2025), CPI-deflated, entered together in a single column; "board size" is the size of the board from BoardEx. Columns (1)–(6) and (8) each enter one size measure on its own, column (7) enters the two intangible-capital stocks together, and column (9) includes all of them jointly. Industry fixed effects are based on 4-digit SIC industries. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log number of layers								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log assets	0.036***								-0.023*
	(0.0040)								(0.012)
Log employees (CS, 1000s)		0.049***							0.073***
		(0.0044)							(0.012)
Log revenue			0.033***						0.0049
			(0.0038)						(0.0093)
Log firm age (yrs.)				0.031***					-0.0098
				(0.0099)					(0.012)
Log business segments					0.044***				0.0054
					(0.010)				(0.013)
Log (# US subsidiaries)						0.025***			-0.00041
						(0.0037)			(0.0048)
Log org. stock							0.013***		-0.0040
							(0.0028)		(0.0038)
Log knowledge stock							0.0011		-0.010***
							(0.0027)		(0.0035)
Log board size								0.12***	-0.055*
								(0.025)	(0.031)
Observations	22,229	20,960	21,328	13,286	22,142	22,240	21,525	19,070	11,015
# firms	3,296	3,278	3,204	2,008	3,287	3,297	3,156	3,007	1,817
Mean dep. var. (log)	1.81	1.82	1.82	1.78	1.81	1.81	1.81	1.82	1.79
Mean layers (level)	6.46	6.49	6.50	6.26	6.47	6.46	6.47	6.52	6.32
R^2	0.20	0.22	0.20	0.18	0.16	0.17	0.18	0.17	0.24
Within-R^2	0.04	0.07	0.04	0.01	0.01	0.01	0.01	0.01	0.06
Year FE?	Y	Y	Y	Y	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 5: Differences in hiring, promotions, and tenure

This table reports regression estimates using observations at the firm-year level. In all columns, the right-hand side variable is "log # layers $(t-1)$," which is the logarithm of the number of hierarchical layers in the previous year. In column 1, the dependent variable is the fraction of new LinkedIn employees during the year ("% new"). In column 2 ("% promote, layer"), the dependent variable is the fraction of existing LinkedIn employees promoted during the year, where a promotion is a move to a role in a strictly higher layer of the firm's estimated hierarchy. Lateral moves and demotions do not count. In columns 3 and 4, the dependent variable is the logarithm of the average tenure of the firm's employees (in years). In column 3, this variable is constructed excluding employees who joined during the year ("Non-new"), while in column 4 it includes all employees ("All"). In column 5, "ESOP" is an indicator variable equal to one if the firm has an employee stock option plan or employee stock bonuses. In column 6, "Internal CEO" is an indicator variable equal to one if the firm appoints a new CEO who had worked at the firm before starting the CEO position (the sample consists of all CEO turnover events). In column 7, "CEO pay" is the logarithm of the total compensation (salary and bonus) of the CEO. In column 8, "CEO pay ratio" is the ratio of the CEO's pay to the median worker's pay, winsorized at the 5% and 95% (by year). Industry fixed effects are based on 4-digit SIC industries. Size decile fixed effects are based on the firm's total assets in the previous year. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Employees		Tenure (log years)		ESOP?	Internal CEO	CEO pay	CEO pay ratio
	% new	% promote layer	Non-new	All				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log # layers $(t-1)$	-0.0080**	0.0069***	0.057***	0.073***	0.095***	0.028	0.064	35.0***
	(0.0036)	(0.0016)	(0.022)	(0.026)	(0.024)	(0.055)	(0.064)	(5.77)
Observations	22,257	22,257	22,257	22,257	13,533	1,210	11,180	7,266
# firms	3,304	3,304	3,304	3,304	2,071	862	1,470	1,861
Mean dep. var.	0.17	0.04	1.72	1.52	0.14	0.59	8.60	109.26
R^2	0.02	0.01	0.07	0.07	0.10	0.01	0.30	0.29
Size decile FE?	Y	Y	Y	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y	Y	Y	Y

Table 6: Operating performance

This table reports regression estimates using observations at the firm-year level. In columns 1-4, the dependent variables are measures of profitability and productivity. "ROA" is the ratio of operating income before depreciation (Compustat item "oibdp") to total assets (item "at"); "ROE" is the ratio of net income (item "ni") to common equity (item "ceq"); "Net margin" is the ratio of net income to total revenue (item "revt"); "Rev/emp." is the ratio of total revenue to the number of employees (item "emp"). The dependent variable in column 5 is a cost measure: "Wages" is the ratio of staff expense (item "xlr") to total assets. All these variables are winsorized at the 5% level. In all regressions, the right-hand side variable is "log # layers $(t-1)$," which is the logarithm of the number of hierarchical layers in the previous year. Industry fixed effects are based on 4-digit SIC industries. Size decile fixed effects are based on the firm's total assets in the previous year. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	ROA	ROE	Net margin	Rev/emp.	Wages
	(1)	(2)	(3)	(4)	(5)
log # layers $(t-1)$	0.0376***	0.0263	0.184**	-110.4***	0.0113
	(0.00809)	(0.0176)	(0.0808)	(24.98)	(0.00873)
Observations	17,608	22,220	21,453	21,604	4,646
# firms	3,029	3,302	3,218	3,297	680
Mean dep. var.	-0.01	-0.06	-0.53	618.01	0.09
R^2	0.55	0.14	0.41	0.63	0.74
Within-R^2	0.23	0.03	0.08	0.05	0.12
Size decile FE?	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y

Table 7: Firm patenting

This table reports regression estimates using observations at the firm-year level. In column 1, the dependent variable is the ratio of R&D expenses (Compustat item "xrd") to total assets (item "at"), winsorized at the 5% level. In columns 2 and 3, the dependent variable is the number of patents applied for by the firm in a given year. In columns 4 and 5, the dependent variable is the number of citations received by these patents. Citations are computed as the average number of citations received by the end of the sample across all patents of the firm in a given year. Columns 2 and 4 provide estimates from Poisson regressions. Columns 3 and 5 provide estimates using the IHS (inverse hyperbolic sine) transform of the dependent variable. In all regressions, the right-hand side variable is "log # layers $(t-1)$," which is the logarithm of the number of hierarchical layers in the previous year. Industry fixed effects are based on 4-digit SIC industries. Size decile fixed effects are based on the firm's total assets in the previous year. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)
	R&D	# patents		# cites	
		Poisson	IHS	Poisson	IHS
log # layers $(t-1)$	-0.021***	1.23***	0.25***	0.068	-0.013
	(0.0074)	(0.26)	(0.076)	(0.27)	(0.061)
Observations	13,177	20,046	22,257	5,559	5,596
# firms	2,035	2,972	3,304	1,195	1,208
Mean dep. var.	0.13	18.83	0.71	1.89	0.77
Pseudo-R^2		0.83		0.29	
R^2	0.53		0.44		0.35
Size decile FE?	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y

Table 8: Stock and firm volatility

This table reports regression estimates using observations at the firm-year level. In columns 1 and 2, the dependent variable is the square root of the sum of squared daily stock returns over the year. The raw measure is used in column 1, while it is winsorized at the 5% level in column 2. In columns 3 and 4, the dependent variables are operating asset volatility and cash flow volatility, respectively. These measures are computed as in Gormley and Matsa (2016). In column 5, the dependent variable is the debt-to-asset ratio (leverage), which is computed as the ratio of total debt (Compustat item "dt") to total assets (item "at"). In column 6, the dependent variable is the cash-to-asset ratio, which is computed as the ratio of cash and marketable securities (item "che") to total assets. Both ratios are winsorized at the 5% level. In all regressions, the right-hand side variable is "log # layers $(t-1)$," which is the logarithm of the number of hierarchical layers in the previous year. Industry fixed effects are based on 4-digit SIC industries. Size decile fixed effects are based on the firm's total assets in the previous year. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Stock σ		Operations and financial			
		Wins.	σ(OA)	σ(CF)	$\frac{Debt}{Assets}$	$\frac{Cash}{Assets}$
	(1)	(2)	(3)	(4)	(5)	(6)
log # layers $(t-1)$	-0.044***	-0.042***	-0.033**	-0.0063	-0.0024	-0.025***
	(0.011)	(0.010)	(0.016)	(0.0047)	(0.011)	(0.0083)
Observations	20,071	20,071	20,025	15,273	19,284	21,916
# firms	3,146	3,146	3,145	2,572	2,966	3,302
Mean dep. var.	0.54	0.53	0.40	0.04	0.28	0.23
R^2	0.19	0.23	0.10	0.14	0.05	0.10
Size decile FE?	Y	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y	Y

Table 9: Resilience to shocks: Covid-crash stock returns and product market fluidity

This table reports regression estimates. In columns 1-4, the unit of observation is a firm in fiscal year 2020 and the dependent variable is the firm's cumulative daily excess stock return over the Covid collapse period from February 3 to March 23, 2020, which is computed as $\exp\left(\sum_t \log(1 + r_{it} - r_{ft})\right) - 1$, where r_{it} is the daily stock return, r_{ft} is the daily one-month Treasury-bill rate, and the sum runs over the trading days in the collapse period. The cumulative excess return is winsorized at the 5th and 95th percentiles. Financial firms (SIC 6000-6999) and utilities (SIC 4900-4999) are excluded. "Log # layers" is the logarithm of the number of hierarchical layers in fiscal year 2020; in columns 3 and 4 it is interacted with an indicator for the Covid-exposed industries identified by Fahlenbrach, Rageth, and Stulz (2021) ("Covid industry") and its complement ("other industry"); the industry indicator itself is absorbed by the industry fixed effects. "Cash / assets" is cash and short-term investments over total assets and "Book leverage" is long-term plus current debt over total assets, both winsorized at the 5th and 95th percentiles. In columns 5 and 6, the unit of observation is a firm-year and the dependent variable is the logarithm of the number of hierarchical layers. "Product market fluidity" is the firm-year measure of Hoberg, Phillips, and Prabhala (2014), matched to the firm by gvkey and the calendar year of its fiscal year end, and available through 2023. Columns 5 and 6 include fiscal-year fixed effects; column 6 additionally includes industry fixed effects. Size decile fixed effects are pooled deciles of contemporaneous total assets; industry fixed effects are 4-digit SIC. Columns 1-4 use a common sample with cash and leverage observed; firms that are alone in their 4-digit SIC are dropped by the fixed-effects estimator. "–" denotes not applicable. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Cum. excess return, Feb 3 - Mar 23 2020				Log # layers	
	(1)	(2)	(3)	(4)	(5)	(6)
Log # layers	-0.014	-0.012				
	(0.016)	(0.016)				
Log # layers × Covid industry			-0.073**	-0.074**		
			(0.036)	(0.035)		
Log # layers × other industry			-0.0015	0.0015		
			(0.017)	(0.017)		
Cash / assets		0.040		0.042		
		(0.027)		(0.027)		
Book leverage		-0.18***		-0.18***		
		(0.025)		(0.025)		
Product market fluidity					-0.0080***	-0.0039**
					(0.0013)	(0.0016)
Size decile FE?	Y	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	N	Y
Year FE?	–	–	–	–	Y	Y
Observations	1,454	1,454	1,454	1,454	20,407	20,403
# firms					3,289	3,285
Mean dep. var.	-0.36	-0.36	-0.36	-0.36	1.81	1.81
Adj. R^2	0.23	0.26	0.23	0.27	0.06	0.19

Table 10: Firm response to AI: number of hierarchical layers

This table reports regression estimates using observations at the firm-year level. The dependent variable is the logarithm of the number of hierarchical layers in year t. The right-hand side variables are various measures of AI adoption based on job posting data, obtained from Babina, Fedyk, He, and Hodson (2024), and each enters lagged by one year $(t-1)$. They include the share of AI job postings, broadly (column 1) and narrowly (column 2) defined, the average share of AI skills (column 3), and the average maximum share of AI skills (column 4). Size decile fixed effects are pooled deciles of contemporaneous total assets, matching the timing of the dependent variable. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log # layers			
	(1)	(2)	(3)	(4)
Share job postings AI $(t-1)$	-0.42***			
	(0.15)			
Share job postings AI (narrow) $(t-1)$		-0.36**		
		(0.15)		
Average share (skills) $(t-1)$			-0.38***	
			(0.13)	
Avg max share $(t-1)$				-0.072***
				(0.025)
Observations	6,386	6,386	6,386	6,386
# firms	1,729	1,729	1,729	1,729
Mean dep. var.	1.83	1.83	1.83	1.83
Within R^2	0.02	0.02	0.02	0.02
Size decile FE?	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y
Firm FE?	Y	Y	Y	Y

Appendix

A1 Estimating the number of organizational layers

This appendix describes how the layer count is constructed from the job titles observed in the LinkedIn data. The procedure takes a firm's observed promotions, estimates a ranking of its roles, and selects the number of distinct layers that the ranking can support given the noise in the data. The steps are as follows.

1. **Titles to roles.** Each worker-year's raw job title is cleaned and mapped to a role. The mapping strips geographic and seniority-free modifiers, merges synonyms, and groups title variants that describe the same position. A string or bi-gram can serve as a role base only if it ranks in the top half of the firm's strings or bi-grams by frequency and falls within the most frequent 95 percent of string occurrences (80 percent for bi-grams); bi-grams take precedence, and the single most frequent item is always kept. For example, consider a firm with only two titles, "Software Engin, UX" and "Software Engin, backend." Among single strings, "Software" and "Engin" each appear twice while "UX" and "backend" appear once; among bi-grams, "Software Engin" appears twice while "Engin UX" and "Engin backend" appear once each. Because bi-grams take precedence, both titles map to the single role "Software Engin." Within what becomes the firm's largest internal labor market (step 4), roles that share the same O*NET occupation code and the same seniority level are then merged into a single role; roles without a reliable occupation code are left as mapped. All subsequent steps operate on roles, not raw titles.

2. **The promotion network.** For each firm, a directed edge connects role a to role b when a worker moves from a to b in consecutive observations at the same firm. Each edge carries a weight equal to the number of workers who make that transition, and opposing flows between the same pair of roles are netted: the surviving edge carries $\max(w_{ab} - w_{ba}, 0)$, and pairs whose flows offset exactly are dropped. Consider a firm in which associates are promoted to vice president, vice presidents move to either principal or manager, and both principals and managers are promoted to managing director. The same firm also employs clerks, who are promoted to senior clerk. These two promotion paths are not connected to each other, and the difference in their size determines which one the procedure ranks, as described in step 4 below.

3. **Pruning.** Before the largest component is selected, the network is pruned following the leave-10-percent-out rule of Algorithm 1 in Huitfeldt, Kostøl, Nimczik, and Weber (2023): a node that is an articulation point of the worker-role graph, meaning a node whose removal disconnects it, is removed if the connected piece that the node holds together covers less than 10 percent of the firm's roles. This drops incidental bridges, such as a single worker whose idiosyncratic job change is the only link between two otherwise unrelated groups of roles, while rare but genuine moves in large groups are kept. Huitfeldt et al. introduced the step for measurement error in administrative occupation codes; with self-reported LinkedIn titles such one-off transitions are more common, not less. Our implementation applies the cut to role nodes as well as workers, and only inside small pieces: in no sample firm does pruning change the largest component. The ILM counts in Table 2 are taken before pruning.

4. **The largest internal labor market.** The analysis restricts attention to the largest weakly connected component of the pruned promotion network, which we call the firm's largest internal labor market (ILM). In the example, the associate-to-managing-director component is the largest ILM; the clerk pair forms a separate, smaller ILM.

5. **The ranking objective.** A candidate ranking orders the roles of the largest ILM from bottom to top. A violation is a promotion that moves *down* the candidate ordering, and each violation is weighted by the number of worker transitions on the violated edge: reversing a move made by one worker is not equivalent to reversing a move made by a hundred. The estimand is the ranking that minimizes the total weight of violations: if the ordering is right, promotions should flow up.

6. **The search.** Roles are initially ordered by the weight of the promotions they send, with the role sending the most promotions placed at the bottom. The search then repeatedly proposes swaps of roles, recomputes the violation weight, and accepts any swap that does not increase it. Because many orderings can achieve the same minimum (roles that never interact are incomparable), the search averages each role's position over the orderings it visits at the lowest violation weight found so far, restarting the average whenever a strictly better ordering is found. This average is one replication's estimated ranking.

7. **Replications and rank uncertainty.** The search is replicated 25 times per stage

(50 at the deepest stage), each from the same initial ordering with an independently seeded sequence of swap proposals; for firms whose promotion networks contain many tied out-degrees, additional replications use alternative starting orders. The search runs in three stages. A stage ends when the search has gone a set number of swap proposals without improving the violation weight. That patience is twice the firm's number of role pairs in the first stage, floored at 5,000 and capped at one million, and doubles at each later stage. A fourth stage repeats the deepest search with 50 replications. Across replications, each role i receives a mean rank $\bar{r}_i$ and a rank standard deviation s_i. The dispersion s_i measures how precisely the data pin down role i's position.

8. **The noise floor.** The tolerance for distinguishing layers is built in three steps. First, roles whose rank dispersion is statistically indistinguishable from a uniformly random placement, $s_i \geq c \cdot N/\sqrt{12}$ with $c = 0.5$, are excluded; $N/\sqrt{12}$ is the standard deviation of a uniformly random rank among the firm's N ranked roles, so these are roles the data cannot place at all.[32] Second, the tolerance is $A = \frac{1}{N'} \sum_i s_i^2$, the average squared rank dispersion across the N' surviving roles. Third, the tolerance is scaled by γ, the ratio of LinkedIn-covered workers to Compustat employment, capped at one. Incomplete coverage can only tighten the standard a firm's ranking must meet.

9. **Selecting the number of layers.** The layer count is the smallest K such that the within-layer variance of the surviving roles' mean ranks, computed by exact one-dimensional k-means, is no greater than γA. Intuitively, the procedure adds layers until the remaining within-layer dispersion of role ranks is small relative to the noise in the ranks themselves.

10. **Convergence and the alternative layer measure.** The staged search produces a per-firm signal of whether the ranking settled. A firm is flagged when its layer count rises across the stages, the noise floor falls, and the rank order barely changes between stages (a rank correlation above 0.95). The extra layers then come from the shrinking floor, not from a better ordering. For firms whose search converged, the noise-floor rule applies. The search for some firms does not settle within the compute budget. These are typically the largest firms that have dense, tie-heavy promotion networks.

[32] A rank drawn uniformly at random from $\{1, \ldots, N\}$ has mean $(N+1)/2$ and second moment $(N+1)(2N+1)/6$, so its variance is $\frac{(N+1)(2N+1)}{6} - \frac{(N+1)^2}{4} = \frac{N^2-1}{12}$ and its standard deviation is $\sqrt{(N^2-1)/12} \approx N/\sqrt{12}$. The exclusion rule uses the approximation $N/\sqrt{12}$.

The noise floor count for these firms depends on the budget rather than on the firm's characteristics. If we search too little, the replication samples a very wide range of orderings. The result is a very high noise floor, and the rule collapses the firm to two to four layers. However, if we add search and compute, the replications agree on an arbitrary ordering of incomparable roles. The floor in this case shrinks and the rule over-resolves to the firm having too many layers. These firms instead use an alternative layer measure, an agglomerative rule: with the roles sorted by mean rank, adjacent groups are merged, least distinguishable pair first, until every remaining pair of neighbors differs by more than $z(N)$ standard errors, where $z(N) = \sqrt{2\ln((N-1)/\alpha)}$ and $\alpha = 0.05$. Two neighboring groups are compared by the gap between their mean ranks divided by a two-sample standard error, built from the average squared rank dispersion within each group. The pair with the smallest ratio is merged first.[33]

11. **Assignment and the firm-year count.** Whichever rule selects K, the assignment of roles to layers is identical for every firm: the exact one-dimensional k-means partition of mean ranks into K groups, applied to all roles in the largest internal labor market, including those excluded when K was chosen. The alternative measure therefore contributes only the number of layers, not the grouping; a single assignment procedure keeps the two selection rules from introducing differences in how roles are grouped, only in how many groups there are. A firm's layer count in year t is the number of layers staffed by at least one worker in that year. This firm-year count is the paper's measure.

A2 Sample construction

The sample starts from the 8,128 firms matched to the resume data and applies the conditions below in order. Each count is the number of firms removed by that condition among firms that survived the earlier ones.

- Industry: SIC codes 4900–4999, 5200, 5300–5399, 7100–7299, and 4400–4599 (356 firms).

[33] With N ranked roles there are $N-1$ adjacent boundaries to test. Holding the family-wise error rate at α by the Bonferroni rule allocates size $\alpha/(N-1)$ to each test, and the Gaussian upper quantile satisfies $z_p \approx \sqrt{2\ln(1/p)}$ for small p; substituting $p = \alpha/(N-1)$ yields $z(N) = \sqrt{2\ln((N-1)/\alpha)}$. The alternative clustering method of Huitfeldt, Kostøl, Nimczik, and Weber (2023) motivates the test: they merge adjacent occupations whose rank difference falls below the higher-ranked occupation's estimated standard deviation, capturing the individual rank uncertainty that an overall noise floor ignores. The alternative measure replaces that one-standard-deviation rule with a two-sample standard error and the family-wise threshold above.

- No coverage measure: firms whose resume panel was never matched against Compustat headcount (2,563 firms; excluded from the counts in Section 5).
- Coverage: LinkedIn headcount below 15% of Compustat employees, pooled over the sample years (534 firms).
- Compustat: employees, total assets, and U.S. incorporation and headquarters reported (116 firms).
- Size and entity type: at least ten employees in some sample year; REIT, BDC, and MLP shells, nine hand-audited false matches, and five gig-economy platforms removed (109 firms).
- Worker coverage: at least 25% of worker-years in the largest internal labor market (710 firms).
- Over-attribution: LinkedIn headcount above twice the Compustat headcount, pooled over the sample years (68 firms; because the ratio is pooled, an excluded firm loses all of its years).
- Layer count: firm-years with a single estimated layer are dropped, and firms with no year of measured layers are removed (218 firms leave the sample because none of their years survives; nearly all of these had no measured layer count in any year).

The remaining 3,454 firms form the regression sample.

A3 Known job ladders: investment banks

We rely on LinkedIn's representativeness regarding title distribution and promotion reporting. Ideally, we would compare a firm's final corporate hierarchy—the number of layers and titles within each—to a self-reported organizational chart. Those charts are unfortunately unavailable in any 10-K filing that we have searched or on company websites. Instead, we explore an industry with a fairly standardized job ladder for a major job role: investment banks and bankers. The titles in this career path should be ranked as expected and be separated into different hierarchical layers.

The exercise proceeds in three steps. Firstly, we identify the titles and rank ordering of investment banking career paths of investment banks. Next, we assess whether these titles

(i) have the expected ranks in our data and (ii) are separated into different layers. Thirdly, we consider only the individuals with investment banking roles and re-run the ranking and clustering algorithm. This final step is a check for whether the data and algorithm can see the hierarchy without different firm roles. We implement this test on Jefferies and Evercore.[34]

Jefferies' investment banking titles run from analyst to associate, vice president, director, and managing director. On the full firm, the estimator finds 7 layers, and the rungs rise along the ladder: the investment banking analyst role sits in layer 2, associates in layer 3, vice presidents and directors in layer 5, and managing directors in layer 6. The other roles that carry the word "analyst," such as bare "analyst," "fixed income analyst," and "operations analyst," are back-office jobs and sit in layer 1, below the bankers. Restricting the panel to workers with banking titles (step 3; 29,102 worker-years, 127 roles) yields 9 layers, with analysts in layer 2, associates in layer 3, vice presidents in layer 6, directors in layer 7, and managing directors in layer 8, so every rung now has its own layer.

Evercore adds a senior managing director rung above managing director. On the full firm, the estimator finds 7 layers: the investment banking analyst role sits in layer 2, associates in layer 3, vice presidents in layer 4, and directors, managing directors, and senior managing directors together in layer 5. The remaining "analyst" roles are senior and control-function jobs ("senior analyst," "senior risk analyst," "supervisory analyst"), which sit above the entry rung. Restricting the panel to banking titles (step 3; 15,112 worker-years, 66 roles) yields 8 layers and separates the top of the ladder: analysts in layer 2, associates in layer 4, vice presidents in layer 5, directors and managing directors in layer 6, and senior managing directors in layer 7. One operations-analyst role is ranked near the top of the banker-only network, which is the main departure from the known ladder. In both banks, the estimated layers reproduce the known order of titles, and conditioning on banking roles separates rungs that share a layer in the full firm.

These validation tests reassure that the LinkedIn and promotion data reveal the deeper corporate hierarchy.

[34]Some of the largest and well-known investment banks are also in the data. We find that the ranking of investment bank titles and layers behaves similarly to these two examples. However, the expansive set of non-investment bank roles with similar titles makes it challenging to create a clean sample for a narrow analysis of one job role.

Figure A1: Internal labor markets in sample firms

This figure illustrates the first step of the algorithm, identifying internal labor markets from worker transitions across occupations, for four sample firms. Each bubble is one internal labor market (a weakly-connected component of the firm's promotion network). The horizontal axis is the number of occupations (roles) in the internal labor market; the vertical axis is the average experience of its workers, measured in years since the start of each worker's career (from the full employment history, not the sample window); bubble size is proportional to the internal labor market's total worker-years. Worker-level experience is winsorized at 40 years. A subset of internal labor markets is labeled with the first role name of the component.

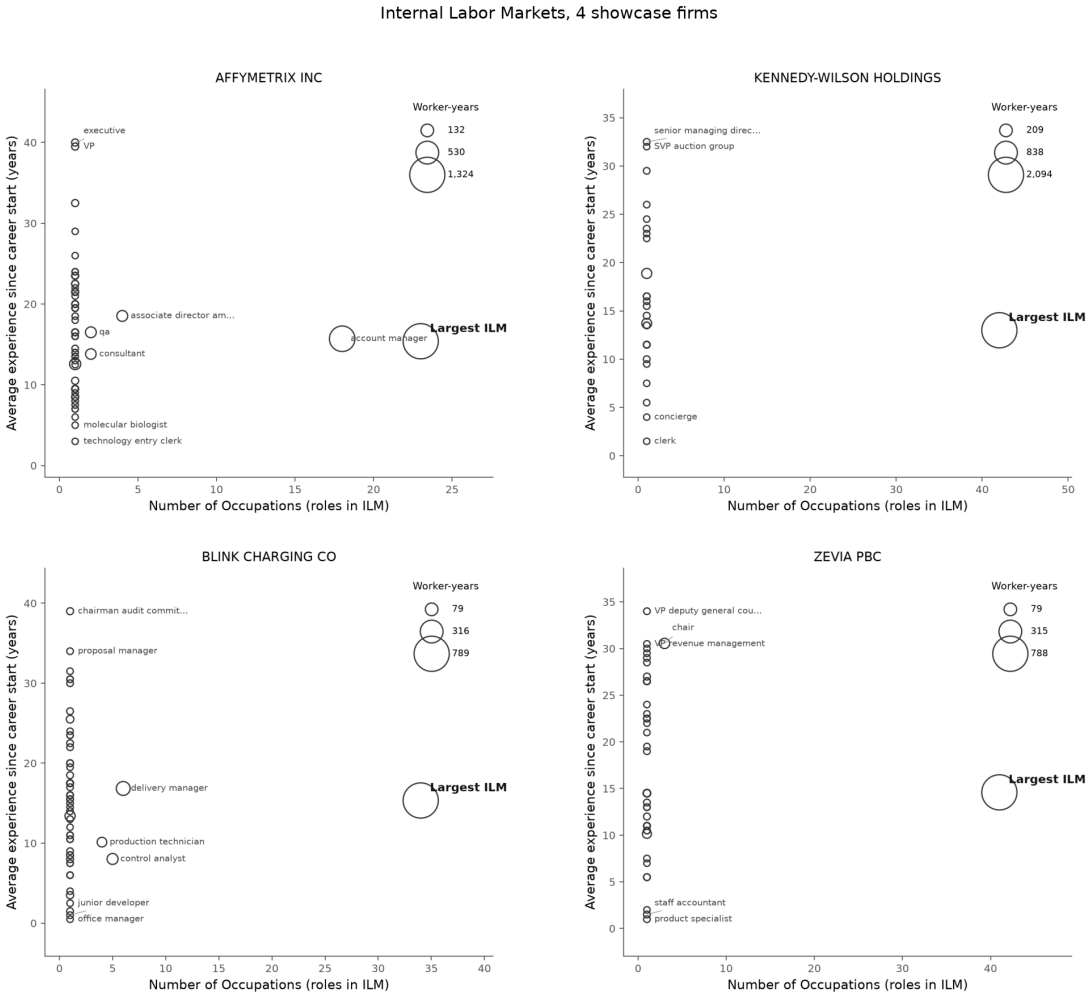


Figure A2: Ranks and hierarchy levels in sample firms

This figure illustrates the second step of the algorithm for four sample firms. Within each firm's largest internal labor market, every role receives an estimated rank (the vertical axis), and roles are clustered into discrete hierarchy levels (the horizontal axis). Points are individual roles; vertical bars are the rank uncertainty; shaded bands denote the hierarchy levels; a subset of roles is labeled. Roles whose estimated rank is statistically indistinguishable from random placement are omitted.

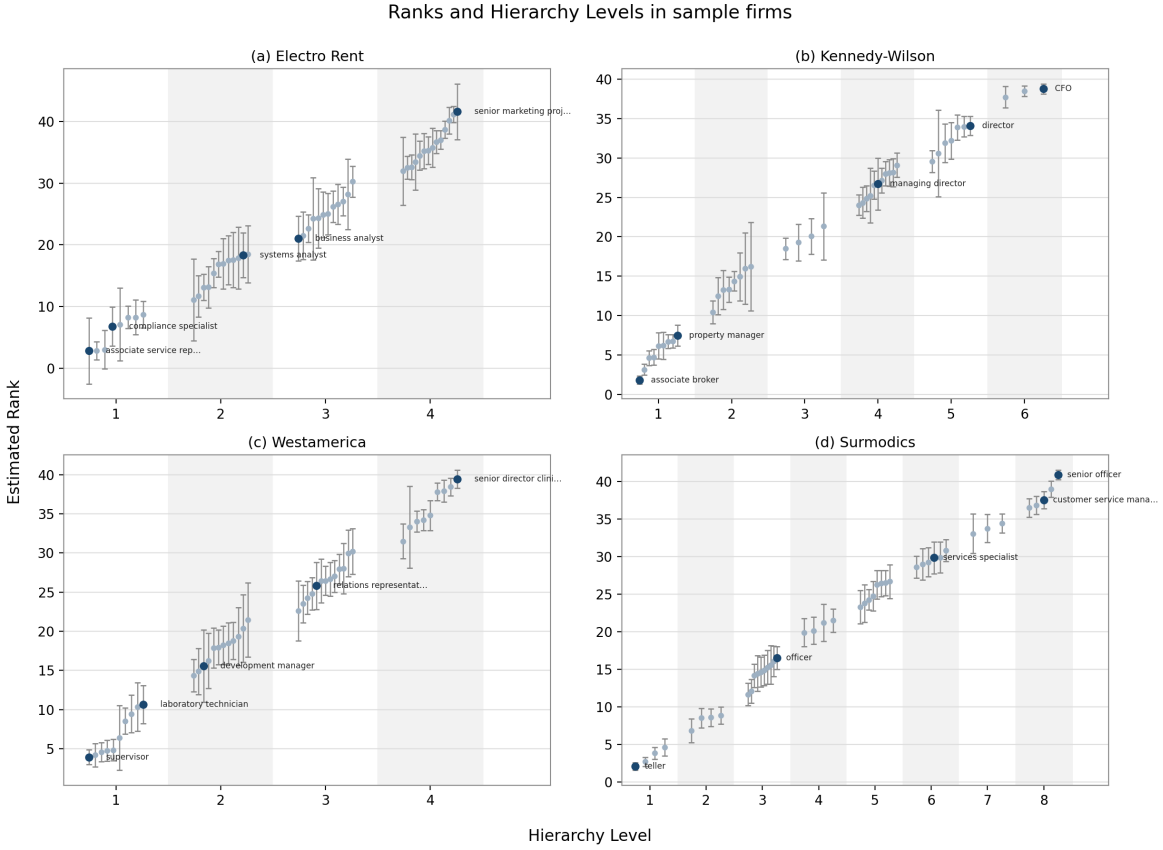
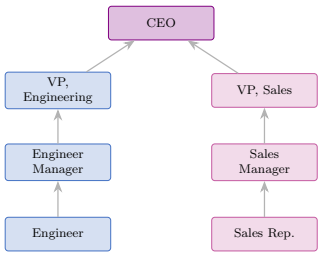
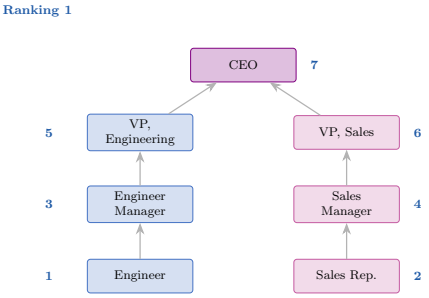


Figure A3: From promotions to layers: a toy example in four steps

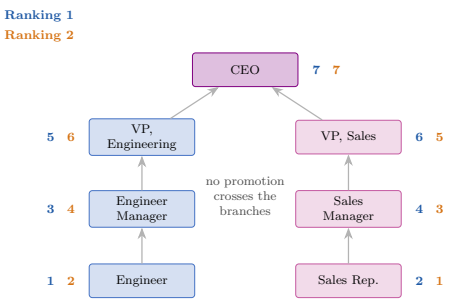
This figure gives a simplified view of the steps that take the raw panel of resume data to a set of hierarchy layers. Step 1 shows the data. Each arrow is an observed promotion, drawn from the role a worker left to the role the worker entered. Step 2 ranks the roles: ranks are assigned so that no promotion points downward, and the blue ranking is consistent with every observed move. Step 3 shows that many rankings fit the same data. The orange ranking is equally consistent, because no worker ever moves between the two branches, so the promotions cannot order the engineering roles against the sales roles. Twenty rankings fit these promotions equally well. Step 4 averages and clusters. Each role's rank is averaged across all consistent rankings, so roles the data cannot tell apart tie. For each candidate number of layers K, the average ranks are then grouped into K layers by one-dimensional k-means, and the smallest K whose within-layer spread falls within the sampling noise is kept, here $K = 4$.



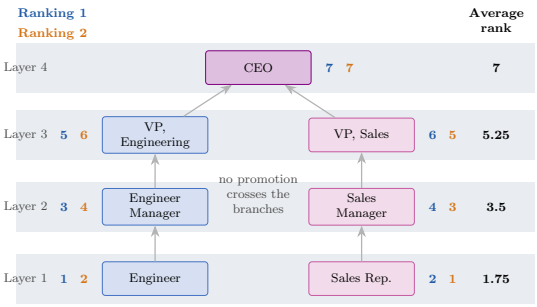
Step 1: the data. Arrows show observed promotions: from the role a worker left to the role they entered.



Step 2: rank the roles. Assign ranks so that no promotion points downward. The blue ranking is consistent with every observed move.



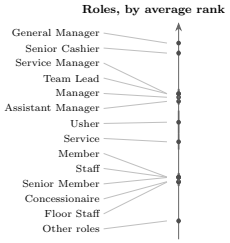
Step 3: many rankings fit. The orange ranking is equally consistent: no worker ever moves between the branches, so the data cannot order Engineering against Sales. Twenty rankings fit these promotions equally well.



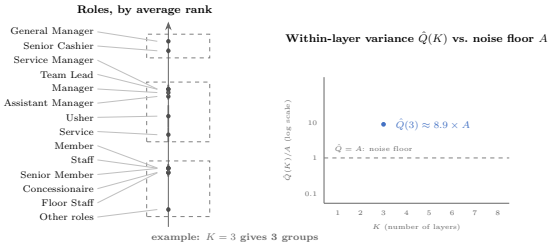
Step 4: average, then cluster. Average each role's rank across all consistent rankings: roles the data cannot tell apart tie. Then, for each candidate K, group the average ranks into K layers (1-D k-means) and keep the smallest K whose within-layer spread is within the sampling noise: here $K = 4$.

Figure A4: From ranks to layers at a real firm: Regal Entertainment

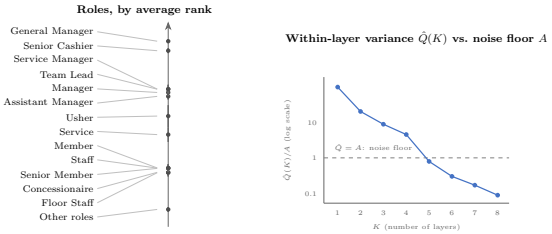
This figure shows the layer count at a single firm, Regal Entertainment Group. Panel (a) plots each role's average rank across the equally good orderings, with whiskers at plus or minus one standard deviation of the rank across repetitions. Panel (b) takes one candidate number of layers K, groups the average ranks into K layers by one-dimensional k-means, and computes the within-layer variance $\hat{Q}(K)$; at $K = 3$ the within-layer variance is roughly nine times the noise floor A. Panel (c) repeats the calculation for every candidate K. Panel (d) applies the noise-floor rule, which keeps the smallest K with $\hat{Q}(K)$ at or below A, giving five layers. Values are the estimated per-role mean ranks and standard deviations.



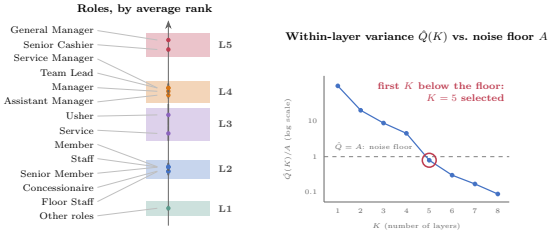
Step 1: estimated ranks. Each dot is one role at Regal Entertainment Group: its average rank across the equally good orderings. Whiskers show ±1 SD of the rank across repetitions, the estimator's uncertainty.



Step 2: one candidate at a time. Pick a K; group the roles into K layers with the fixed grouping rule (1-D k-means on the average ranks); compute $\hat{Q}(K)$, the within-layer variance. At $K = 3$, the dashed grouping leaves $\hat{Q}(3)$ at 8.9 times the noise floor A (dashed line).



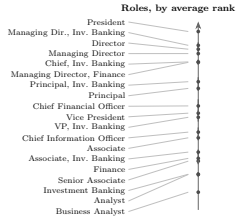
Step 3: repeat for every K. Each dot is one candidate: pick K, group, compute. $\hat{Q}(K)$ falls as K grows, since more layers always fit tighter. The question is where to stop.



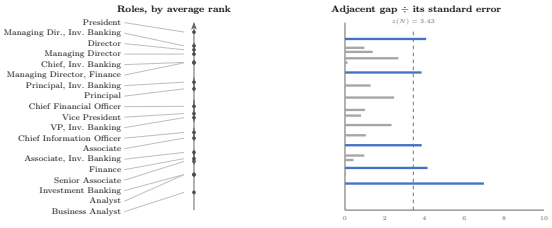
Step 4: stop at the noise floor. The rule keeps the smallest K with $\hat{Q}(K) \leq A$: here $K = 5$. Finer splits would slice noise, not structure. The bands show the five layers, from the $K = 5$ run of the same grouping rule.

Figure A5: From ranks to layers when the ranking search does not settle: Greenhill & Co

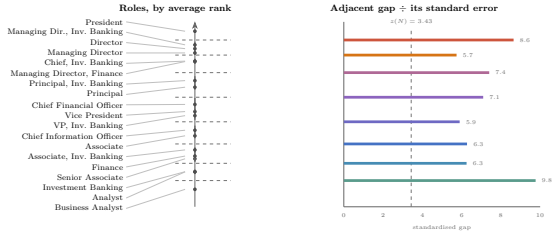
This figure shows the layer count for a firm whose ranking search did not settle, Greenhill & Co (19 roles). For firms whose search settles, the number of layers comes from the noise-floor rule; for the remainder it comes from the alternative layer measure shown here. Panel (a) plots each role's average rank with whiskers at plus or minus one standard deviation. Panel (b) standardizes each adjacent gap by its two-sample standard error and compares it with the threshold $z(N) = \sqrt{2\ln((N-1)/\alpha)} = 3.43$ at $\alpha = 0.05$, the Bonferroni correction for the $N - 1 = 18$ boundary tests at this firm. Panel (c) pools roles whose gap falls below the threshold, smallest standardized gap first, recomputing the gap on the pooled groups until every remaining boundary clears the threshold; eight boundaries survive. Panel (d) reports the resulting $K = 9$ and the assignment of roles to layers, which uses the same rule at every firm: one-dimensional k-means on the average ranks given K. The noise-floor rule would give $K = 10$ for this firm.



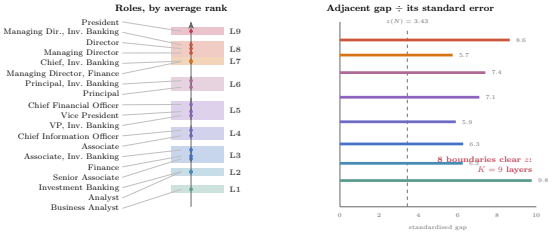
Step 1: estimated ranks. Each dot is one role at Greenhill & Co: its average rank across the equally good orderings. Whiskers show ±1 SD of the rank across repetitions, the estimator's uncertainty. The ranking search for this firm did not settle, so the number of layers is set by the alternative layer measure rather than by the noise-floor rule.



Step 2: test each adjacent gap. Sort the roles by average rank and standardise each adjacent gap by its two-sample standard error, $\sqrt{s_i^2 + s_{i+1}^2}$. A gap counts as a layer boundary only if it clears $z(N) = \sqrt{2\ln((N-1)/\alpha)} = 3.43$ at $\alpha = 0.05$, the Bonferroni threshold for the $N - 1 = 18$ boundary tests run at this firm. Most gaps fall far short.



Step 3: merge what the data cannot separate. Roles separated by a gap below the threshold are pooled, smallest standardised gap first, and the gap is recomputed on the pooled groups; merging stops when every remaining boundary clears z. 8 boundaries survive, each shown with its standardised gap.



Step 4: the selected number of layers. The surviving boundaries give $K = 9$. Given K, roles are assigned to layers by the same rule used for every firm: 1-D k-means on the average ranks. The noise-floor rule would have given $K = 10$.

Figure A6: O*NET work styles across the within-firm hierarchy

This figure validates the layer measure against the occupational content of work. Each SOC-matched LinkedIn title inherits its occupation's importance ratings on the 16 classic O*NET work styles (survey-rated, 1 = not important to 5 = extremely important). Worker-year-weighted mean importance is computed within each quintile of the firm's layer assignment (Q1 = bottom layers, Q5 = top), pooled over the estimation sample, 2013–2024. Cells show the deviation of each quintile's mean from the style's cross-quintile average, so warm cells mark styles over-represented at that level. Because style ratings are constant within a SOC, the gradient reflects occupational composition across layers, not trait differences among individuals in the same job. Source: work-style importance ratings from the O*NET database (release 29.3); LinkedIn titles are matched to SOC occupations as described in the LinkedIn–BLS representativeness analysis of Section 6.4.

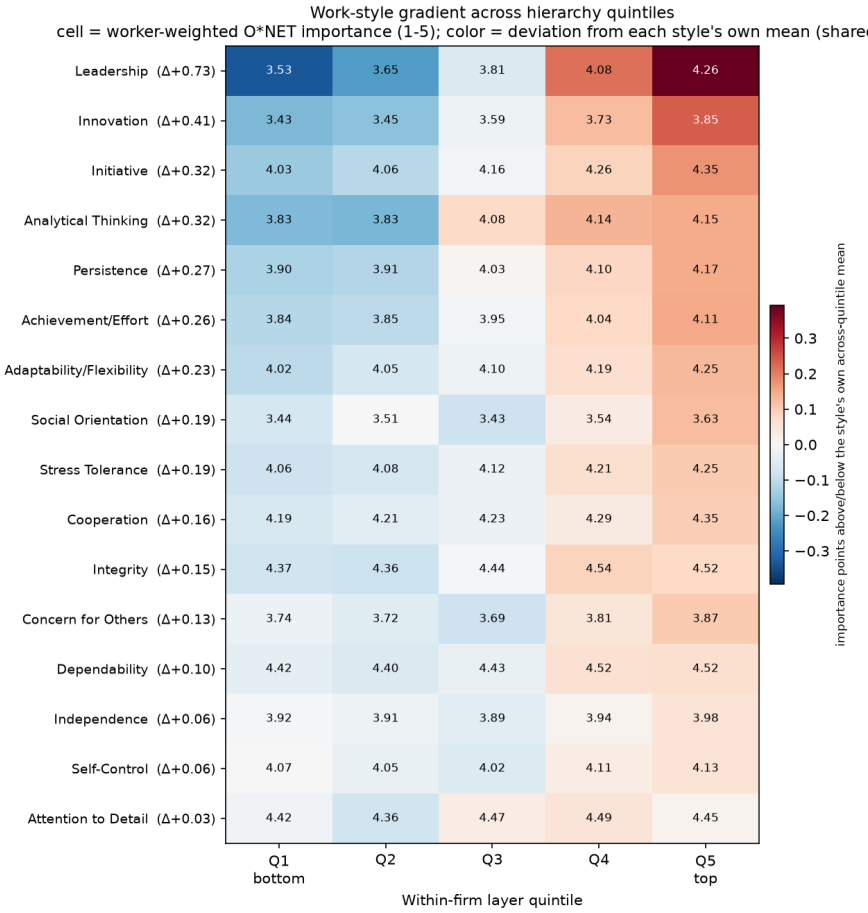


Table A1: Measurement of hierarchy in organizational studies

This table summarizes previous research that measures alternative forms of the corporate hierarchy. "Span of control" refers to the number of direct subordinates that a manager oversees. "Depth" refers to the number of positions between the top executive (e.g., the CEO) and a given position.

Paper	**Measure**	**# Layers**	**Country**	**Method(s)**
Adhvaryu, Bassi, Nyshadham, Tamayo, and Torres (2023)	# Hierarchical layers / span of control	16 (working group)	Argentina (one firm)	Firm-specific organizational chart
Baker, Gibbs, and Holmstrom (1994)	# Hierarchical layers	8	US (one firm)	Job titles with transitions
Bias, Lochner, Obernberger, and Sevilir (2026)	# Hierarchical Layers / hierarchization (HHI of employee share)	4	Germany	Occupational codes
Bloom, Sadun, and Van Reenen (2012)	Decentralization	Any (survey)	Multiple	Survey data
Caliendo, Monte, and Rossi-Hansberg (2015)	# Hierarchical layers	4	France	Occupational codes
Caliendo, Mion, Opromolla, and Rossi-Hansberg (2020)	# Hierarchical layers	4	Portugal	Occupational codes and wages
Friedrich (2022)	# Hierarchical layers	4	Denmark	Occupational codes
Garicano and Hubbard (2016)	Associate-partner ratio in law firms	2	US (law firms)	1992 Census of Services (title codes and wages)
Guadalupe, Li, and Wulf (2014)	Span of control / # managers	N/A (span)	US (300 firms)	Direct observation (Hewitt Associates, HR consulting)
Guadalupe and Wulf (2010)	Span of control / depth	N/A (span)	US (300 firms)	Direct observation (Hewitt Associates, HR consulting)
Gumpert, Steimer, and Antoni (2022)	# Hierarchical Layers / share of managerial occupations	4	Germany	Occupational codes
Huitfeldt, Kostøl, Nimczik, and Weber (2023)	# Hierarchical layers	Max 66, mean 7	Norway	Job titles with transitions
Liberti and Mian (2009)	# Approval levels for loan officers	5	Argentina	Direct observation of reporting
Rajan and Wulf (2006)	Span of control / depth	N/A (span)	US (300 firms)	Direct observation (Hewitt Associates, HR consulting)
Skrastins and Vig (2019)	# Managerial layers	8	India (one bank)	Bank reports
Tag, Astebro, and Thompson (2016)	# Hierarchical layers / span of control	4	Sweden	Occupational codes

Table A2: External pay and skill by position in the estimated hierarchy

This table relates an occupation's position in a firm's estimated hierarchy to external measures of pay and skill. An observation is a firm-occupation cell, weighted by the firm's workers in that occupation. Within each firm, occupations are sorted into worker-weighted quintiles by their average rank (position in the estimated promotion ranking; columns 1–2 and 5–6) or average layer (columns 3–4 and 7–8). Rank and layer come from the LinkedIn promotion data alone, with no BLS or O*NET input. The dependent variable is the occupation's log BLS wage (columns 1–4) or its O*NET Job Zone, a 1–5 index of required education and training (columns 5–8). Coefficients are differences from the bottom quintile (Q_1). Standard errors, clustered by firm, are in parentheses; $^{***}p < 0.01$.

	Log BLS wage				O*NET Job Zone (1–5)			
	Rank		Layer		Rank		Layer	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Q_2	0.245***	0.252***	0.190***	0.199***	0.573***	0.591***	0.458***	0.476***
	(0.020)	(0.020)	(0.022)	(0.022)	(0.062)	(0.063)	(0.064)	(0.065)
Q_3	0.406***	0.409***	0.371***	0.374***	0.900***	0.909***	0.819***	0.828***
	(0.027)	(0.027)	(0.027)	(0.027)	(0.071)	(0.072)	(0.068)	(0.069)
Q_4	0.544***	0.548***	0.509***	0.514***	1.046***	1.058***	0.971***	0.982***
	(0.020)	(0.020)	(0.022)	(0.022)	(0.066)	(0.066)	(0.065)	(0.065)
Q_5	0.740***	0.745***	0.704***	0.711***	1.257***	1.271***	1.199***	1.215***
	(0.026)	(0.026)	(0.028)	(0.028)	(0.062)	(0.062)	(0.064)	(0.063)
Industry FE	Yes	—	Yes	—	Yes	—	Yes	—
Firm FE	—	Yes	—	Yes	—	Yes	—	Yes
Observations	202,893	202,878	202,893	202,878	203,185	203,171	203,185	203,171
R^2	0.640	0.660	0.628	0.649	0.452	0.478	0.441	0.468
R^2 within	0.348	0.363	0.327	0.341	0.207	0.217	0.191	0.200

Table A3: LinkedIn coverage of BLS occupations

This table compares the occupations observed among LinkedIn workers at sample firms with the occupations the Bureau of Labor Statistics (BLS) reports for the same industries (Occupational Employment and Wage Statistics, May 2024). An occupation is a detailed BLS occupation (Standard Occupational Classification, SOC) and is *covered* in an industry if at least one LinkedIn worker at a sample firm in that industry holds a title that maps to it. Panel A reports one count per row, pooled over all industry-occupation pairs BLS reports for the sample's industries; pairs BLS withholds for confidentiality are excluded. "Occupations covered" counts each occupation equally; "employment covered" weights occupations by the number of workers BLS reports in them. Panel B lists the ten industries with the most LinkedIn worker-years at sample firms: "LinkedIn workers" is the number of worker-years at sample firms in the industry whose title maps to an occupation; "BLS occupations" is the number of occupations BLS reports for the industry; "Missed" is the number of those occupations no LinkedIn worker at a sample firm holds.

Panel A. All industries pooled

Quantity	Value
Sample firms matched to a BLS industry	3,432
Number of BLS industries	202
Industry-occupation pairs BLS reports	29,955
LinkedIn worker-years mapped to an occupation	69.8%
Occupations covered (share of all reported occupations)	74.2%
Employment covered (share of BLS employment in those occupations)	91.1%

Panel B. The ten largest industries in the sample

Industry code	Industry	LinkedIn workers	BLS occupations	Missed
522000	Credit Intermediation and Related Activities	8,063,936	157	2
517000	Telecommunications	3,784,176	168	3
518200	Computing Infrastructure Providers, Data Processing, Web Hosting, and Related Services	3,505,795	147	0
722500	Restaurants and Other Eating Places	2,943,407	89	4
513200	Software Publishers	2,575,539	120	1
334500	Navigational, Measuring, Electromedical, and Control Instruments Manufacturing	2,344,536	194	7
325400	Pharmaceutical and Medicine Manufacturing	2,283,289	185	5
523000	Securities, Commodity Contracts, and Other Financial Investments and Related Activities	2,117,976	122	0
455000	General Merchandise Retailers	2,031,618	121	7
541500	Computer Systems Design and Related Services	1,935,398	199	5

Table A4: LinkedIn versus BLS occupation shares

This table compares the share of each occupation in LinkedIn with its share in BLS. For each industry and occupation, BLS reports the occupation's share of the industry's employment; pooling all LinkedIn workers at the sample firms in the industry gives LinkedIn's share of the same occupation (a raw headcount share over the industry's covered occupations, rescaled to sum to one). The percentage difference is (LinkedIn share − BLS share) ÷ BLS share: 100% means LinkedIn shows twice the BLS share, −50% means half, 0% means agreement. Panel A summarizes the distribution of the percentage difference over industry-occupation pairs, weighting each pair by its BLS employment, so the percentiles are over BLS workers. Panel B averages the percentage difference within each occupation across industries (weighted by BLS employment) and lists the ten most over- and ten most under-represented occupations among those with at least 100,000 BLS workers.

Panel A. Distribution of the percentage difference (weighted by BLS employment)

	p10	p25	Median	Mean	p75	p90
% difference, LinkedIn vs. BLS share	−97%	−90%	−61%	3%	14%	112%

Panel B. The occupations LinkedIn over- and under-represents most

Code	Occupation	BLS workers	% difference
Most over-represented in LinkedIn			
11-1011	Chief Executives	134,680	1,273%
11-3121	Human Resources Managers	137,890	553%
11-9199	Managers, All Other	375,150	538%
11-2022	Sales Managers	541,350	524%
13-1041	Compliance Officers	191,700	512%
11-3071	Transportation, Storage, and Distribution Managers	133,450	365%
15-1299	Computer Occupations, All Other	282,740	331%
11-2021	Marketing Managers	307,140	325%
15-2051	Data Scientists	186,720	315%
13-2051	Financial and Investment Analysts	267,350	281%
Most under-represented in LinkedIn			
51-2090	Miscellaneous Assemblers and Fabricators	1,345,110	−99%
31-9096	Veterinary Assistants and Laboratory Animal Caretakers	103,220	−98%
43-6012	Legal Secretaries and Administrative Assistants	129,500	−97%
45-2092	Farmworkers and Laborers, Crop, Nursery, and Greenhouse	245,540	−97%
51-3022	Meat, Poultry, and Fish Cutters and Trimmers	140,400	−96%
31-9092	Medical Assistants	734,860	−96%
31-1120	Home Health and Personal Care Aides	1,369,190	−96%
33-9032	Security Guards	953,100	−96%
47-2061	Construction Laborers	586,550	−94%
53-7061	Cleaners of Vehicles and Equipment	289,050	−94%

Table A5: Characteristics of occupations not observed in LinkedIn

A missed occupation is one that BLS reports for the firm's industry but that no LinkedIn worker at the firm holds; because BLS data are industry-level, the firm itself is not shown to employ it. Panel A regresses the occupation's O*NET Job Zone (a 1–5 index of required education and training; columns 1–2) or log BLS wage (columns 3–4) on "Not observed at firm," an indicator for the firm missing the occupation. There is one observation per firm-occupation pair, and "Implied %" converts the wage coefficients as $e^{\beta} - 1$. Standard errors, clustered by firm, are in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. In Panel B, a firm's promotion ladder is its largest internal labor market: the largest set of its roles linked by observed promotions. A worker is on the ladder if their title maps to a role in that set. Row 1 pools all workers; row 2 weights each firm equally. Rows 3–4 repeat the occupation-share comparison of Appendix Table A4, keeping only ladder workers on the LinkedIn side in row 4. Panel C locates each missed firm-occupation cell at the same-industry peer firms that employ the occupation. Each row gives the distribution of these cells across the peers' within-firm quintiles of hierarchy position (rank or layer), wage, or Job Zone (worker-weighted median across peers) and sums to 100%.

Panel A. Pay and skill of missed occupations, relative to those LinkedIn shows

	Job Zone (1–5)		log wage	
	(1)	(2)	(3)	(4)
Not observed at firm	−0.5239***	−0.6137***	−0.2556***	−0.2991***
	(0.0046)	(0.0047)	(0.0021)	(0.0021)
Implied % (wage)	—	—	−22.6%	−25.8%
Industry FE	Yes	—	Yes	—
Firm FE	—	Yes	—	Yes
Observations	516,354	516,354	520,946	520,946
R^2	0.078	0.086	0.152	0.162
R^2 within	0.048	0.057	0.068	0.079

Panel B. The firm's promotion ladder (its largest internal labor market) and the missed occupations

	Value
How much of the workforce is on the promotion ladder	
Workers on their firm's promotion ladder, pooled across all workers	96.5%
Median firm: share of its workers on its promotion ladder	80.2%
Does restricting to ladder workers change the occupation mix?	
Median % difference vs. BLS, all mapped workers	−61%
Median % difference vs. BLS, ladder workers only	−62%

Panel C. Where the missed occupations sit at the peer firms that have them

	Q1 (bottom)	Q2	Q3	Q4	Q5 (top)
Position in hierarchy (MVR rank)	15%	22%	28%	23%	11%
Position in hierarchy (layer)	17%	20%	29%	22%	11%
Wage	53%	17%	13%	12%	5%
Job Zone (skill)	48%	15%	12%	9%	16%

Table A6: Acquisition activity

This table reports regression estimates using observations at the firm-year level. The dependent variable in column 1 is an indicator of whether the firm has an acquisition (of any type) in SDC during the year. Column 2 only considers acquisitions with reported valuations, which tend to be public targets or larger acquisitions. Column 3 reports Poisson estimates for the total number of acquisitions during the year. In column 4, the dependent variable is the logarithm of the sum of the reported valuations across all acquisitions made by the firm during the year. In column 5, the dependent variable is an indicator variable equal to one if the acquirer's primary SIC code in SDC does not match any of the SIC codes of the target. In all regressions, the right-hand side variable is "log # layers $(t-1)$," which is the logarithm of the number of hierarchical layers in the previous year. Industry fixed effects are based on 4-digit SIC industries. Size decile fixed effects are based on the firm's total assets in the previous year. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	(1) Acq.?	(2) Acq. w/ value?	(3) # Acq. (all)	(4) $\log(\sum V)$	(5) Diversify?
log # layers $(t-1)$	-0.00019	-0.0067	0.40*	0.46*	0.017
	(0.0095)	(0.0065)	(0.20)	(0.24)	(0.062)
Observations	22,257	22,257	14,175	989	1,700
# firms	3,304	3,304	2,175	486	624
Mean dep. var.	0.08	0.05	0.20	4.92	0.62
R^2	0.01	0.01		0.35	0.01
Pseudo-R^2			0.19		
Model	OLS	OLS	Poisson	OLS	OLS
Size decile FE?	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y

Table A7: Acquisition activity and subsequent hierarchy layers

This table reports regressions of the firm's number of hierarchy layers in the year following acquisition activity. Observations are at the firm-year level. The dependent variable in every column is "log # layers $(t+1)$," the logarithm of the firm's number of hierarchical layers in the following fiscal year. Each column reports a separate regression with a single acquisition measure, constructed from SDC and measured in year t. Column 1, "Any acquisition (t)," is an indicator equal to one if the firm completes at least one acquisition during the year. Column 2, "Any valued acquisition (t)," is an indicator equal to one if at least one of the year's acquisitions has a disclosed deal value in SDC; these tend to be public targets or larger acquisitions, but the variable is a disclosure screen, not a deal-size threshold. Column 3, "# acquisitions (t)," is the number of acquisitions completed during the year. Column 4, "Deal value, USD mn (t)," is the sum of disclosed deal values, in millions of nominal U.S. dollars, across the year's acquisitions. Column 5, "# public targets (t)," is the number of the year's targets that are publicly listed. All acquisition measures equal zero in firm-years with no SDC acquisition. All models are estimated by OLS on the paper's main regression sample; the constant is the intercept after absorbing the fixed effects. Year fixed effects and industry fixed effects based on 4-digit SIC industries are included; "size-decile controls" are indicators for deciles of the firm's total assets in year t. Standard errors, clustered at the firm level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Log # layers (t+1)				
	(1)	(2)	(3)	(4)	(5)
Any acquisition (t)	0.0019				
	(0.014)				
Any valued acquisition (t)		-0.011			
		(0.014)			
# acquisitions (t)			0.022**		
			(0.0095)		
Deal value, USD mn (t)				0.000016***	
				(0.0000033)	
# public targets (t)					0.080*
					(0.042)
Constant	1.94***	1.94***	1.93***	1.93***	1.94***
	(0.028)	(0.028)	(0.027)	(0.028)	(0.028)
Observations	22,265	22,265	22,265	22,265	22,265
Firms	3,305	3,305	3,305	3,305	3,305
Adj. R^2	0.189	0.189	0.191	0.191	0.190
Size-decile controls	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Industry (SIC) fixed effects	Yes	Yes	Yes	Yes	Yes

Table A8: Hierarchy layers by product life-cycle stage

We use the life cycle data from Hoberg and Maksimovic (2022). Life cycle is defined by product development, process optimization, mature product, and product decline. The table reports regressions of log layers on the stage weights, where mature product and product decline are combined into one weight and product development is omitted; standard errors clustered by firm in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log number of layers	
	(1)	(2)
Mature/declining product	0.19**	0.13*
	(0.076)	(0.072)
Process optimization	0.19***	0.16**
	(0.074)	(0.072)
Size decile FE?	N	Y
Year FE?	Y	Y
Industry FE?	Y	Y
Observations	21,444	21,438
# firms	3,360	3,360
Mean dep. var. (log)	1.81	1.81