Corporate Hierarchy

Michael Ewens and Xavier Giroud*

October 2025

Abstract

We introduce a novel measure of corporate hierarchies for over 3,100 U.S. public firms. This measure is obtained from online resumes of 7 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 ten hierarchical layers and a pyramidal organizational structure. More hierarchical firms have a more educated workforce, higher internal promotion rates, and longer employee tenure. Their operating performance is higher, but they face higher administrative costs. They are more active acquirers and produce more patents, but not higher-quality patents. They exhibit lower stock return volatility and operating asset volatility. We also examine how companies adjust their hierarchies in response to demand and knowledge shocks. We find that pharmaceutical companies increased their number of layers following the Covid-19 pandemic, while companies flattened their hierarchies following the adoption of artificial intelligence (AI) technologies. These findings are consistent with the theoretical predictions of existing models of corporate hierarchies.

*Ewens: Columbia Business School and NBER, michael.ewens@columbia.edu. Giroud: Columbia Business School, NBER, and CEPR, xavier.giroud@gsb.columbia.edu. 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.

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; 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, 2020; Gumpert, Steimer, and Antoni, 2022; Bias, Lochner, Obernberger, and Sevilir, 2024).[1] In contrast, much less is known about

[1]Moreover, most of this work characterizes organizational forms based on occupation codes (rather than

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 of over 16 million workers at U.S. public firms from 2016 to 2023. The measure we obtain, which captures the number of hierarchical layers, indicates that the average firm has ten layers, and that hierarchies have become flatter over time.

job titles or ranks), which need not capture hierarchies per se.

We start our analysis by validating our hierarchy measure. We validate it in several ways. First, we show that the layers group workers into reasonably homogeneous groups based on their title, work experience, and education. Second, while the Huitfeldt, Kostøl, Nimczik, and Weber (2023) methodology does not restrict the layers' shape, we nevertheless find that the average firm's hierarchy is pyramidal (i.e., more workers on the bottom layers). This is reassuring, as most models of firm hierarchy predict this shape. Third, the fraction of workers with "manager" 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. Overall, these patterns reassure us that our new measure captures the inherent features of corporate hierarchies.

We then show that firms with more hierarchical layers have a more educated workforce, and that advanced degrees are more prevalent the higher the hierarchical layer. This echoes the models of knowledge hierarchies, in which the layers add more "problem solvers" that require higher training costs. We further show that employees in high-layers firms have longer tenure, with promotions playing a greater role in transitions relative to outside hires, consistent with the notion that hierarchies provide employees with a clearer career trajectory and promotion path.

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. They also have higher overhead costs (SG&A) and wages, consistent with the higher costs of maintaining hierarchies. We further show that high-layer firms generate more patents and do more acquisitions. However, there is no significant difference in patent quality (as measured by patent citations) and their propensity to make diversifying acquisitions.

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 echo well 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.

In the last part of the paper, we flip the regression specification and examine the determinants of corporate hierarchies. One prediction of the models of knowledge hierarchies is that companies would add layers in response to positive demand shocks (e.g., Garicano, 2000; Caliendo, Monte, and Rossi-Hansberg, 2015). The intuition is that additional layers add managers (that is, problem solvers), allowing lower-level employees to narrow the knowledge they must master in production. To test this prediction, we examine the evolution of hierarchies at bio-pharmaceutical companies around the Covid-19 pandemic, which induced a spike in demand for their products. Consistent with theory, we find that bio-pharmaceutical companies significantly increased their number of hierarchical layers following the pandemic. Finally, a common prediction of the models of corporate hierarchies is that a decrease in the cost of knowledge acquisition and information processing reduces the need for layers (e.g., Radner, 1993; Bolton and Dewatripont, 1994; Garicano, 2000). Consistent with this prediction, we find that companies decrease their layers after adopting artificial intelligence (AI) technologies. Overall, these results are consistent with the models of corporate hierarchies, and indicate that the cross section of hierarchies observed in the data is a complex outcome of firm-specific shocks.

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 offers conclusion.

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 Zandt (2012) and Malenko (2024).

2.1 Models of corporate hierarchy

It is useful to start with a summary of how economists view the organization. Caliendo (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 loast 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). In particular, one subset of these governance models, which we refer to as “incentive hierarchy,” argue that, because much of the 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 (e.g., Lazear and Rosen, 1981; Malcomson, 1984).

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, 2024). 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

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

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. Our analysis will focus on the largest ILM in the firm. Fortunately, this ILM comprises over 70% 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. 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 receive. 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.

Once roles are ranked, we put them into hierarchical levels. The estimated mean rank and rank standard deviation inform the clustering of roles and the number of layers. The algorithm clusters ranks and uses the estimation uncertainty to select the number of clusters K. A high standard deviation suggests that a role could move up or down the ranking in a series of optimal rankings. This flexibility is a sign of lateral moves between roles and, thus,

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

a layer. Technically, as Huitfeldt, Kostøl, Nimczik, and Weber (2023) write (p. 11):

> Given a number of clusters [layers] K, the k-means algorithm assigns each occupation [role] to a hierarchy level such that the sum of the within-level squared distances in estimated ranks is minimized. The choice of the number of hierarchy levels for each internal labor market is determined by the uncertainty in the rank estimation.

We describe our final measure of hierarchical layers in Section 6 after discussing the data inputs in the next section.

5 Data

5.1 LinkedIn data

We leverage data on worker resumes provided by CoreSignal.[4] The data provider extracts resume data—including work histories, job titles, employers, and educational background—from publicly available profiles of U.S.-based workers from 2016 to 2023. While our analysis focuses on workers who are active during the sample period, we use their complete work history that is available on LinkedIn. Thus, we have a much longer view of the firms' job ladders than the final sample period.

We exclude temporary or part-time workers using titles such as "intern," "consultant," etc. 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. We speculate that the managerial class and upper management are more likely to create and manage their profiles.

A comparison of headcounts between our data and Compustat reveals that coverage is below 50% for most firms (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.[5] Rather, all we require is that the titles occur

[4] See `https://coresignal.com/`.

[5] Our data cleaning of the LinkedIn data removes extremely rare titles and transitions that would have very high rank uncertainty and thus have large impacts on the layer partition.

at least once in the data and that a promotion or demotion exists between titles.[6] In the final sample described below, we have 8,753,665 job transitions from 2017 to 2023.[7] 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 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 1,033 unique titles (median 254) 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 929 (median 217).

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 captures the top 25% of strings and bi-grams by frequency for each firm, 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.

[6] Due to this criterion, we lose over 150 firms when we implement the algorithm because there are too few transitions between jobs in the resume data.

[7] We need to start in 2017 because a transition is measured relative to the previous year.

This process results in a meaningful reduction in the number of titles, now classified as roles, with an average of 259 roles (median 99) per firm. Each role encompasses approximately three raw titles (2.3 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 availability of the company's website (URL), which is reported in both datasets. 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 7,277 firms to the resume site data (across all years and industries) before filters and additional data cleaning. These firms have over 16 million unique workers in all years (again, before cleaning). After filtering from data cleaning and implementing the algorithm, we are left with nearly seven 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 (retail, 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. With these criteria, we are left with 3,128 firms (though the sample size varies depending on the regression). Over 2,000 firms outside the filtered industry codes lack a hierarchical layer measure, have no employees listed on the resume website, or have insufficient coverage. 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). These are all characteristics of large firms, which confirms that our sample tilts toward larger companies.

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 149 ILMs. While this may seem large, this count is driven by rare one-off titles or job cycles that are typical in the user-inputted LinkedIn data. The hierarchical layers measure focuses on the largest ILM, which contains 129 unique roles on average (recall that titles are mapped to roles). On average, there are 196 roles in a firm.

Two statistics provide comfort that the largest ILM captures most roles and workers. Specifically, for the median firm, 60% of the roles and 73% of the worker-years are found in the largest ILM. That is, nearly three-quarters of the workers 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 49% of the number of Compustat employees for the median firm. While this is an inherent limitation of the LinkedIn data, this is less of 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 9.5 (median 8). This exceeds the number of layers—typically 4—obtained in the studies reviewed in Appendix Table A1.[8] 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

[8]The exception is Adhvaryu, Bassi, Nyshadham, Tamayo, and Torres (2023), who identify 16 layers in their study of an Argentinean automotive working group.

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 (2024) that captures the relative share of employees in each layer.[9] 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 ONet occupation codes, respectively). The third measure uses our definition of layers. The average of 0.75 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 (2024), 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 average and median number of layers for U.S. public firms has declined over the sample period. Thus, the flattening of the firm—first documented in Rajan and Wulf (2006) and Guadalupe and Wulf (2010) for the top of the hierarchies—applies to later years among a much larger sample of public firms, taking 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 flattening of their hierarchies, in keeping with the broader evidence from Figure 1.

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 2016–2023, only 18% of firms experienced a change in layers,

[9] 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.

and this occurred in 3.9% of firm-years. When layers do change, they increase three-quarters of the time, and the median change is a one-layer increase. This stability limits our ability to study within-firm changes and thus leaves much of what follows as cross-sectional results.[10]

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.

In Table 3, we report the most commonly occurring titles in the top and bottom layers of the hierarchy. While firms with different total layers are difficult to compare, it is reassuring to see that the titles closely capture those that would naturally be associated with the top (e.g., "Vice President," "Director") and bottom (e.g., "Sales Associate," "Driver," "Cashier") levels of the hierarchy.

Relatedly, Appendix Table A2 shows that titles that include the strings "director" and "president" are increasingly more prevalent higher up in the hierarchy, while titles that include the strings "sales" and "engineer" are increasingly more prevalent in lower layers. The hierarchical layer's rank also positively correlates with other measures of levels. Appendix Figure A1 reports the average ranking for titles by layer after mapping the titles to the ONet classification system.[11] As can be seen, we observe a weakly positive relationship between a title's hierarchical layer and the one to five ranking in ONet.

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

[10]Note that we may still detect meaningful changes in the aggregate, given the changing number of public firms over the sample period.

[11]In Appendix Figure A1, 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. These cutoffs correspond to the 25th and 75th percentiles of the distribution of the number of layers across all firms in our sample.

their first job—increases with the hierarchical layer. 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.

The theory literature reviewed in Section 2 guides the remainder of the validation tests 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 2023 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 connection between firm size—measured in eight different ways (using total assets, employment, revenues, the number of business segments, the number of U.S. subsidiaries, the number of executives, board size, as well as firm age)—and the number of layers. 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 can be seen, all size proxies strongly correlate with a firm's layer count. Indeed, the R-squared in several regressions exceeds 70%. The fit is highest for employment, as theory would predict—as firms add employees (without a commensurate increase in output), they are more likely to alter their organizational structure through additional layers. Importantly, the regression in the last column, which includes all size controls, shows that size does not fully explain the variation in hierarchical layers. That is, there is enough variation in hierarchies that cannot be explained by size alone. This is the variation that 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. Moreover, this strong association motivates the inclusion of size controls in all 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 will 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).[12]

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 employees in a layer, or in a layer quantile, measured using the number of employees in our LinkedIn data. As discussed in Section 5.1, LinkedIn does not capture all employees, and we speculate that LinkedIn over-samples middle to upper-middle management. Such over-sampling does not impact the hierarchy variable construction per se, but does impact employee shares. Thus, estimating firm shapes entails some measurement error.

There are two ways in which we can compare firms' shapes. We can condition on the number of layers, or we can split the firm hierarchy into quintiles (or similar) to pool firms in one analysis. Figure 5 follows the second approach, showing the share by hierarchy quintile for firms with at least five layers. In each industry sector, we observe a monotonic decrease in the share of employees as one moves into higher layers. Figure 6 follows the first approach, reporting the average shape of firms with five (Panel A) and eight layers (Panel B) where the shares are averaged across all firm-years. As can be seen, the evidence is again generally consistent with a pyramidal shape.

Finally, an indirect view of a firm's shape can be obtained from how workers move between layers. Promotion in our data is a change in a worker's role to a new one with a higher rank. Recall that layers are collections of roles, and both layers and roles are ranked. In a pyramidal organization where workers can be promoted, there are fewer positions at the next (higher-ranked) layer than at the current layer. This should result in a longer time to promotion

[12]Relatedly, Hart and Moore (2005) shows that the pyramidal form is optimal.

higher up in the hierarchy (along with any additional experience requirements as a worker moves up the firm). This is indeed what we observe in Appendix Figure A2, which shows that the time to next-layer promotions tends to increase as we move up the firm's hierarchy.

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 the string "manager" in their title for groups of firms that have the same number of layers. Conditioning on the number of layers ensures we are comparing the same layers across firms. As can be seen, the share of managers tends to increase as we move up the firm's hierarchy.[13] 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 Additional validation

In the Appendix, we provide additional validation. In Appendix A1, we show that our measure of hierarchical layers performs well when benchmarked against a small set of firms with known hierarchical layers. In Appendix A2, we provide support for the algorithm's assumptions about the nature of job transitions.

6.2.6 Summary

In sum, the validation tests provided in this section indicate that publicly available resume data can be used to build meaningful proxies for corporate hierarchies consistent with a large body of theoretical work on organizations. In Sections 7 and 8, we use our measure to explore the role of hierarchies for corporate decision-making and a wide range of outcomes.

[13]Note that the pattern is not monotonic. Toward the top of the hierarchy, the share of managers decreases, which likely reflects the different terminology used to refer to upper-echelon managerial roles (e.g., "director," "president").

6.3 Title measurement issues

Before moving to the empirical analysis, we discuss potential measurement issues in job titles. 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 (2023) 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 we found the opposite pattern in Figure 1. 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.

7 Corporate hierarchy and firm outcomes

Now that we have demonstrated that our measure of hierarchical layers captures expected patterns in firm organization, we ask 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.[14] When estimating these regressions, we cluster standard errors at 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 firms with more layers to have a more educated workforce. To examine whether this is the case, we use the educational histories of employees from their online resumes.[15]

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

[14] 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.

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

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 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.[16]

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.[17] 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, in keeping with our finding of higher internal promotion rates and tenure, we find that high-layer firms are more likely to appoint internal CEOs. Turning to CEO compensation, we find that CEOs of high-layer firms have higher compensation. Moreover, multi-layer firms have 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). These findings are consistent with several models of corporate hierarchy. In particular, they are consistent with the tournament-type models, in which higher rewards are provided

[16]In the Appendix, we complement these findings in two ways. First, Appendix Figure A3 plots the average tenure in a given layer quintile for groups of firms that have the same total number of layers. As can be seen, the average tenure within a layer quintile tends to increase as we move up the hierarchy. Moreover, firms with more layers tend to have higher tenure by layer across the hierarchy. Second, Appendix Table A3 provides additional perspective on the hiring and promotion patterns. Specifically, we consider all role switches and all new external hires as the set of employee-year observations. We then show that, within a firm, the probability of external hiring is monotonically decreasing as we move up the hierarchy. (The estimates capture the probability of an external hire in a given layer quintile relative to the lowest layer quintile as base group.) This pattern is consistent with the predictions of Baker, Gibbs, and Holmstrom (1994) who argue that, as higher layers demand more firm-specific human capital, external hires are more difficult to find.

[17]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.

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).

7.2 Firm productivity and costs

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 Zandt, 2012; Caliendo, 2020).

The left-hand panel of Table 6 explores the relationship between the number of layers and four measures of productivity. Columns 1-3 show that an increase in the number of layers is associated with a higher return on assets (ROA), gross profit margin, and net profit margin. These findings indicate that multi-layer firms achieve higher operating performance. Interestingly, column 4 shows that high-layer firms have a 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.

The right-hand panel of Table 6 explores cost measures (scaled by total assets). In column 5, we consider selling, general, and administrative expenses (SG&A). SG&A is often considered a measure of organizational capital. As can be seen, SG&A is substantially higher for high-layer firms, consistent with the higher cost of maintaining more complex hierarchies. Finally, in column 6, we examine wages. The point estimate is positive and significant, indicating that firms with more hierarchical layers have more expensive human capital. This again 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.

Next, we examine two major types of firm investment: acquisitions and innovation.

7.3 Acquisitions

One of the most significant investments that public firms make over their life cycle are acquisitions of other firms. Empirical evidence suggests that acquisitions are often value-destroying (e.g., Alexandridis, Antypas, and Travlos, 2017), while there is mixed evidence on the value gains from diversifying acquisitions (e.g., Eckbo, 2014).

Why would we expect a firm's hierarchical structure to matter for the firm's acquisition activity? To our knowledge, no model of hierarchies offers direct prediction about acquisition activities. From a theoretical perspective, the prediction could go either way. On 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.

Table 7 studies whether our measure of hierarchical layers predicts acquisition activity and the type of acquisitions. We measure acquisition activity using SDC Platinum and the Compustat-SDC merge from Ewens, Peters, and Wang (2024). In column 1, the dependent variable is an indicator variable equal to one if the firm had at least one acquisition during the year. As can be seen, firms with more layers are significantly more likely to engage in acquisitions. In column 2, we obtain similar results if we only consider "large" deals (that is, deals with reported valuations). These results also hold if we consider the count of acquisitions (column 3) and the sum of the (reported) valuations across all acquisitions made by the firm (column 4). Overall, these results indicate that high-layer firms are significantly more active in the acquisition market.

Column 5 of Table 7 examines whether the type of acquisition differs. Following Gormley and Matsa (2016), we define a diversifying acquisition as one where none of the target's SIC code matches the acquirer's primary SIC code. As can be seen, we find no evidence that high-layer firms differ along this dimension. This indicates that, although high-layer firms are more likely to engage in acquisitions, they do not favor diversifying acquisitions, as those might be harder to integrate into a multi-layered structure.

Overall, we conclude from these results that the number of hierarchical layers has meaningful predictive power for acquisition activity.

7.4 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 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 8 provides regression estimates of patenting activity and quality. Investments in innovation often start with R&D. Column 1 shows that firms with more layers invest more in R&D (as a ratio of total assets), although the point estimate is not significant at conventional levels. The Poisson estimate in column 2 shows that this higher R&D investment corresponds with a significant increase in the number of 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.[18] 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.[19]

Overall, these innovation outcomes show 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.

7.5 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 and adjust to unexpected shocks.

[18] 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 (2019) for more details.

[19] 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.

This, in turn, would reduce the volatility of the firms' cash flows.

In columns 1 and 2 of of Table 9, 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.[20] 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 has little predicting power for these variables. If at all, the evidence in column 5 suggests that high-layer firms use higher leverage at the margin (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.

7.6 Stock returns

So far, our analysis shows that high-layer firms have higher operating performance (Section 7.2) and lower risk (Section 7.5). Another dimension of firm performance is the firm's stock returns. To examine the stock market performance of high- vs. low-layer firms, we build long-short portfolios using the hierarchical layers as the sorting characteristic for 2016 to 2023. Each year, firms are sorted, and the portfolio goes long in the top quartile and short in the bottom quartile. The portfolio is rebalanced each year. Table 10 reports the excess returns (alpha) that are achieved by this trading strategy under different asset pricing models. As can be seen, we find evidence for an insignificant alpha in all columns.

The insignificant alpha indicates that stock market investors correctly anticipate the benefits and costs of having multiple layers.[21] Leaving the alpha aside, a look at the factor loadings

[20] 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.

[21] Consistent with this interpretation, when looking at analysts' EPS (earnings per share) forecasts from the IBES (Institutional Brokers' Estimate System) database, we find no evidence that analysts systematically over- or under-estimate the performance of high- vs. low-layer firms.

is informative. In particular, the profitability factor loads positively, which is consistent with our profitability results (Table 6). Moreover, the full 6-factor model in column 6 shows that high-layer stocks behave like large-cap, "quality" stocks. The latter feature signals stability and resiliency, which echoes well with what we found in the volatility analysis (Table 9). While we are unaware of organizational theories tied to these results, it is reassuring to see that they match the correlational tests above.

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 Corporate life cycle

As a starting point, 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 (2024), 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.[22]

[22] We caution that these results are correlational due to the endogenous nature of the decision to go public.

8.2 The Covid pandemic

The Covid pandemic provides a source of variation for studying how firms respond to shocks. Consider the set of firms in the pharmaceutical industry. For these firms, the pandemic represents a twofold shock. First, they were forced to solve the vaccine problem for a novel disease in a rapid amount of time. Second, they likely anticipated a huge spike in demand (conditional on a successful vaccine).

In the models of Garicano (2000) and Caliendo, Monte, and Rossi-Hansberg (2015), firms respond to demand shocks and shocks to "complexity" by adding layers. The rationale is that additional layers add managers (problem solvers), allowing lower-level workers to narrow the knowledge they must master in production. This, in turn, increases the organization's ability to address higher demand and the higher complexity of their environment. Accordingly, through the lens of these models, we would expect pharmaceutical firms to add layers following the Covid shock.

To study this question, we consider the set of firms in the 4-digit SIC code 2834 (that is, firms that manufacture, fabricate, or process drugs in pharmaceutical preparations—in short, pharmaceutical firms). The dependent variable is the current year's log number of hierarchical layers. The analysis considers firm-year observations from 2017-2022 (that is, around the Covid shock). As control firms, we consider either (i) all firms outside SIC code 2834 or (ii) all firms outside SIC code 2834 and outside the so-called "Covid industries" that were directly impacted by Covid such as lodging or cruise ships (the list of Covid industries is obtained from Fahlenbrach, Rageth, and Stulz, 2021). We then estimate a difference-in-differences specification in which "Covid" and "post" are indicator variables equal to one for the treated firms (SIC code 2834) and the post-Covid years (that is, 2020-2022), respectively. The specification further includes firm and year fixed effects. The firm fixed effects are important since we are interested in within-firm changes in hierarchical layers around the Covid shock. As we noted in Section 6.1.1, the number of layers has limited within-firm variation. Hence, a caveat of including firm fixed effects is that it reduces the power of our tests. Nevertheless, given the magnitude of the Covid shock for pharmaceutical firms, we might still detect significant changes in the number of layers around that specific shock.[23]

Table 11 presents the results. Column 1 shows that the number of layers increased by about 4% among the treated firms in the post-Covid years. Column 2 introduces a dynamic

[23]Note that, since the treatment assignment is at the industry level, we cluster standard errors along this dimension (4-digit SIC).

specification. As can be seen, the response starts in 2020 and is stronger in 2021-2022, when most vaccines were released. In columns 3 and 4, we find that the results are similar if we include size decile fixed effects. Finally, in columns 5 and 6, our results continue to hold if we exclude firms in the Covid industries from the control group. Note that, in all specifications, the pre-treatment coefficients are small and insignificant, which mitigates the possibility of pre-trends that could confound our estimates.[24]

Overall, the evidence from Table 11 indicates that firms respond to demand and complexity shocks by adding hierarchical layers.[25] These responses further suggest that a firm's hierarchical structure is shaped by the shocks the firm has faced.

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 Zandt, 2012). 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).

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 12. We consider four different metrics of AI adoption: the share of AI job postings, broadly (column 1) and narrowly (column 2) defined, the average

[24]The only exception is the 2017 pre-treatment coefficient in column 2. However, this estimate is not robust. Once we account for size decile fixed effects in column 4 and the finer control group in column 6, the corresponding estimate is small and insignificant.

[25]One potential concern with this analysis is that, as pharmaceutical firms hire more employees to accommodate the demand for vaccines, this might mechanically trigger a higher number of layers. However, when we re-run the regressions in Table 11 controlling for contemporaneous changes in employment, we find that the point estimates remain large and statistically significant. That is, the increase in the number of layers is not merely reflecting an increase in the firms' workforce.

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 reduction in the number of layers in all four columns, as predicted by theory. Although the tests are under-powered, we find that the point estimates are significant at the 10% level regardless of the metric of AI adoption.[26]

To obtain more within-firm variation in the number of layers, we also ran a specification using long differences, that is, regressing the change in log(# layers) for a given firm over all available years on the corresponding change in AI (as well as year and size decile fixed effects). Doing so yields estimates that are similar to those in Table 12.

9 Conclusion

We propose a measure of corporate hierarchy for U.S. public firms that is obtained from online resumes of seven 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 a more educated workforce, higher internal promotion rates, and longer employee tenure. Their operating performance is higher, but they face higher administrative costs. They are more active acquirers and produce more patents, but not higher-quality patents. They exhibit lower stock return volatility and have lower business risk (i.e., operating asset volatility). Turning to the determinants of corporate hierarchies, we find that pharmaceutical companies increased their number of layers following the Covid-19 pandemic, while companies flattened their hierarchies following the adoption of AI technologies. These findings are consistent with the theoretical predictions of existing models of corporate hierarchies, especially those that conceptualize hierarchical layers as a way to bring in "problem solvers" whose knowledge helps solve the more complex problems faced by lower-layer workers.

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

[26] The evidence in Table 12 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.

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 exciting topic.[27]

[27] 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.

Alexandridis, George, Nikolaos Antypas, and Nickolaos Travlos, 2017, Value creation from M&As: New evidence, *Journal of Corporate Finance* 45, 632–650.

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, 2019, 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, 2024, Going public and the internal organization of the firm, Working Paper.

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, 2020, Productivity and organization in Portuguese firms, *Journal of Political Economy* 128, 4211–4257.

——— , 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, 2023, Too many managers: The strategic use of titles to avoid overtime payments, NBER Working Paper 30826.

Eckbo, B. Espen, 2014, Corporate takeovers and economic efficiency, *Annual Review of Financial Economics* 6, 51–74.

Ewens, Michael, and Matt Marx, 2024, Startup patents, Working Paper.

Ewens, Michael, Ryan H. Peters, and Sean Wang, 2024, 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.

Fama, Eugene F., and Kenneth R. French, 2018, Choosing factors, *Journal of Financial Economics* 128, 234–252.

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, 2012, 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.

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, 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, 2018, 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.

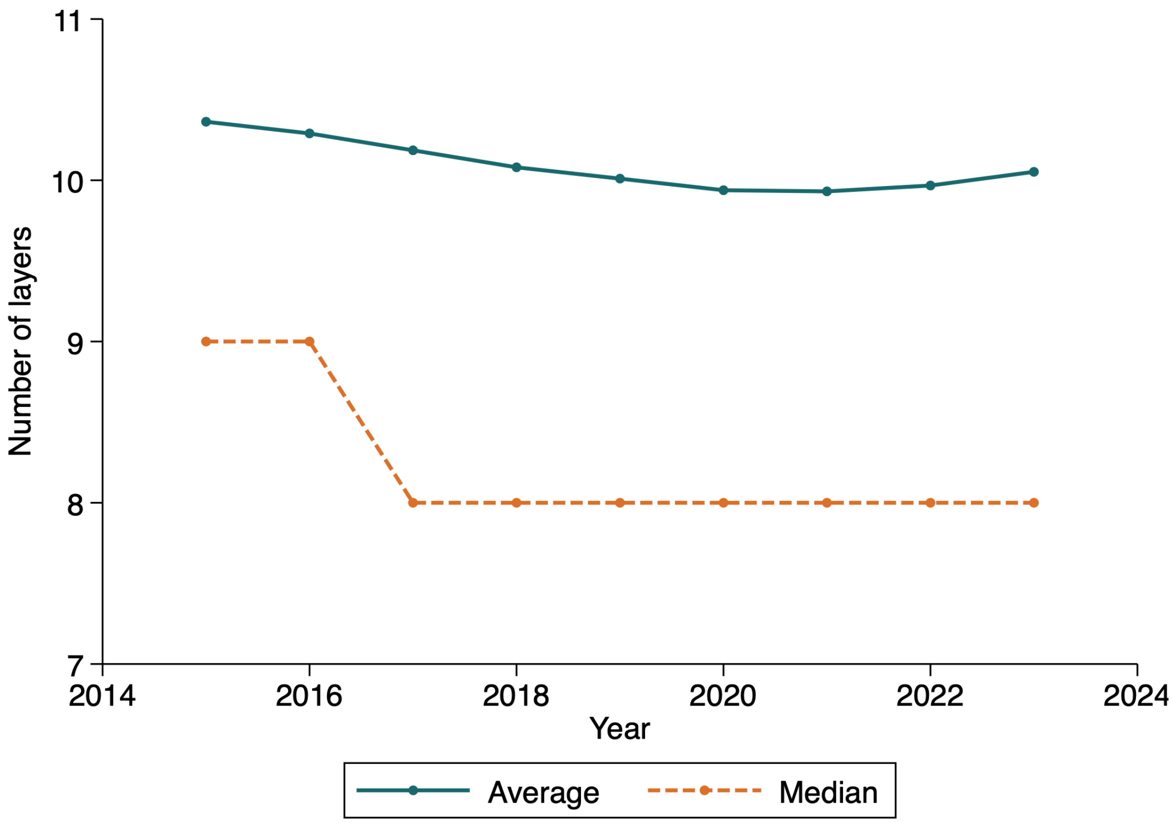

Figure 2: Organizational layers by industry

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

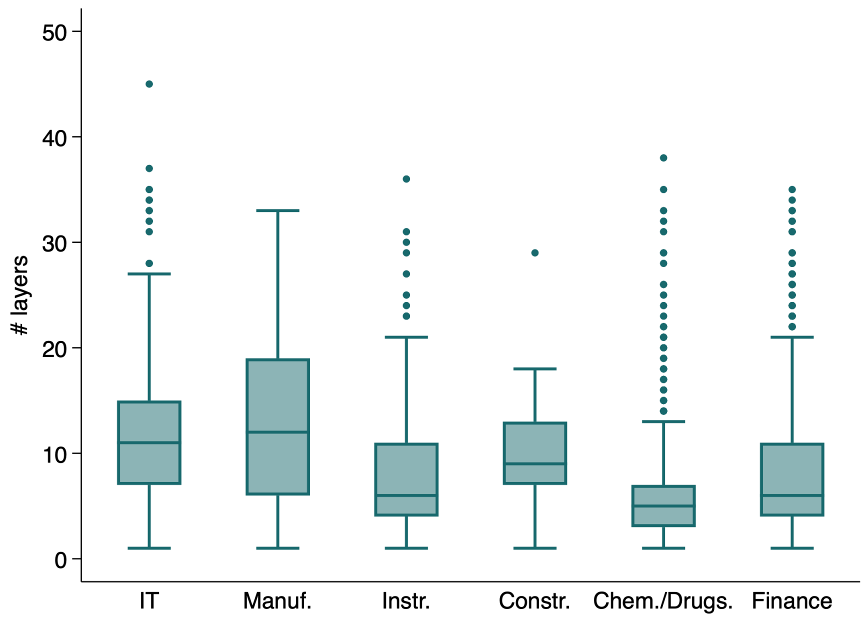

Panel B: Hierarchical layers by major industry over time

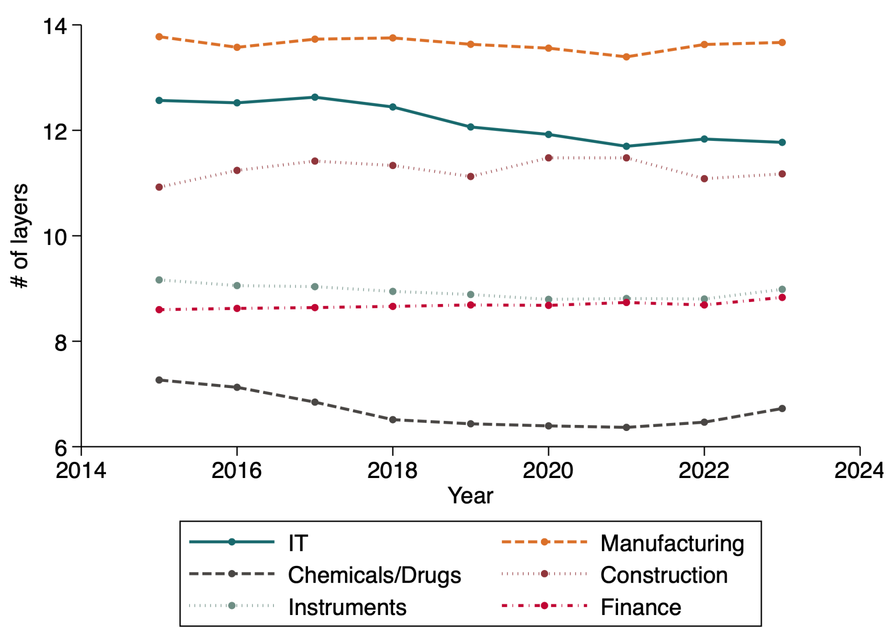

Figure 3: Work experience across the firm hierarchy

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 corresponds to the interquartile range across all firms.

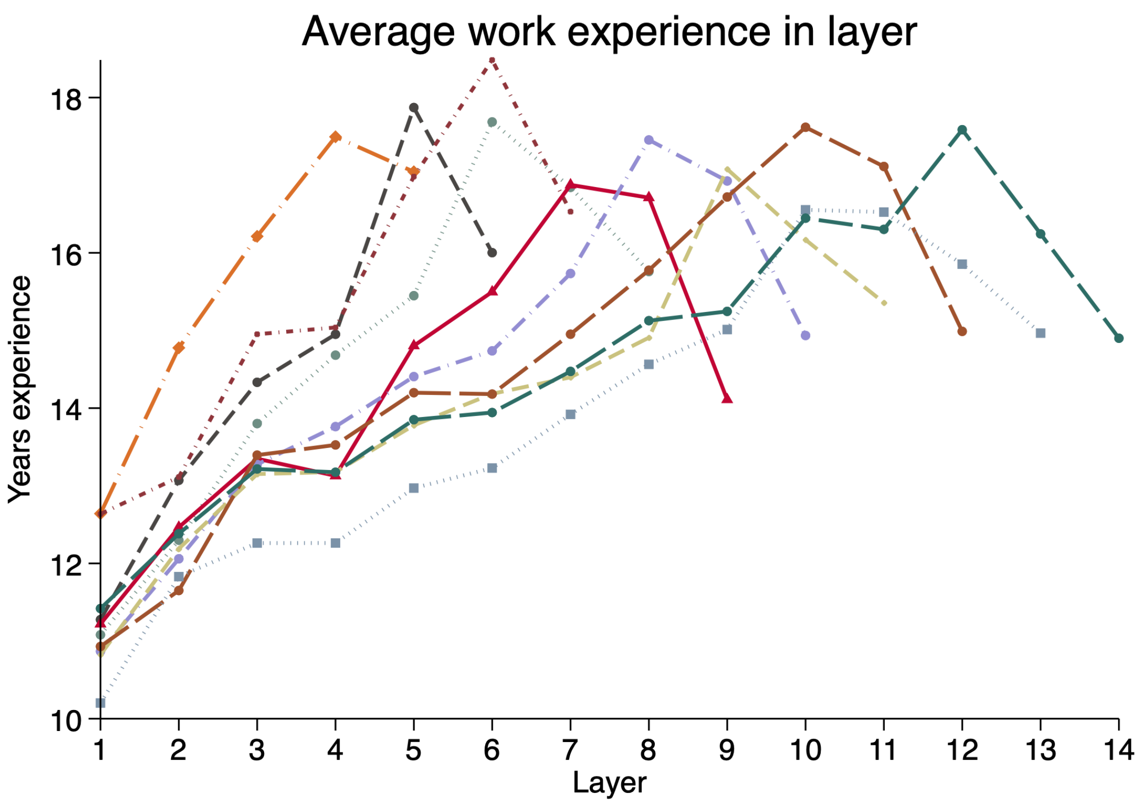

Figure 4: Variation in organizational layers by size

This figure displays the box plot of the number of hierarchical layers by size quintile for the sample of manufacturing firms in 2023. 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

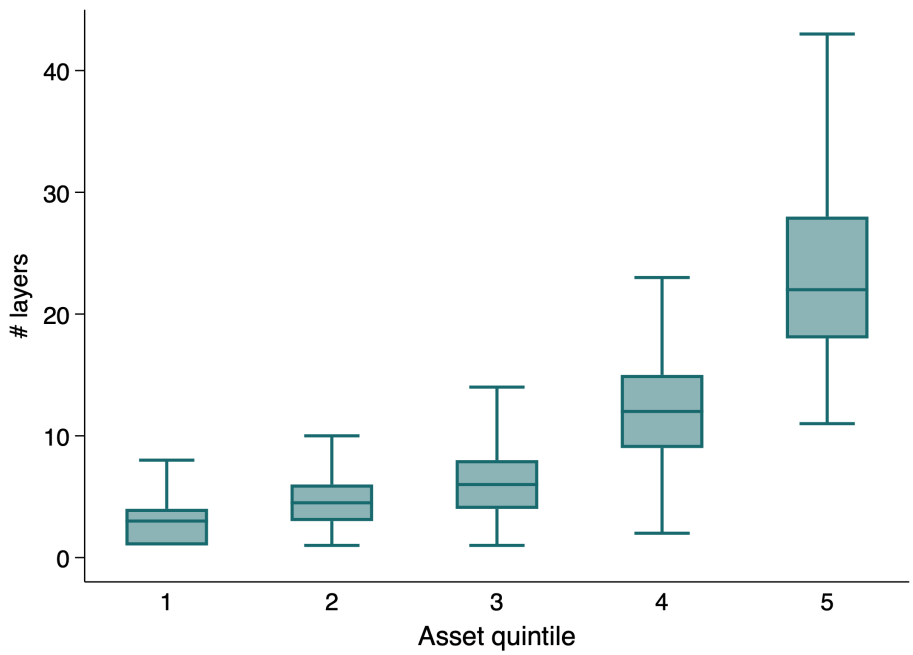

Panel B: Employee count quintiles

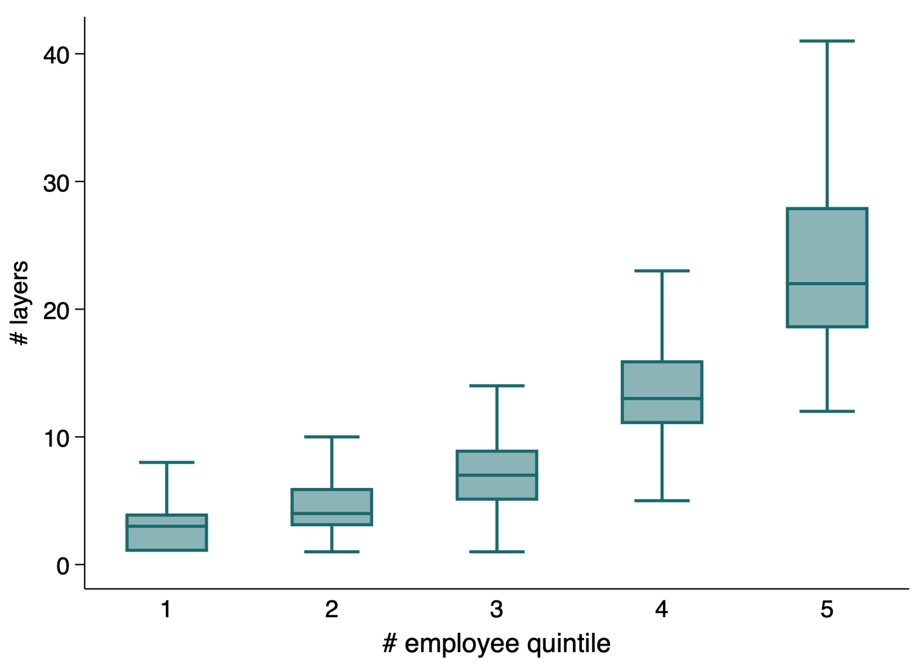

Figure 5: Share of employees by hierarchy layer quintile

This figure reports the average share of employees in each layer quintile across firms for the major industries in the sample. Only firms with at least five layers are considered. The share of employees is defined as the unique number of employees in the LinkedIn data whose title is associated with each layer, divided by the total number of employees on LinkedIn. The figure reports the average across all relevant firm-year observations.

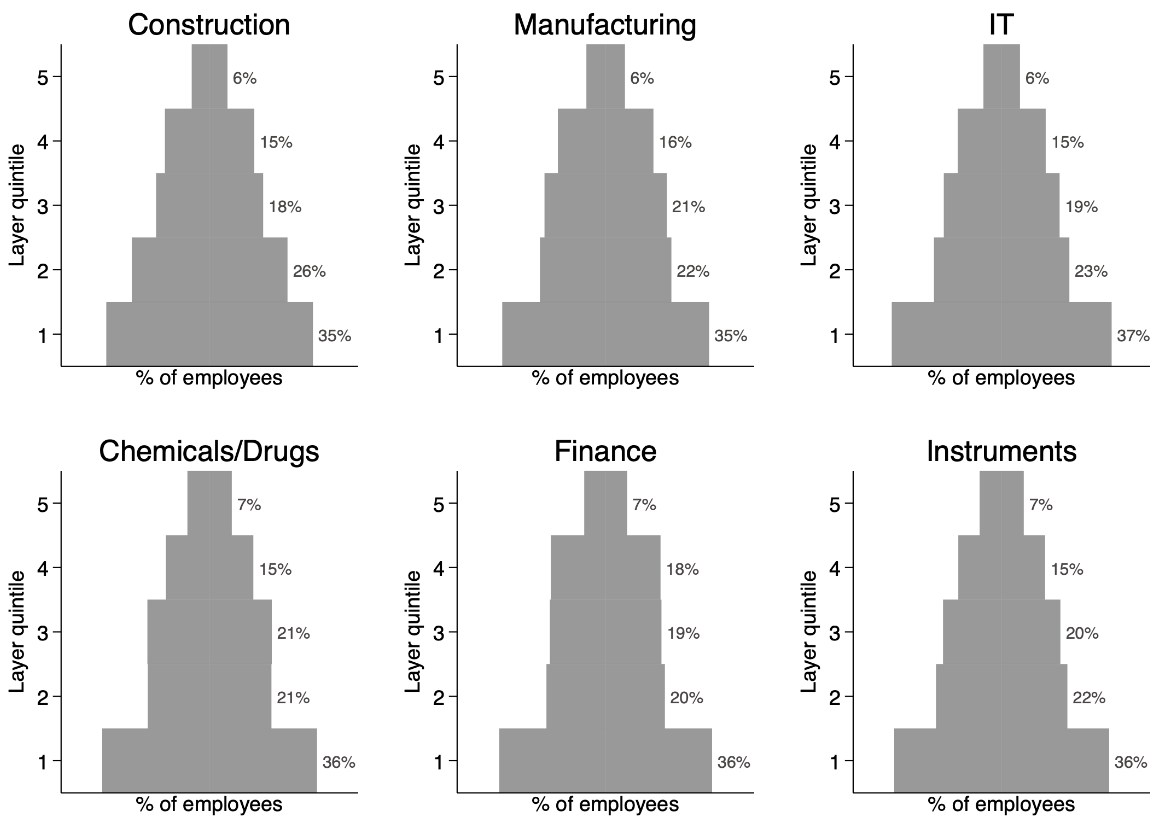

Figure 6: Share of employees by hierarchy layer

This figure reports the average share of employees in each layer for firms with five (Panel A) and eight (Panel B) layers. The share of employees is defined as the unique number of employees in the LinkedIn data whose title is associated with each layer, divided by the total number of employees on LinkedIn. The figure reports the average across all relevant firm-year observations.

Panel A: Firms with five layers

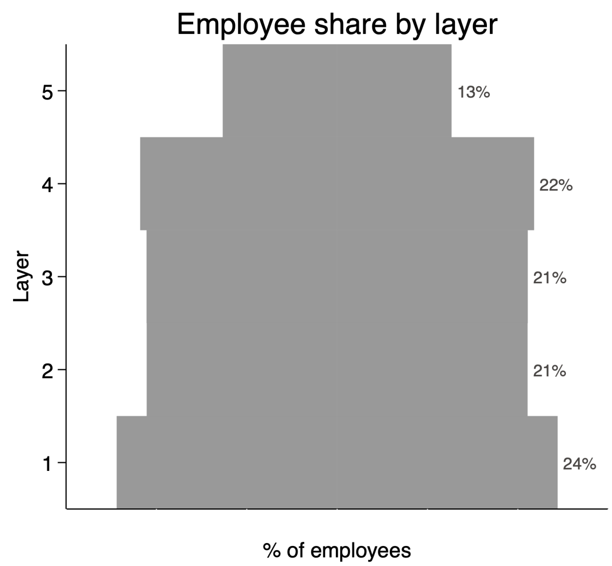

Panel B: Firms with eight layers

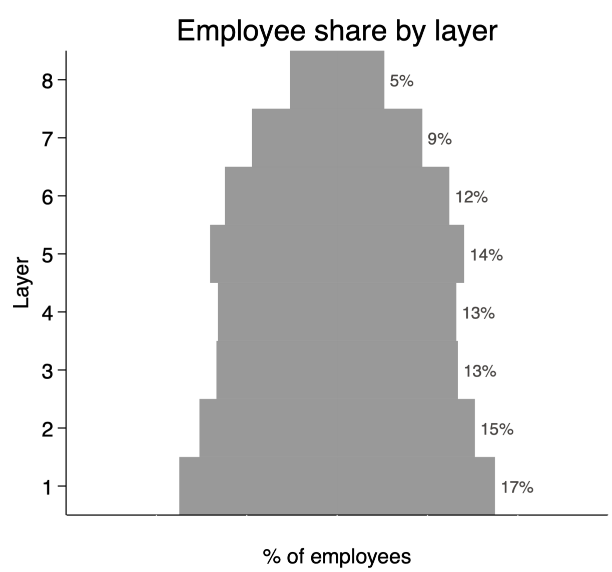

Figure 7: Managers by hierarchical layer

This figure reports the fraction of employees with the string "manager" in their title 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 corresponds to the interquartile range across all firms.

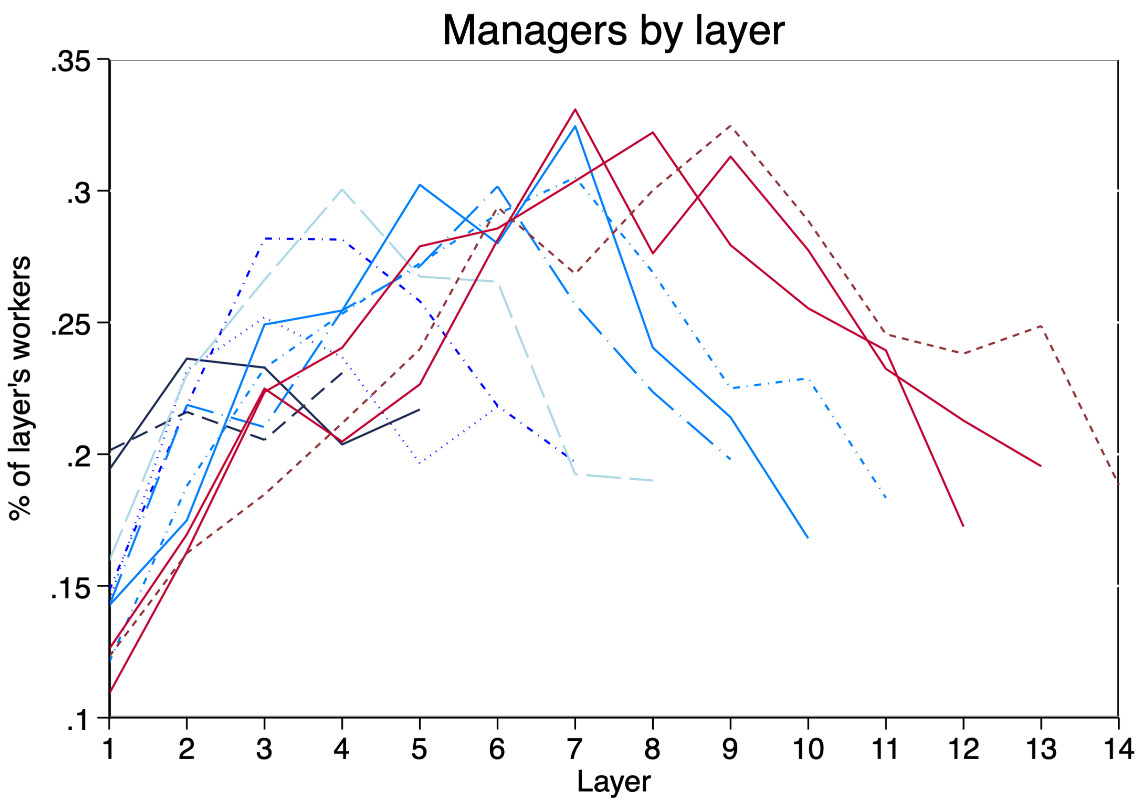

Figure 8: Advanced degrees by hierarchical layer

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 in our sample whose number of layers is between 4 and 14, which corresponds to the interquartile range across all firms.

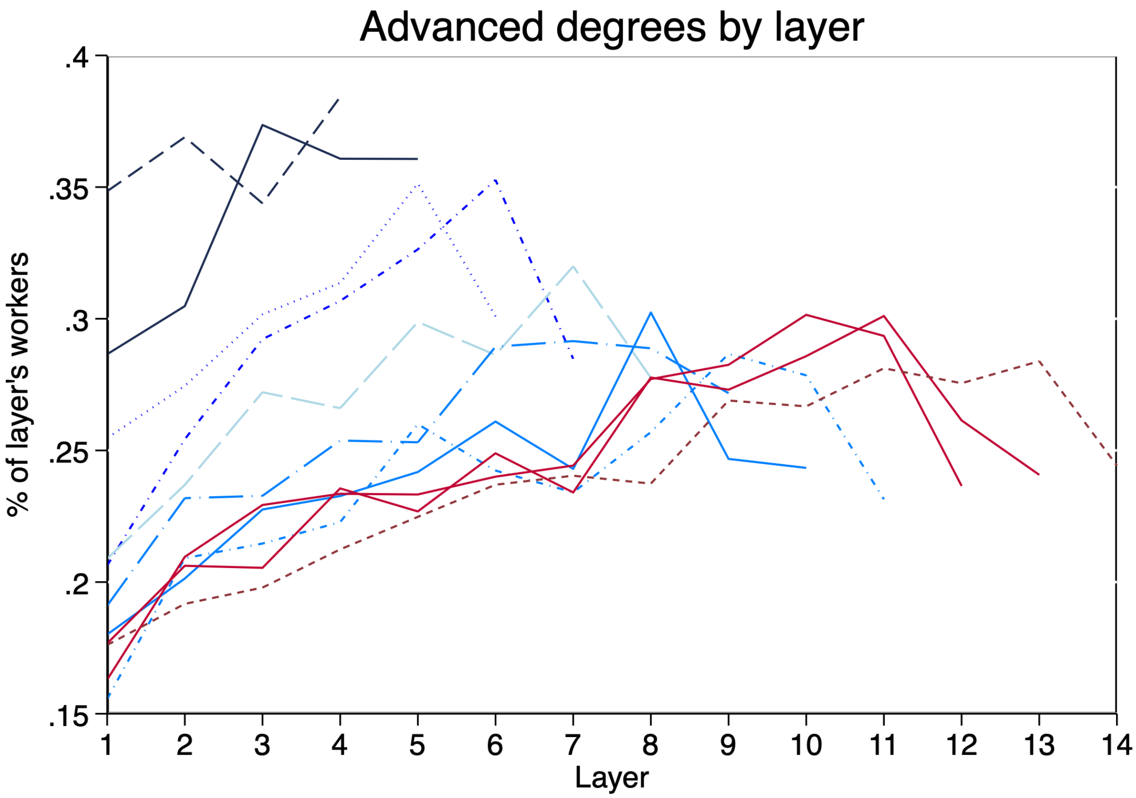

Figure 9: Change in the number of hierarchical layers around the IPO

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.

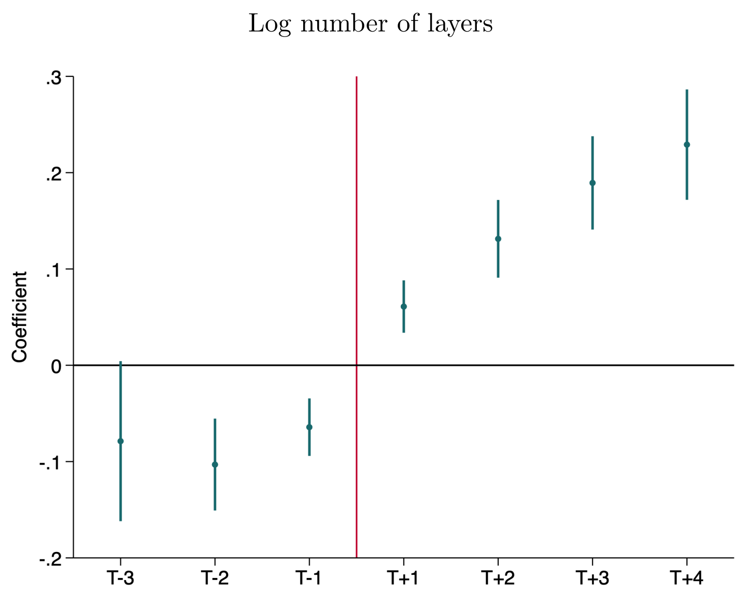

Table 1: Comparison of firms in and out of the sample

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 recoded once in 2015 or in the first available year during our sample period (2016–2023). 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	16626.42	5166.11	11460.30***	(2471.57)	7,147
Total liabilities	13855.65	3823.59	10032.05***	(2300.21)	7,138
PPE (net)	1413.81	811.69	602.12***	(183.41)	6,980
Debt / equity	2.37	1.79	0.58***	(0.09)	7,127
Cash	1581.45	400.10	1181.36***	(388.86)	6,524
Cash / assets	0.25	0.25	0.00	(0.01)	6,480
Revenue (m)	4114.69	2461.89	1652.79***	(429.23)	7,094
Net income	84.36	46.48	37.88***	(6.17)	6,474
EBIT	636.69	297.26	339.42***	(100.73)	6,472
Net income / revenue	-0.08	-0.10	0.02*	(0.01)	4,879
ROA	-0.13	-0.36	0.23***	(0.02)	6,265
Firms	3,128	4,019			

Table 2: Summary statistics

This table reports summary statistics for the hierarchy variables. "Number of internal labor markets" is the count of internal labor markets (ILMs) within the firm that are identified by the algorithm described in Section 4. "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 (2024), computed using layers based on the occupation codes of the German Classification of Occupations, the occupation codes of ONet, 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 internal labor markets	148.69	168.61	1.00	99.00	2403.00	3,128
Size of largest ILM (roles)	129.27	202.54	1.00	53.00	1916.00	3,128
Number of unique roles	196.12	239.65	2.00	116.00	2066.00	3,128
% role in largest ILM	57.98	20.12	3.85	59.90	100.00	3,128
% worker-years in largest ILM	67.09	24.16	1.34	73.43	98.31	2,961
Employees (Compustat, 1000s)	6.82	35.51	0.00	0.71	1525.00	2,982
Number employees (LinkedIn, '000s)	2.02	7.06	0.00	0.35	209.95	3,128
Number hierarchy layers	9.54	6.99	1.00	8.00	45.00	3,128
Hierarchization (German)	0.61	0.06	0.00	0.62	0.75	3,128
Hierarchization (ONet)	0.66	0.06	0.00	0.67	0.80	3,128
Hierarchization (this paper)	0.75	0.22	0.00	0.83	0.96	3,128
	Panel B: Firm-year					
	Mean	SD	Min	Median	Max	N
Number of internal labor markets	155.01	173.86	1.00	105.00	2403.00	21,954
Size of largest ILM (roles)	143.63	217.08	1.00	62.00	1916.00	21,954
Number of unique roles	212.55	255.30	2.00	127.00	2066.00	21,954
% role in largest ILM	60.70	19.43	3.45	62.77	100.00	21,954
% worker-years in largest ILM	69.19	23.08	1.34	75.42	98.31	20,891
Employees (Compustat, 1000s)	7.53	33.47	0.00	0.78	1608.00	19,966
Number employees (LinkedIn, '000s)	2.41	8.18	0.00	0.38	270.92	21,954
Number hierarchy layers	10.09	7.28	1.00	8.00	45.00	21,954
Hierarchization (German)	0.61	0.06	0.00	0.62	0.75	21,954
Hierarchization (ONet)	0.66	0.06	0.00	0.67	0.80	21,954
Hierarchization (this paper)	0.76	0.21	0.00	0.83	0.96	21,954

Table 3: Common titles by hierarchy layer

This table reports the most common titles across all firms and years in their top layer (left-hand panel) and bottom layer (right-hand panel).

Top layer	**Bottom layer**
Vice President	Sales Associate
Product Manager	Driver
Director	Customer Service Representative
Data Engineer	Associate
Product Owner	Warehouse Associate
Managing Director	Personal Banker
Senior Software Engineer	Fulfillment Associate
Senior Product Manager	Package Handler
Project Manager	Bank Teller
Software Engineer	Cashier

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 year since the firm went public; "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); "# execs in Execucomp" is the count of executives in Execucomp; "board size" is the size of the board from BoardEx. 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	0.28***								-0.0031
	(0.0042)								(0.016)
Log employees (CS, 1000s)		0.33***							0.21***
		(0.0042)							(0.022)
Log revenue			0.26***						0.071***
			(0.0046)						(0.022)
Log firm age (yrs.)				0.23***					0.0041
				(0.020)					(0.022)
Log business segments					0.27***				-0.019
					(0.025)				(0.018)
Log (# US subsidiaries)						0.16***			-0.0097
						(0.0084)			(0.0063)
Log # execs in Execucomp							0.47***		0.054
							(0.049)		(0.034)
Log board size								1.22***	0.16***
								(0.048)	(0.048)
Observations	18,493	16,827	17,460	10,474	18,317	18,510	7,882	15,359	4,319
# firms	2,984	2,865	2,857	1,713	2,954	2,987	1,113	2,662	632
Mean dep. var.	2.04	2.09	2.10	2.11	2.05	2.04	2.53	2.11	2.48
R^2	0.71	0.80	0.70	0.43	0.32	0.35	0.42	0.46	0.82
Within-R^2	0.60	0.72	0.60	0.08	0.04	0.08	0.05	0.23	0.72
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 column 1, the dependent variable is the fraction of new LinkedIn employees during the year ("% new"). In column 2, it is the fraction of existing LinkedIn employees who are promoted during the year ("% promote"). In columns 3 and 4, it is the logarithm of the average tenure of the firm's employees (in year). 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. 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.

	Employees		Tenure (log years)		ESOP?	Internal	CEO	CEO
	% new	% promote	Non-new	All		CEO	pay	pay ratio
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log # layers ($t-1$)	-0.046***	0.048***	0.053**	0.086***	0.059***	0.18***	0.37***	33.8***
	(0.0058)	(0.0052)	(0.024)	(0.026)	(0.019)	(0.053)	(0.051)	(4.73)
Observations	9,023	9,020	9,020	9,023	8,845	749	7,989	6,162
# firms	1,354	1,354	1,354	1,354	1,685	571	1,125	1,554
Mean dep. var.	0.14	0.18	1.72	1.63	0.14	0.59	8.60	101.54
R^2	0.04	0.05	0.03	0.04	0.09	0.02	0.37	0.33
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: Firm productivity and costs

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"); "Gross margin" is the ratio of gross profits (item "gp") to total revenue (item "revt"); "Net margin" is the ratio of net income (item "ni") to total revenue; "Rev/emp." is the ratio of total revenue to the number of employees (item "emp"). The dependent variables in columns 5 and 6 are cost measures. "SG&A" is the ratio of selling, general & administrative expense (item "xsga") to total assets; and "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 (1)	Gross margin (2)	Net margin (3)	Rev/emp. (4)	SG&A (5)	Wages (6)
log # layers $(t-1)$	0.0227***	0.0700***	0.339***	-171.9***	0.0804***	0.0146**
	(0.00679)	(0.0253)	(0.0705)	(21.95)	(0.00930)	(0.00574)
Observations	15,670	15,473	15,473	17,323	14,080	4,062
# firms	2,670	2,883	2,883	2,882	2,658	625
Mean dep. var.	-0.02	0.27	-0.69	585.90	0.31	0.08
R^2	0.58	0.46	0.44	0.64	0.59	0.75
Within-R^2	0.25	0.04	0.11	0.08	0.31	0.11
Size decile FE?	Y	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y	Y

Table 7: 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 ($t-1$)	0.010*	0.0083**	0.21*	0.33**	0.043
	(0.0055)	(0.0042)	(0.11)	(0.14)	(0.042)
Observations	18,489	18,489	11,656	998	1,689
# firms	3,009	3,009	1,958	511	657
Mean dep. var.	0.09	0.06	0.24	4.80	0.62
R^2	0.02	0.01		0.37	0.01
Pseudo-R^2			0.17		
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 8: 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 ($t-1$)	0.0053	2.44***	0.19***	0.29	0.13*
	(0.0069)	(0.59)	(0.032)	(0.19)	(0.070)
Observations	10,325	13,479	16,462	2,597	2,650
# firms	1,803	2,447	3,000	587	607
Mean dep. var.	0.15	16.10	0.48	0.97	0.54
Pseudo-R^2		0.86		0.32	
R^2	0.52		0.39		0.36
Size decile FE?	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y

Table 9: 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 $(t-1)$	-0.042***	-0.042***	-0.051***	-0.0038	0.015*	-0.0060
	(0.010)	(0.0089)	(0.015)	(0.0041)	(0.0081)	(0.0068)
Observations	16,095	16,095	14,318	12,357	14,780	16,092
# firms	2,766	2,766	2,750	2,269	2,638	2,991
Mean dep. var.	0.55	0.55	0.47	0.05	0.27	0.26
R^2	0.20	0.24	0.12	0.15	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 10: Stock market performance of hierarchy-sorted portfolios

This table presents multi-factor regressions of the monthly returns of portfolios sorted by the lagged number of hierarchical layers from 2016 to 2023 (96 months). In each year, we sort firms by the number of layers and construct a long-short portfolio that goes long in the top quartile and short in the bottom quartile of firms. The right-hand side variables include a constant (α) and the six factors of Fama and French (2018), which are obtained from Kenneth French's website. $R^m - r^f$ is the market risk factor, "SMB" is the size factor, "HML" is the growth factor, "RMW" is the profitability factor, "CMA" is the investment factor, and "Momentum" is the momentum factor. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
α	-0.00020	-0.0026	-0.0033	-0.0024	-0.0026	-0.0031
	(0.0039)	(0.0027)	(0.0027)	(0.0025)	(0.0027)	(0.0026)
$R^m - r^f$	-0.084	0.12	0.090	0.10	0.12	0.085
	(0.11)	(0.069)	(0.068)	(0.065)	(0.072)	(0.067)
SMB		-0.89***	-0.79***	-0.90***	-0.89***	-0.78***
		(0.090)	(0.11)	(0.090)	(0.093)	(0.12)
HML		-0.090	-0.13*	-0.035	-0.088	-0.057
		(0.076)	(0.075)	(0.14)	(0.078)	(0.15)
RMW			0.22**			0.24**
			(0.11)			(0.12)
CMA				-0.12		-0.15
				(0.18)		(0.20)
Momentum					0.011	0.047
					(0.067)	(0.075)
Observations	96	96	96	96	96	96
R^2	-0.01	0.53	0.54	0.52	0.52	0.53

Table 11: Hierarchical layers in the pharmaceutical industry following Covid-19

This table reports regression estimates using observations at the firm-year level. The sample period spans the years 2017-2022 (three years before and three years after the Covid-19 pandemic). The dependent variable is the logarithm of the number of hierarchical layers. On the right-hand side, "Covid" is an indicator variable equal to one for firms in SIC code 2834 (pharmaceutical companies), and "Post" is an indicator variable equal to one for the post-Covid years (2020-2022). In columns 1-4, the control group includes all firms outside SIC code 2834. In columns 5 and 6, it includes all firms outside SIC code 2834 and outside the "Covid industries" identified by Fahlenbrach, Rageth, and Stulz (2021) that were directly impacted by Covid. In columns 2, 4, and 6, "Post" is replaced with a set of year dummies (using 2019 as base year) that characterize the dynamics of the treatment effect. Size decile fixed effects are based on the firm's total assets in the previous year. Standard errors, clustered at the industry level (4-digit SIC), are reported in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Log # layers					
	(1)	(2)	(3)	(4)	(5)	(6)
Covid × Post	0.040***		0.028***		0.025**	
	(0.0090)		(0.0100)		(0.011)	
Covid × 2017		-0.019**		-0.0080		-0.0068
		(0.0090)		(0.0085)		(0.0097)
Covid × 2018		0.00069		-0.00066		-0.00044
		(0.0058)		(0.0048)		(0.0056)
Covid × 2020		0.020***		0.013***		0.011**
		(0.0030)		(0.0042)		(0.0043)
Covid × 2022		0.034***		0.022***		0.020**
		(0.0055)		(0.0077)		(0.0079)
Covid × 2023		0.051***		0.043***		0.041***
		(0.0065)		(0.0076)		(0.0083)
Observations	14,931	14,931	14,681	14,681	12,375	12,375
# firms	2,827	2,827	2,809	2,809	2,364	2,364
Mean dep. var.	2.03	2.03	2.03	2.03	1.96	1.96
Within-R^2	0.00	0.00	0.03	0.03	0.04	0.04
Size decile FE?	N	N	Y	Y	Y	Y
Firm FE?	Y	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y	Y

Table 12: 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. The right-hand side variables are various measures of AI adoption based on job posting data, obtained from Babina, Fedyk, He, and Hodson (2024). 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 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.

	Log # layers			
	(1)	(2)	(3)	(4)
Share job postings AI	-0.61*			
	(0.34)			
Share job postings AI (narrow)		-0.75*		
		(0.40)		
Average share (skills)			-0.68*	
			(0.36)	
Avg max share				-0.14*
				(0.074)
Observations	5,505	5,505	5,505	5,505
# firms	1,491	1,491	1,491	1,491
Mean dep. var.	2.41	2.41	2.41	2.41
Within R^2	0.01	0.01	0.01	0.01
Size decile FE?	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y
Firm FE?	Y	Y	Y	Y

Appendix

A1 Known job ladders: investment banks

We rely on LinkedIn's representativeness regarding title distribution and promotion reporting. Ideally, we could 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 Jeffries and Lazard.[28]

Jeffries' investment bank titles are analyst, associate, director, and president. The data correctly ranks these roles with analyst at the bottom and president at the top. In the final data, 74% of titles with "analyst" are in the first layer (as expected). The remaining 26% include operations, business, and other non-banking analysts in the second layer. Most individuals (87%) with "associate" in their title are in the eighth layer. The large gap between analysts and associates is because of the rich set of roles at the lower layers in the firm that are not directly working in banking. Indeed, when we narrow the set of analyst jobs to those with clear investment banking roles (step 3), the associates are in layer three. We find directors are in layer 9 (one above associates) and presidents in layer 10. After conditioning on investment banking roles, these layers are five and six. Thus, the data and clustering algorithm correctly ranks and separates the investment bankers. We repeated the same exercise for Lazard and found similar results. Most importantly, when we condition on investment bankers and rerun the algorithm, Lazard has five layers with "analyst" at the lowest level and "president" at the top.

[28]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.

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

A2 Job transitions

An assumption of our algorithm is that observed role changes in a worker's job history are a promotion or a lateral change in their position at the firm. Confirming this is challenging and thus rests on the assumption that individuals rarely stay at firms where they are demoted. If this assumption is valid, individuals who move between firms should be more likely to move to a lower position on the hierarchy than a higher one. To test this and allow for comparisons of ranks across firms, we split firm hierarchies into quintiles. For all workers who move between firms (between the largest internal labor market of each), we identify the quintile they move to relative to the quintile of their previous position. Appendix Table A4 presents the transition matrix. As can be seen, for those outside of the first quintile, more than half of the moves are to the same quintile and below in each transition. In unreported results, we also find that movers who experience promotions across firms move to larger firms, while those who have no change or are demoted move to smaller firms. We interpret these patterns as support for our central assumption that within-firm role changes are promotions.

A3 Figures and Tables

Figure A1: ONet "ranking" and hierarchy layers

This figure reports the average ONet job zone (see `https://www.onetonline.org/find/all`) by hierarchy layer. A job zone "group[s] occupations into one of five categories based on levels of education, experience, and training necessary to perform the occupation." 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 corresponds to the interquartile range across all firms.

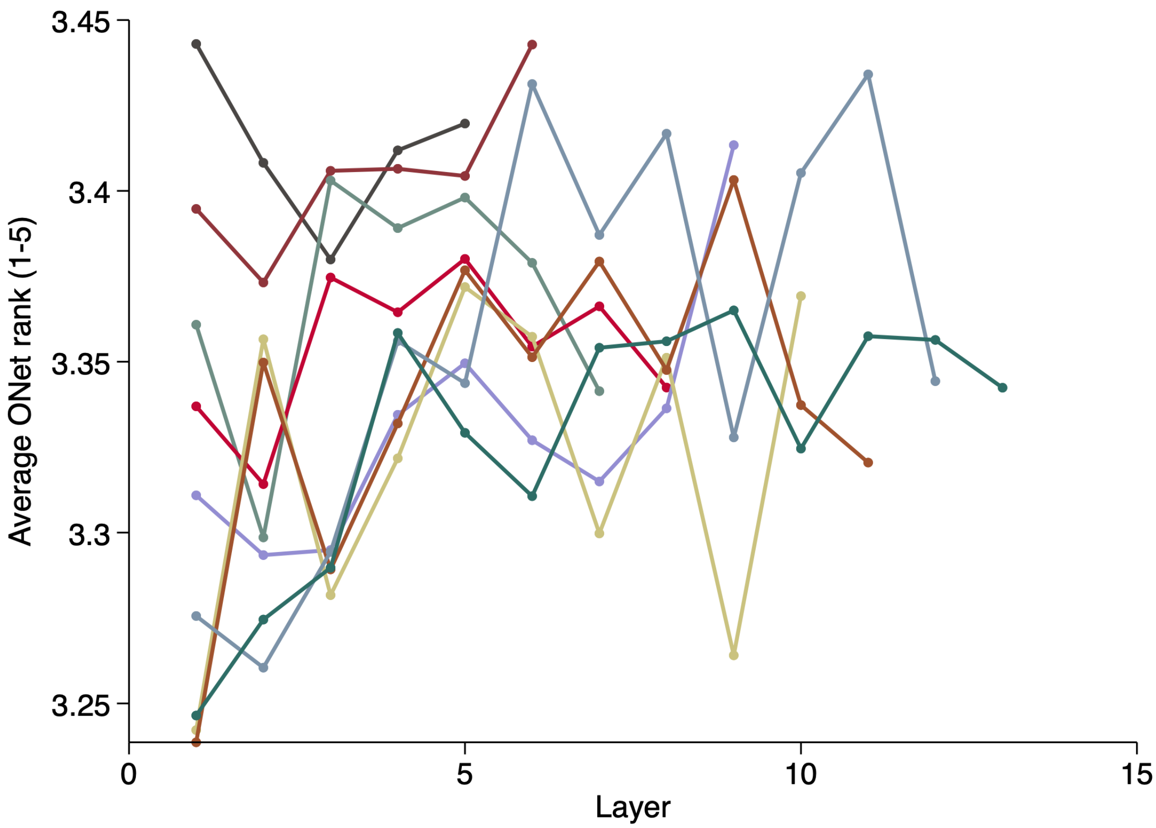

Figure A2: Years to promotion by hierarchy layer

This figure reports the average number of years that a promoted worker spends in a given layer before the promotion. 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 corresponds to the interquartile range across all firms.

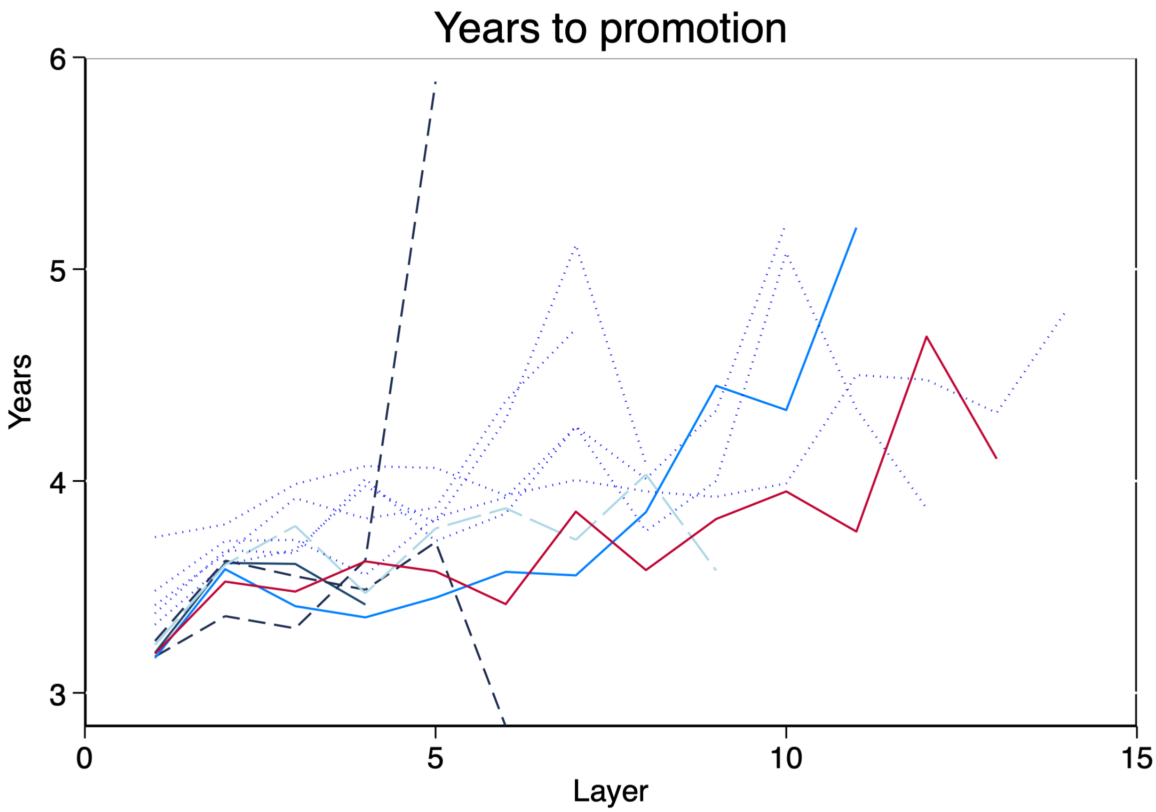

Figure A3: Tenure by layer quintile

This figure reports the average number of years in a role for all workers who switch roles within the firm over the sample period. The averages are reported within a layer quintile. Each line represents a different group of firms that has the same total number of layers.

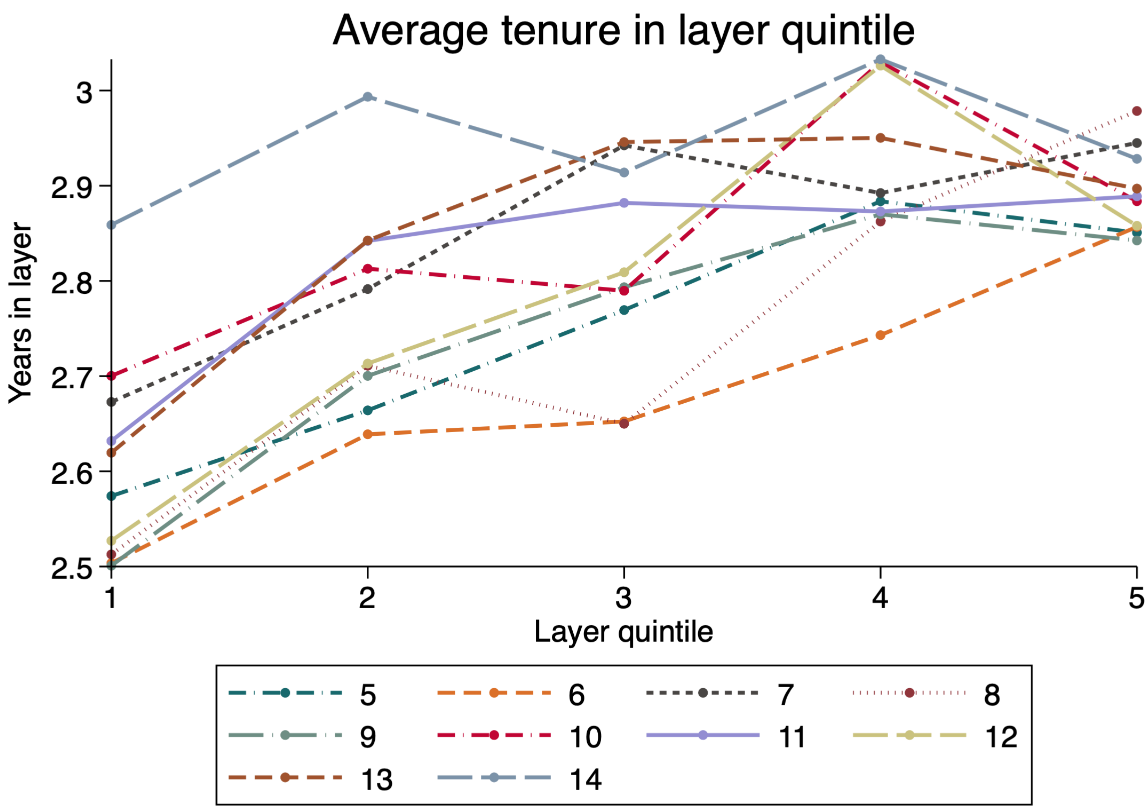

Table A1: Measurement of hierarchy in organizational studies

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 (2024)	# 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 (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 (2018)	# Managerial layers	8	India (one bank)	Bank reports
Tag, Astebro, and Thompson (2016)	# Hierarchical layers / span of control	4	Sweden	Occupational codes

Table A2: Top 10 most common titles in each quartile

This table reports the 10 most common strings in job titles by hierarchical layer quartile (for firms with at least four layers), along with their frequency percentages within the top 10. The last row reports the percentage of all job titles in the respective quartile that have strings in the top 10.

Quartile 1		**Quartile 2**		**Quartile 3**		**Quartile 4**	
manager	18.57%	manager	31.05%	manager	34.09%	manager	20.67%
engineer	13.63%	senior	13.50%	senior	13.13%	vice	15.73%
senior	12.43%	engineer	12.47%	engineer	11.35%	president	15.69%
specialist	10.39%	analyst	7.47%	director	11.09%	director	10.73%
sales	9.71%	sales	6.94%	software	6.95%	senior	10.65%
analyst	9.58%	specialist	6.55%	sales	5.17%	engineer	8.12%
associate	9.55%	director	6.16%	sr	4.81%	software	5.86%
representative	5.56%	business	5.40%	product	4.79%	sales	4.42%
service	5.53%	operations	5.34%	development	4.40%	vp	4.15%
consultant	5.05%	lead	5.12%	operations	4.23%	program	3.98%
Top 10 cover: 27.37%		Top 10 cover: 27.86%		Top 10 cover: 29.42%		Top 10 cover: 32.66%	

Table A3: Where external employees join in the hierarchy

This table reports regression estimates where the dependent variable is an indicator for whether a worker joined the firm from another firm in the last year (columns 1-2) or the last two years (columns 3-4). The sample includes all such workers, as well as those workers who stayed at the firm but switched layers. The unit of observation is a worker-year in which any of these moves occurred. The right-hand side variables are indicator variables for the quintile of the hierarchy where the individual moved to or was promoted to. The excluded quintile (base group) is the bottommost quintile. 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.

	Joined $t-1$ to t		Joined $t-2$ to t	
	(1)	(2)	(3)	(4)
Quintile 2	-0.029***	-0.029***	-0.081***	-0.081***
	(0.0028)	(0.0028)	(0.0084)	(0.0084)
Quintile 3	-0.038***	-0.038***	-0.11***	-0.11***
	(0.0022)	(0.0022)	(0.0073)	(0.0074)
Quintile 4	-0.041***	-0.041***	-0.11***	-0.11***
	(0.0034)	(0.0034)	(0.015)	(0.016)
Quintile 5	-0.050***	-0.050***	-0.15***	-0.15***
	(0.0034)	(0.0034)	(0.014)	(0.014)
Observations	1,502,037	1,493,506	1,657,456	1,648,142
# workers	1,258,578	1,252,008	1,303,957	1,298,748
# firms	1,692	1,692	1,693	1,693
Mean dep. var.	0.09	0.09	0.33	0.32
R^2	0.00	0.00	0.01	0.01
Firm FE?	Y	Y	Y	Y
Size decile FE?	N	Y	N	Y
Year FE?	Y	Y	Y	Y

Table A4: Transition matrix between firms

This table reports the transition probabilities across layer quintiles between firms. A cell reports the fraction of jobs transitions between firms from a given layer quintile of the previous firm (marked from 1 to 5 in the first column) to a given layer quintile of the new firm (marked from 1 to 5 in the first row).

	Layer quintile of new firm				
	1	2	3	4	5
1	**50.8**%	21.5%	13.5%	9.7%	4.4%
2	39.3%	**24**%	17.2%	13.1%	6.3%
3	33.8%	23.8%	**18.3**%	16.3%	7.8%
4	30%	22.5%	18.5%	**19.5**%	9.5%
5	28.1%	22.1%	17.6%	19%	**13.2**%