Corporate Hierarchy

Michael Ewens and Xavier Giroud*

March 2025

Abstract

We introduce a novel measure of corporate hierarchies for over 2,600 U.S. public firms. This measure is obtained from online resumes of 16 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 more stable cash flows. We also examine how companies adjust their hierarchies in response to demand and knowledge shocks. We find that biotech companies increased their number of layers following the Covid-19 pandemic, while companies flatten 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 seminar participants at the University of Houston, Rice University, SKEMA, INSEAD, HEC, University of Oxford, and the University of Warwick, as well as Daniel Bias and Lukas Freund for helpful comments and suggestions. We thank Shiyang Chu 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 allow them to track 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 three 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. 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 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), 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 and cash flow volatility rather than increased 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.

Third, we examine the implications for equity prices. Interestingly, long-short portfolios sorted based on our hierarchy measure (long in high-layer firms, short in low-layer firms) generate a statistically significant, negative alpha. This finding is robust across various asset pricing models, including the 6-factor model of Fama and French (2018). This negative alpha

could reflect the lower risk of high-layer firms (and, conversely, the higher risk of flatter firms), as a source of risk that is priced by the market but not captured by other factors.[2]

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. Consistent with this prediction, 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 theory, we find that companies decrease their layers after adopting artificial intelligence (AI) technologies, which we capture based on the number of AI job postings. 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 described 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]The negative alpha could also indicate that investors overestimate the benefits of corporate hierarchies. However, when looking at analysts' EPS (earnings per share) forecasts, we find no evidence that analysts overestimate the performance of high-layer firms.

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 most firms' production depends on the skills of its 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) summarizes, when we observe an organizational structure, we can infer the agents' knowledge, the order of information and problems flow 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 trades off specialization for production and additional worker coordination. 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 model 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 set of 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 predictions for firm decisions and outcomes.

2.2 What is a layer?

Before introducing our hierarchy measure, we define a hierarchical layer used throughout the paper. The empirical methods described below (Section 3) consider a layer as a set of jobs with similar complexity, tasks, and responsibilities. In turn, jobs or roles in a layer should have a similar wage profile. Caliendo and Rossi-Hansberg (2012) and Caliendo, Monte, and Rossi-Hansberg (2015) define a layer as a group of employees with similar knowledge levels and wages performing tasks at a comparable level of authority. How information flows between roles can also reveal levels. For example, Bolton and Dewatripont (1994) workers in layers send messages to those above and never send them down a layer. In the Garicano (2000) knowledge hierarchy, a layer is a set of workers with the same scope of knowledge and who communicate problems to the layer above them. We do not observe communication nor worker responsibilities, but believe the levels we measure capture similar duties in production and close ranking within the firm.

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 with 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; Bias, Lochner, Obernberger, and Sevilir, 2024; Gumpert, Steimer, and Antoni, 2022). 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 range of possible levels using occupation codes ranges from 4 to 8. 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).[3]. There are two critical pieces of intuition for this method. First, job transitions ("worker flows") reveal the ranking of jobs (Clauset, Arbesman, and Larremore, 2015). Second, a layer in a firm is a collection of roles that are more likely to have lateral moves between them rather than promotion or demotion (Bonhomme, Lamadon, and Manresa, 2022).[4]

The first step is to identify the connected jobs in the firm, or internal labor markets (ILMs). The method requires a ranking of roles, but that can only be achieved when one can observe flows between roles. 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 an average of over 70% of the worker-years in our sample (see Section 6.1). One cost of this focus is in terms of multi-divisional firms, since 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 roles B (50%) and role C (50%), while row C only transitions to A 5% of the time, the ranking would put B and C above A. Huitfeldt, Kostøl, Nimczik, and Weber (2023) introduce a minimum violation ranking (MVR) to account for the noise in resume data and job transitions. The algorithm starts by ranking roles by promotions to it. 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

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

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

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 rank standard deviation from the second step is a signal of a layer. 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, 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 after discussing the data inputs in the next section.

5 Data

5.1 LinkedIn data

We leverage data on worker resumes provided by CoreSignal.[5] The data provider extracts resume data—work histories, titles, employers, and education—for publicly available profiles of U.S.-based workers from 2016 to 2023. While we only focus on workers active at the firm in this sample period, we use their reported work history, which is available on LinkedIn. Thus, we have a much longer view of the firm's job ladders than the final sample period. We exclude temporary or part-time workers using titles like "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. 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 headcount between our data and Compustat (see Section 6.1) reveals that coverage is below 50% for most firms. However, Compustat covers temporary workers

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

and non-U.S. employees. The algorithm to measure hierarchical layers described below is not impacted by over- or under-representation of a given title. Rather, we require that all titles occur at least once in the data and that a promotion or demotion exists between titles.[6] Again, we do not have a means of directly validating whether we have complete coverage of titles or transitions. In the final sample described below, we have 5,551,952 job transitions from 2016 to 2023. To nevertheless validate our data, we rely on an extensive set of validation tests (see Section 6.2) that show that the data capture features of hierarchies predicted by theory and practice.

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. While our methodology does not map jobs to occupational codes, this large number of titles makes it challenging to classify transitions. 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., reduce 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,

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

the algorithm maps it to a common string. 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 that of public firms in Compustat. This merge is facilitated by the availability of the company's website (URL), which is reported in both datasets. There are still many false positives and negatives, so we supplement the merge with a fuzzy string comparison of firm names and an extensive manual check of hundreds of potential matches. We match 7,277 total firms to the resume site data 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 7 million unique workers.

The final sample of firms only includes those in Compustat on or after 2015 with at least one year of measured hierarchical layers. We exclude firms in the following SIC codes: 6000-6999 (financials), 4900-4999 (utilities), 5200 and 5300-5399 (retail), 7100-7299 (service firms with poor LinkedIn coverage), and 4400-4599 (airlines and shipping). Any remaining firms after these filters whose LinkedIn employee count is less than 15% of Compustat employees (Compustat item "emp") or with less than 25% of worker-years in the largest internal labor market are dropped. We are left with 2,611 firms (though the sample size changes depending on the regression). Over 6,000 firms outside the filtered industry codes either lack a hierarchical layer measure, have no employees on the resume website, or have too low coverage. Table 2 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 larger 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 3 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. Table 3 shows that the average firm has 253 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 192 unique roles on average (recall that titles are mapped to roles). On average, there are 321 roles in a firm.

Two statistics provide comfort that the largest ILM captures most roles and workers. Specifically, 54% 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 44% 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 10.9 (median 8). This exceeds the number of layers—typically 4—obtained in the studies reviewed in Appendix Table A1.[7] The higher number of layers at U.S. public firms is not surprising, since they are larger than most firms considered in the prior literature. Importantly, this higher number of hierarchical layers suggests that the existing methods (e.g., those based on occupational codes that have a built-in cap on the number of layers) may be too coarse to capture the rich organizational structure of publicly traded firms.

[7]The exception is Adhvaryu, Bassi, Nyshadham, Tamayo, and Torres (2023), who identify 14 layers in their study of an Argentinean automotive working group.

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. (See the definition in Table 1.) The first two measures use their 5-layer 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 0.76 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.

Finally, Panel B reports the statistics from Panel A for the complete set of firm-year observations. The evidence mirrors the one 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. Figure 2 presents the kernel estimations of layers over time, and confirms this pattern. 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, we ask whether hierarchical layers vary across industries and within industries over time. Panel A of Figure 3 shows meaningful variation across industries: manufacturing has many layers and a substantial variation (the interquartile range is 9 to 25 layers). Panel B of Figure 3 shows that the information technology (IT) sector has relatively high layer counts, which have declined over the sample period. The chemicals and drugs industry experienced a significant increase in layers in 2020, likely driven by the pandemic. We explore this specific change in Section 8.

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, we find that only 3.4% of firms have change in

their layers; 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.[8]

6.2 Validation

In this section, we present a series of validation tests to assess the strengths and weaknesses of 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 validations checks whether layers have the expected titles, work experience, and connections via promotion. Consider the titles in Table 4, which shows the most commonly occurring titles by layer. While firms with different total layers are difficult to compare, these titles suggest that the algorithm and ranking capture hierarchy as expected (i.e., similar job functions and skills). Appendix Table A2 shows that titles with "director" and "senior" are increasingly more prevalent higher up in the hierarchy, while "sales" and "engineer" are increasingly more prevalent in lower layers. The hierarchical layer's rank also positively correlates with other measures of levels. Figure A1 reports the average ranking for titles by layer after mapping the titles to the ONet classification system. We see 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 4 shows that work experience—measured since the worker graduates or 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 demonstrate 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.

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

6.2.2 Firm size

First, models of 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 5 shows how the distribution of hierarchical layers varies across size quintiles. We focus on the manufacturing sector, but the patterns are similar across all other sectors. As is 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 5 investigates the connection between firm size—measured in seven different ways (using total assets, employment, revenues, the number of business segments, the number of U.S. subsidiaries, the number of executives, and board size)—and the number of layers. The dependent variable is the log of the number of layers in $t + 1$, while all right-hand variables are measured at time t. 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), the more likely they are to alter their organizational structure via 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 theories of corporate hierarchy, as well as the results from Table 5, motivate the inclusion of size controls in all our regressions. In our preferred specification, we will include asset decile fixed effects.

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, each higher layer will have a smaller share of employees. For example, this follows from the models of information hierarchy where hierarchies divide workers into specialization vs. coordination (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).[9]

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

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 with the number of employees in our LinkedIn data. We know from Table 3 that LinkedIn source 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 but does impact employee shares. Thus, estimating firm shapes entails some measurement error.

There are two ways how 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 6 follows the second approach, showing the share by hierarchy quintile for firms with more than four layers. In each industry sector, we observe a monotonic decrease in the share of employees as layers increase. Figure 7 follows the first approach, reporting the average shape of firms with four and eight layers where the shares are averaged across all firm-years. As can be seen, the evidence is again generally consistent with a pyramidal shape.[10]

An indirect view of a firm's hierarchy 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 jobs at the next layer than at the current layer. This restriction implies that we should observe a decreasing promotion rate from layer i to $i+1$. Figure A4 confirms this prediction: as we move to higher layers, we observe a monotonic decrease in next-layer promotion rates. A look at transitions between layer quintiles in Appendix Table A3 confirms that transitions between layers decrease with layer and that most workers remain in their layer quintile in each year (the diagonal).An alternative view of these patterns can be found in analyzing the time needed for any promotion. The lower number of spaces above a given layer should result in a longer

[10]Figure 7 also plots the "span of control" across layers. The span of control of layer i is the share of employees in that layer divided by the share in layer $i+1$. A lower span of control shifts the firm's shape from pyramidal to a rectangular. In Figure 1, we observed that the number of layers had weakly declined over the sample period. Appendix Figure A3 repeats this analysis for the span of control. Specifically, we split the hierarchy into quintiles for all firms with at least five layers. The figure then reports the average span of control for the bottom four quintiles by year. The dynamics in "4th quintile" are consistent with the flattening of the top of the firm—or the increase in the number of executives who report to the CEO found in Rajan and Wulf (2006) and Guadalupe and Wulf (2010). Combined with Figure 1, this evidence indicates that the average public firm has fewer layers over time and a less pyramidal shape.

time to promotion as layers increase. Indeed, we find in Figure A5 a 14% increase in the years to promotion between layer one and layer 10.

6.2.4 Worker characteristics

The last set of validation tests attempts to confirm that worker characteristics across layers vary as we would expect from theory. For example, 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 (i) a prevalence of managers in higher layers and (ii) an increase in the share of managers as we move from the bottom layers. Figure 8 reports the share of workers in each layer with "manager" or "supervisor" in their title. We see a significant increase—16% to 25%—as we move from layer one to layer five. The share of managers then stabilizes at layer six. Table 6 explores the differences across firms by regressing the share of workers who are managers on the number of layers. Firms with more layers have a larger share of managers. (They have a larger share of supervisors as well, although the point estimates are not significant.) 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 the remainder of this paper, we use our measure to explore the role of hierarchies for corporate decision-making and a wide range of outcomes.

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, and asset decile fixed effects. That is, we compare firms in the same year, same industry (based on 4-digit SIC codes), and of similar size, that differ based on their number of layers. 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.[11]

Table 7 reports regression results for the fraction of employees that have at least a bachelor's (column 1) and a master's degree (column 2), respectively. The results confirm that firms with more hierarchical layers tend to have a more educated workforce. While the point estimate is positive in both columns, it is only significant for bachelor's degrees. This suggests that undergraduate degrees are especially sought after among multilayered firms.

Next, we examine the relationship between the number of layers and the extent to which

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

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 8, we examine several dimensions related to the firm's human capital. The first two columns show that high-layer firms have higher internal promotion rates (as a percent of active workers) and are less likely to bring in external hires. 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 and 4, 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.

An important part of the firm's internal labor market is the top executive team. In the last three columns, 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 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 in 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).

Panel A of Table 9 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 that are not directly involved in the production process.

Panel B of Table 9 explores cost measures (scaled by total assets). In column 1, 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 2, we consider wages. The point estimate is positive, suggesting that multi-layer hierarchies have more expensive human capital. However, we caution that the estimate is not significant at conventional levels.[12]

We next study 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

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

targets and potential synergies. On the other hand, the multi-layer structure may complicate the integration of other firms and their workforce.

Table 10 studies whether our measure of hierarchical layers predicts acquisition activity and the type of acquisitions. The unit of observation is the firm-year. We measure acquisition activity using SDC 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 more likely to engage in acquisitions. A one standard deviation increase in layers corresponds to a 15% higher probability of an acquisition (relative to the mean). 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 log of the total (reported) valuations paid (column 4). Overall, these results indicate that high-layer firms are significantly more active in the acquisition market.

The final column in Table 10 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. Column 5 shows no evidence that high-layer firms differ on 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 11 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). This higher R&D investment translates into a larger number of patents. Specifically, the estimates in columns 2 to 4 show that a one standard deviation increase in layers corresponds to a 22% increase in patenting (all regressions are Poisson). Do these patents differ in quality? The last three columns suggest they do not. In each regression, the dependent variable is the average number of citations per patent. As can be seen, we find no significant difference in this proxy of patent quality.[13]

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

In Panel A of Table 12, 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 (form CRSP) over the year. Columns (1) and (2) refer to the raw measure, while it is winsorized at the 5% level in columns (3) and (4). The odd columns refer to our baseline specification, while the even columns also include firm fixed effects. As can be seen, all estimates indicate that more layers are associated with lower stock volatility.

In Panel B, we refine this analysis by asking whether the lower stock volatility reflects lower business and/or financial risk. In columns 1 and 2, we consider measures of operating asset volatility and cash flow volatility, respectively.[14] As is shown, firms with more layers have significantly lower business risk based on these two metrics. In the last two columns, we examine leverage (the debt-to-asset ratio) and cash holdings (the cash-to-asset ratio) and find that the number of layers is not a significant predictor of these variables. This indicates that the lower stock volatility of high-layer firms is due to business risk rather than financial risk.

[13] In additional analyses, we also considered novelty and the number of right-tail patents. Again, we found no significant difference between firms based on their hierarchical layers.

[14] Both measures are computed as in (Gormley and Matsa, 2016).

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 its stock return. 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. The excess return of this portfolio averages -1.5% per annum. Table 13 presents the multi-factor regressions of the monthly returns. The estimates from the single-factor model in column 1 show evidence of a weak negative market loading. The results above related to firm size, profitability, and volatility differences by hierarchical layers suggest that additional factors are required. Column 2 adds the remaining two Fama-French factors. Here, we see that the portfolio behaves like a large-cap growth portfolio.

The statistically significant alpha (1% annual return) shows there is still a premium to the low-layer stocks. Column 3 adds the profitability factor, and as expected from the productivity results, it loads positively. The following two columns add the investment (column 4) and momentum factors (column 5). The negative alpha remains significant across specifications, but the growth factor loses significance. The full 6-factor model (column 6) shows that high-layer stocks behave like large-cap, "quality" stocks. The latter feature signals stability and resiliency, proxies for which we found in the volatility analysis. While we are unaware of organizational theories tied to these results, it is reassuring to see that they match the correlational tests above.

Interestingly, the negative alpha indicates that high-layer firms trade at a discount. This could reflect the lower risk of high-layer firms. Arguably, to the extent that high-layer firms have better risk management abilities (as discussed in Section 7.5), they might face lower risk along dimensions that are priced by the market but not fully captured by the 6-factor model.[15]

[15] An alternative explanation for the negative alpha is that high-layer firms may systematically underperform relative to the market's expectations. However, we find little support for this argument. When looking at analysts' EPS (earnings per share) forecasts, we find no evidence that analysts systematically overestimate the performance of high-layer firms.

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.

Panel A of 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.[16]

In Panel B, we consider the span of control (defined as in Section 6.2.3) in the middle of the hierarchy for firms with at least five layers. As can be seen, we find that the span of control decreases post-IPO. This lower span of control suggests that firms become less pyramidal. Taken together, these results indicate that, after going public, firms increase their number of layers while adopting a more rectangular hierarchical shape.

[16] We caution that these results are correlational due to the endogenous nature of going public decisions.

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 of SIC code 2834 or (ii) all firms outside of 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.

Table 14 presents the results. Column 1 shows that the number of layers weakly increased by 3% among the treated firms in the post-Covid years. Column 2 introduces a dynamic specification. As can be seen, the response starts in 2021, about 9 months after the pandemic's

start and the year in which most vaccines were released. The pre-treatment coefficients are all small and insignificant, which mitigates the possibility of pre-trends that could confound our estimates. In columns 3 and 4, we find that the results are similar if we include asset 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. Overall, the evidence from Table 14 indicates that firms respond to demand and complexity shocks by adding hierarchical layers. 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).[17] The regressions also include firm, year, and asset 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 above, 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 15. We consider four different metrics of AI adoption: the share of AI job postings, broadly (column 1) and narrowly (column 2) defined, the average share of AI skills (column 3), and the average maximum share of AI skills (column 4). As can be seen, we find evidence for a reduction in the number of layers in all four columns, as predicted by theory. That being said, since the tests are under-powered due to the persistent nature of the number of hierarchical layers, we caution that the estimates are statistically in-

[17] Following Babina, Fedyk, He, and Hodson (2024), we exclude firms in NAICS code 51 and 54 ("Information" and " Professional, Scientific, and Technical Services").

significant based on conventional thresholds, but close to the 10% significance level in columns 1, 2, and 3.

9 Conclusion

We propose a measure of corporate hierarchy for U.S. public firms that is obtained from online resumes of 16 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 more stable cash flows. Turning to the determinants of corporate hierarchies, we find that biotech companies increased their number of layers following the Covid-19 pandemic, while companies flatten their hierarchies following the adoption of AI technologies. These findings echo well with models 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 of hierarchies, in keeping with the set of facts documented in this paper. 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.[18]

[18]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.

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.

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.

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.

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.

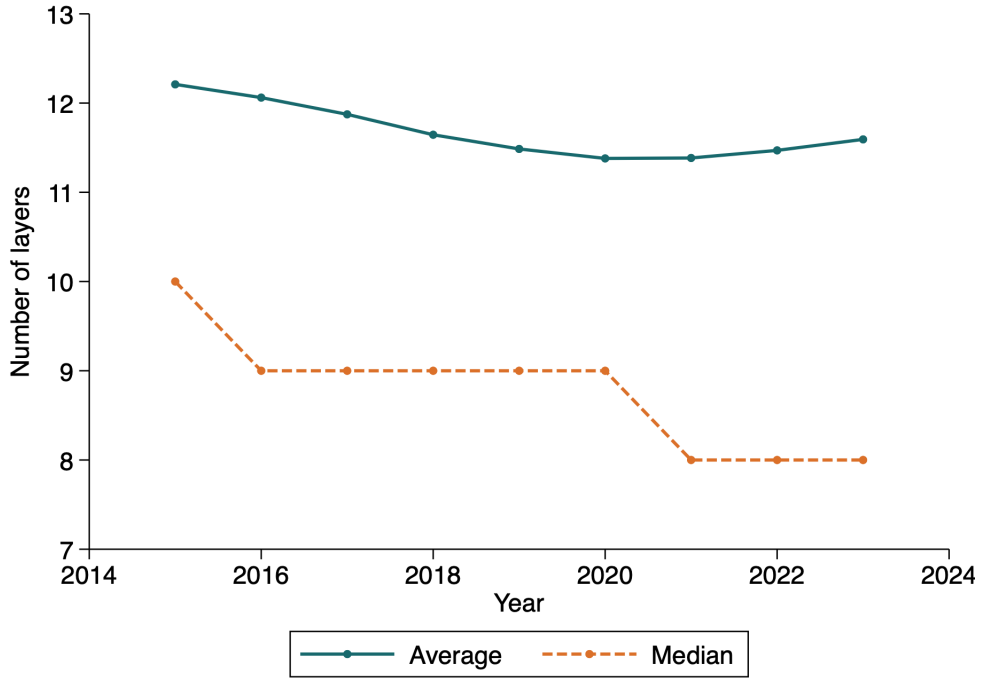


Figure 2: Full distribution of organizational layers by year

This figure reports the kernel densities of the number of hierarchical layers by year for the firms in our sample.

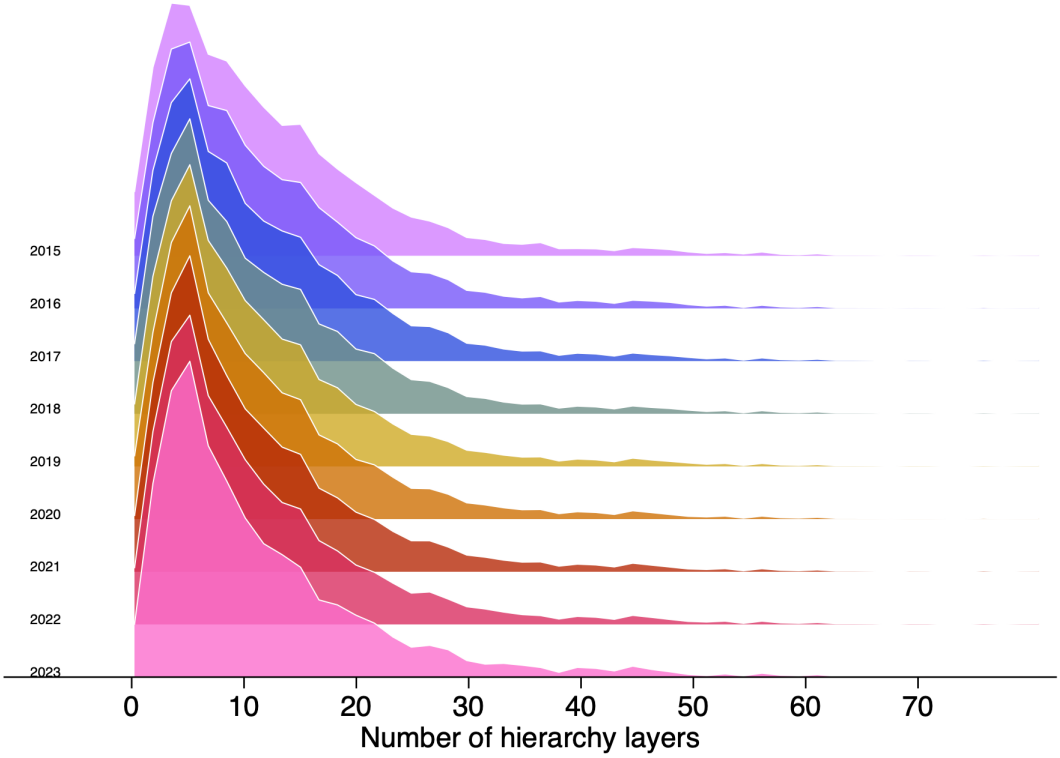
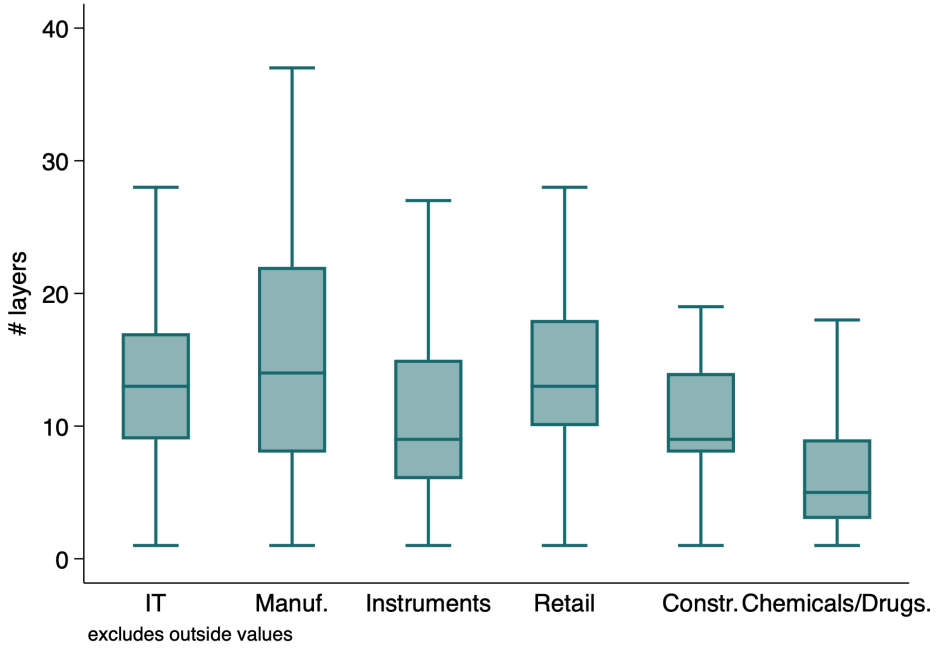


Figure 3: Organizational layers by industry

Panel A figure reports the box plot (see Figure 1 for definition) for six major industries in the sample. All layers are measured in the last year, during which the firm's hierarchy is measured in the data. Panel B reports the average of the number of hierarchical layers by industry-year.

Panel A: Hierarchical layers by major industry



Panel B: Hierarchical layers by industry-time

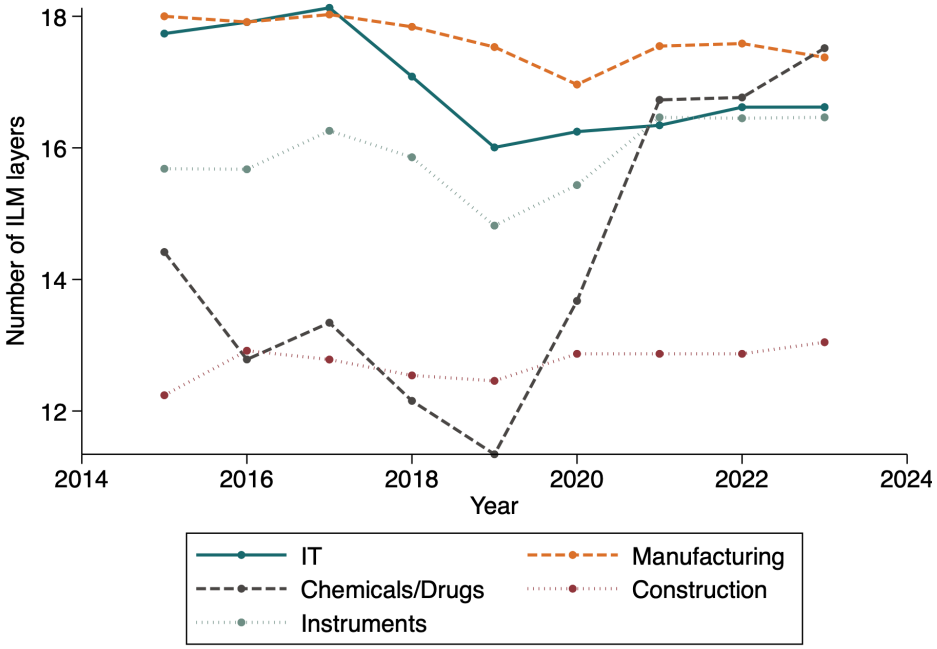


Figure 4: Work experience across the firm hierarchy

The figure reports the median years of work experience across all jobs (i.e., years since college or first job).

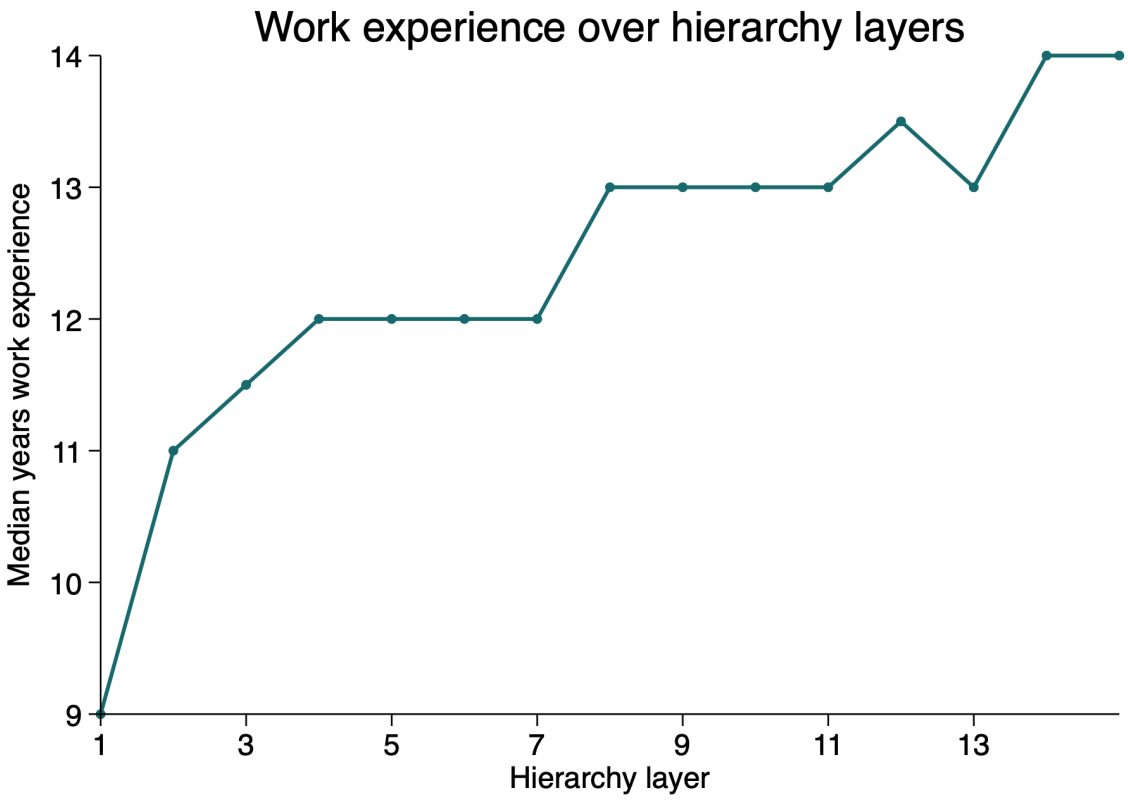
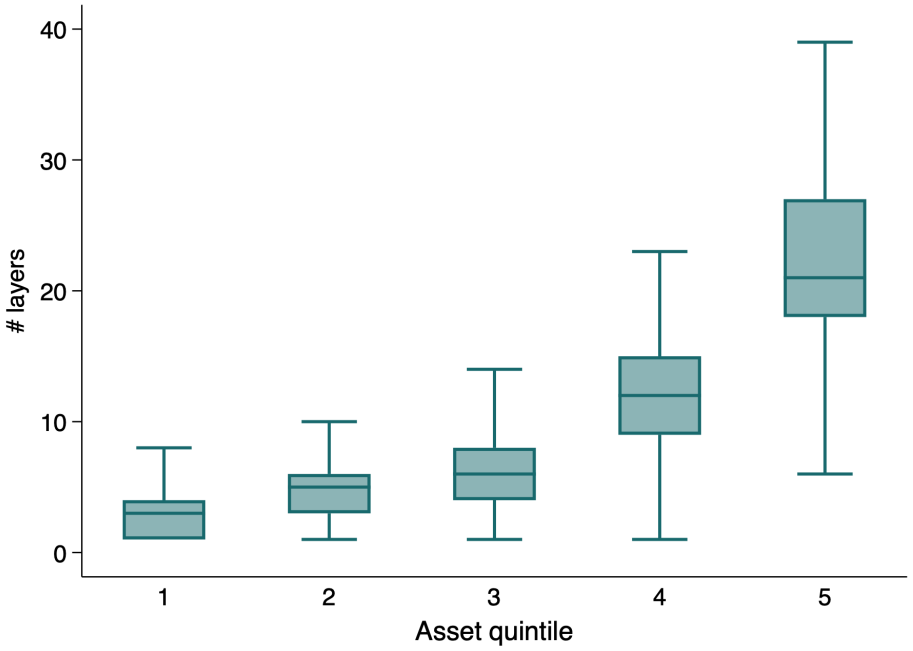


Figure 5: Variation in organizational layers within size buckets: Manufacturing

Panel A splits the sample of manufacturing firms in 2023 by asset quintiles and presents the box plot distributions of the number of hierarchical layers. Panel B splits the sample of manufacturing firms by employee counts as reported by Compustat, plotting the box plot of hierarchy layers.

(a) Panel A: Assets quintiles (Manufacturing)



(b) Panel B: Employee count quintiles (Manufacturing)

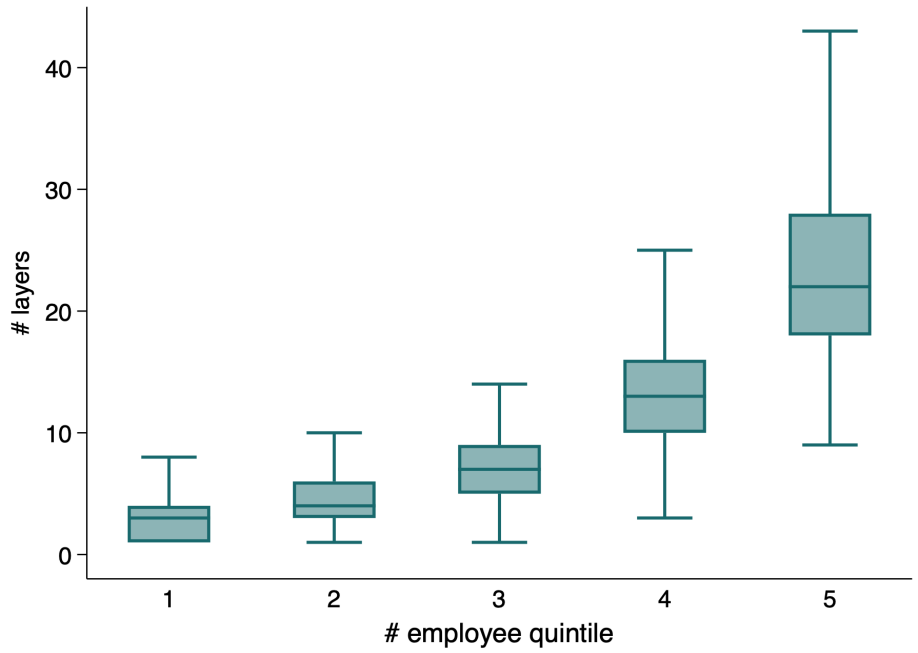


Figure 6: Share of employees by hierarchy layer quintile: industries

The figures reports the share of employees across the firm in each layer quintile. It presents these for the four major single digit SIC codes. 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.

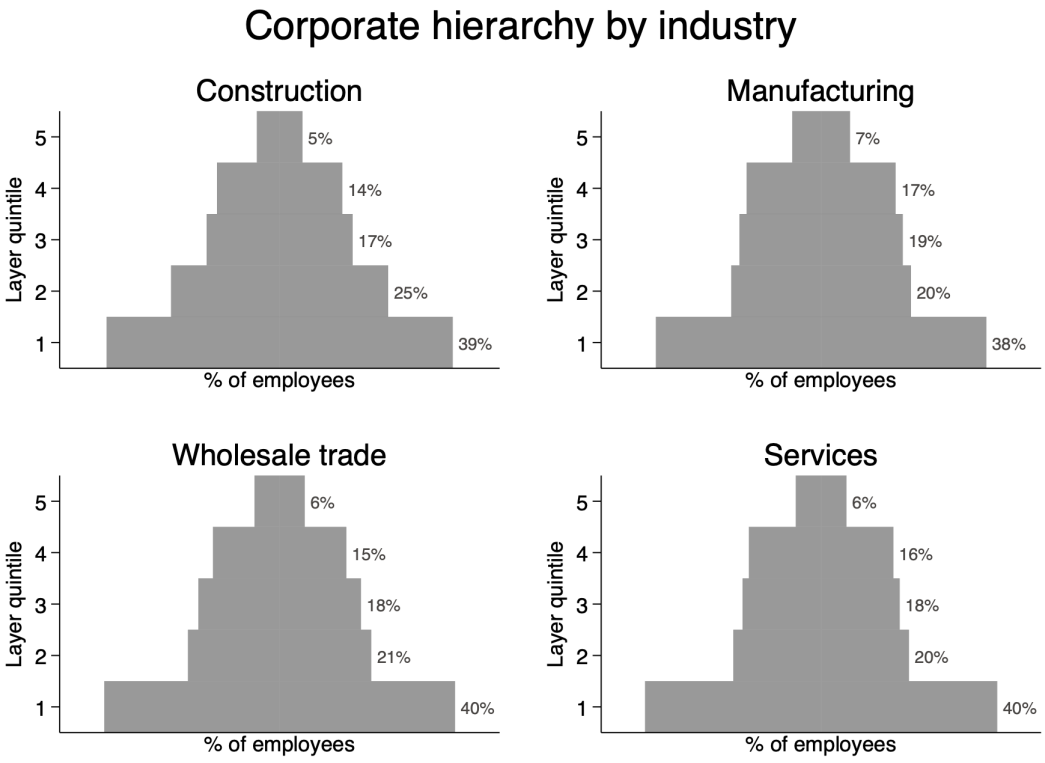
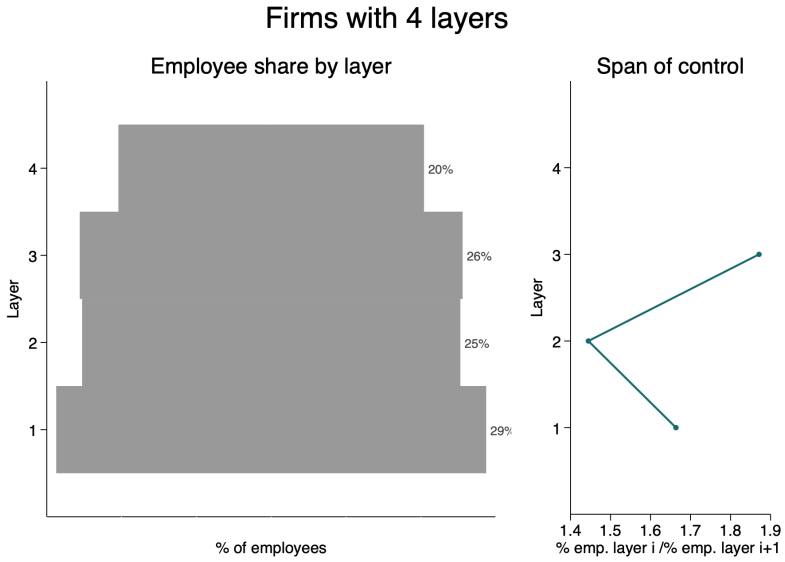


Figure 7: Share of employees by layer and span of control

The figures reports the share of employees across the firm in each layer (left figure) and the span of control by layer. It presents these for firms with four (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 figures report the average of these shares and span of control across all firm-years. Span of control for layer i is defined as the share of employees in layer i divided by the share of employees in the next layer $i + 1$.

(a) Panel A: Firms with 4 layers



(b) Panel B: Firms with 8 layers

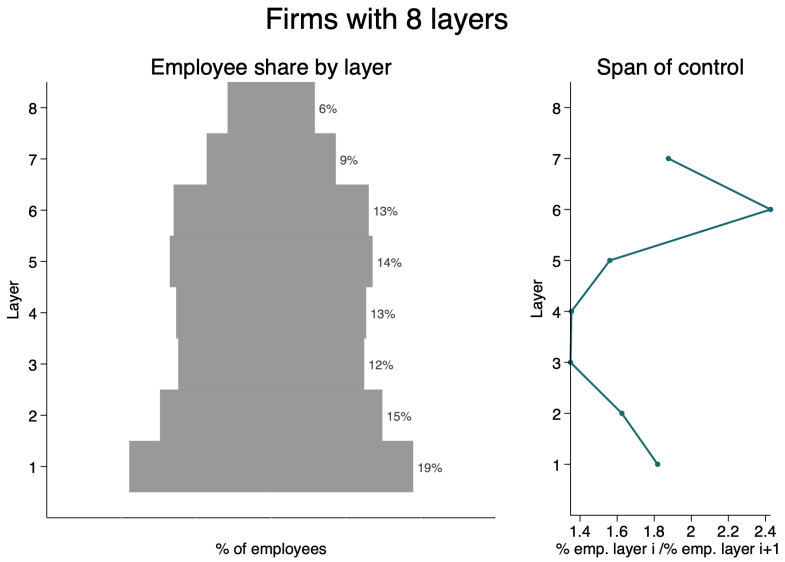


Figure 8: Managers and supervisors across the firm hierarchy

The figure reports the fraction of employees in each hierarchical layer with "manager" or "supervisor" in their title.

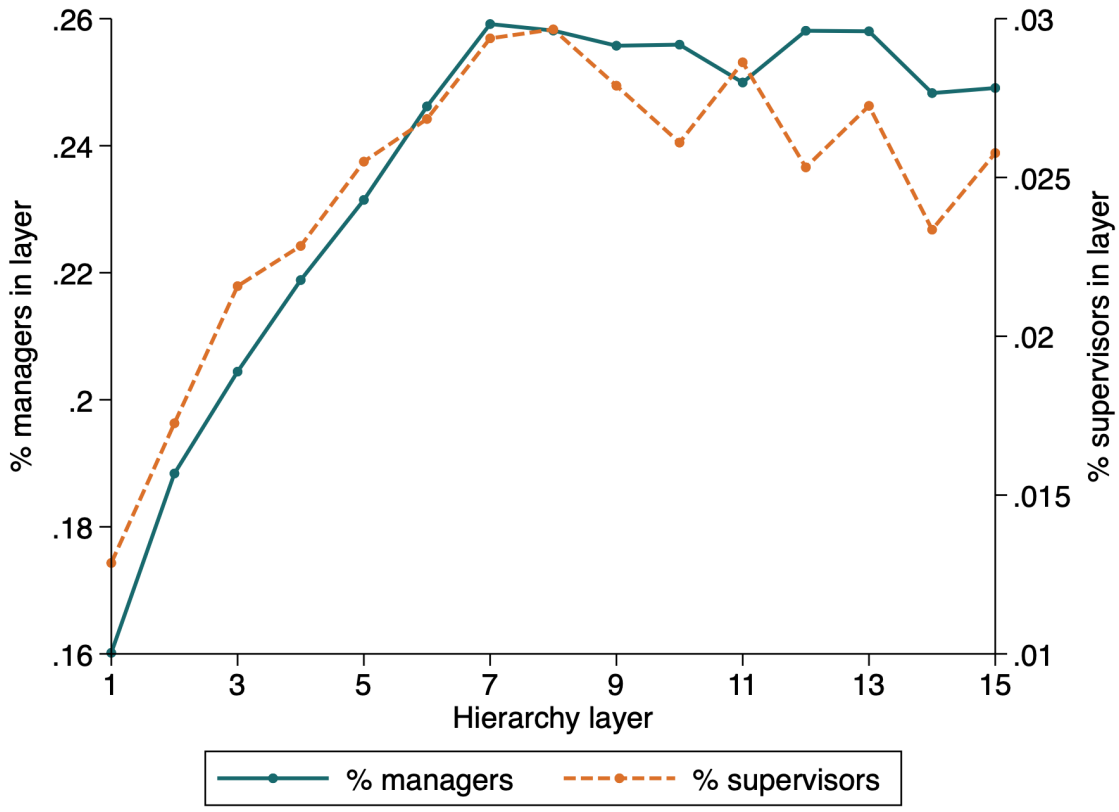


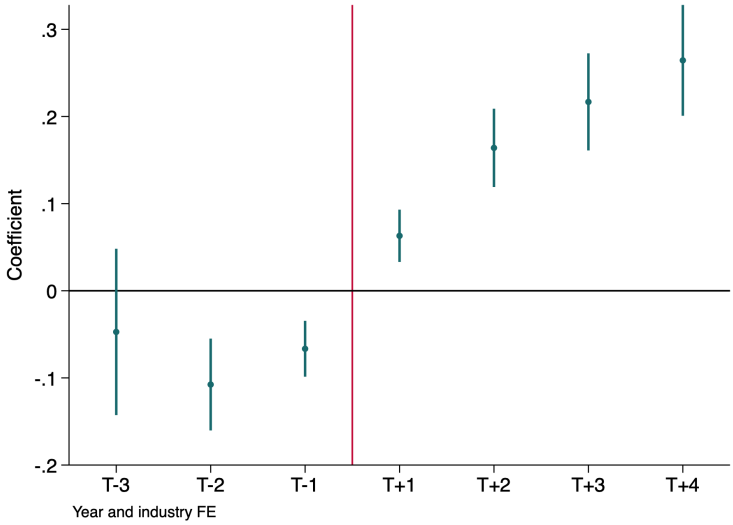
Figure 9: Change in the log of organizational layers around the IPO

The figure reports the coefficient estimates $\hat{\beta}_i$ from the following regression:

$$Y_{it} = a + \sum_{k=-3}^{k=4} \beta_k * \text{k years since IPO} + \gamma_t + \delta_i + \epsilon_{it}$$

Here i is firm and t is year for the set of firms that had an IPO in our sample period. Y_{it} is $\log \#$ layers$_{it}$ in the first panel: the log number of hierarchical layers. The second panel reports the span of control – share of employees in the middle quintile over the share in the 4th quintile. "k years since IPO" is a dummy variable that is one if the year is k years since (to) the firm's IPO. γ_t is a year fixed effect and δ_i is an industry fixed effect. The excluded year is the year of the IPO. Standard errors are clustered at the company-level, and the figure reports the 95% confidence interval.

(a) Panel A: Log number of layers



(b) Panel B: Span of control middle quintile

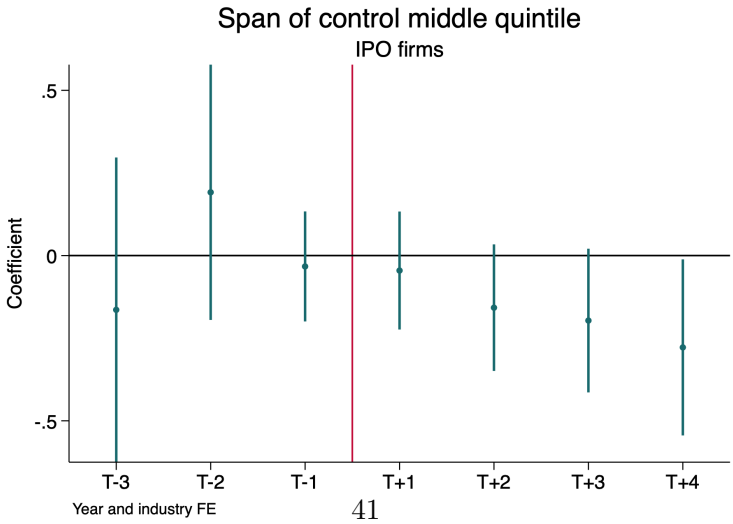


Table 1: Variable definitions

Notes: The table defines the main variables used throughout the paper.

Variable	Definition
Log # of layers	The log of the number of hierarchical layers for a firm created using the career history data in LinkedIn resume data and the algorithm from Huitfeldt, Kostøl, Nimczik, and Weber (2023). We estimate this using the largest internal labor market (see below).
Number of internal labor markets (ILM)	The count of the number of components of the firm's workers connected via promotion. Two jobs are connected if at least one promotion or demotion exists between them. An ILM is a set of jobs that are connected by some such path.
Role	We aggregate the raw job titles found in LinkedIn using a text clustering technique. These roles are used for the promotion analysis and creating hierarchical layers.
Size of largest ILM	The number of unique roles in an internal labor market.
Span of control	For a layer i of a firm's hierarchy, this measures the share of all employees in the firm that are in layer i divided by the same fraction for the next layer $i+1$.
% of role in largest ILM	The fraction of all roles in the firm found in the largest internal labor market.
CEO layer / max layer	If we have the CEO in the job data, then this reports the layer of the CEO divided by the total layers in the firm.
Fraction worker-years in ILM	The fraction of all worker-years whose title (and thus role) is in the largest ILM in each year. If all jobs are in the largest ILM, then this number is 1.
Number hierarchy layers	The count of the number of hierarchical layers in the firm.
"Hierarchization"	The measure of hierarchy from Bias, Lochner, Obernberger, and Sevilir (2024), which we generalize given our layers definitions. Suppose a firm has K layers and β_{ij} is the share of employees in layer j for firm i. Hierarchization is $1 - \sum_{j=1}^{K} \beta_{ij}^2$. Thus, a flat firm has a value of 0 and a firm with equal shares across K layers approaches 1 as K increases.
Number of unique roles	The number of unique roles in the firm across all years.

Table 2: Comparison of firms in and out of the sample

Notes: The table reports the differences in observables between the firms in the final sample ("In sample") and thus not ("Out of sample"). The set of firms only includes those in Compustat on or after 2015, and not in the following SIC codes: 6000-6999 (financials), 4900-4999 (utilities), 5200,5300-5399 (retail), 7100-7299 (service firms with poor LinkedIn coverage), and 4400-4599 (airlines and shipping). *** denotes p-value $< .01$, ** denotes p-value $< .05$, and * denotes p-value $< .1$ from a two-side t-test for equivalence of means.

	In sample	Out of sample	Diff.	s.e.	obs.
Assets	5293.40	4885.28	408.11	(625.34)	8766
Total liabilities	3294.11	2912.23	381.88	(372.01)	8753
PPE (net)	1469.51	1542.41	-72.90	(223.74)	8763
Debt / equity (winsorized)	0.86	0.62	0.24***	(0.06)	8740
Cash (winsorized)	280.30	205.94	74.36***	(10.64)	8754
Cash / Assets	0.30	0.31	-0.01	(0.01)	8695
Revenue (m)	3562.54	3336.00	226.55	(440.86)	8701
Net income (winsorized)	75.20	58.87	16.32***	(6.14)	8689
EBIT	469.65	356.32	113.32*	(64.41)	8688
Net income / revenue	-0.13	-0.14	0.00	(0.01)	5890
ROA	-0.19	-0.21	0.02**	(0.01)	7911
Employees	6.62	10.37	-3.75***	(1.00)	7481
Firms	2,611	6,145			

Table 3: Sample summary statistics

Notes: The table reports the summary statistics for the variables defined in Table 1. Panel A considers the cross-section of firms using the last year in which their corporate hierarchy was observed. Panel B considers all firm-years.

	Panel A: Cross-section					
	Mean	SD	Min	Median	Max	N
Number of internal labor markets	252.82	440.02	1.00	107.00	5224.00	2,605
Size of largest ILM (roles)	192.48	415.16	1.00	60.00	6045.00	2,605
Number of unique roles	321.38	635.41	2.00	124.00	9364.00	2,605
% role in largest ILM	54.20	15.56	4.00	56.43	100.00	2,605
% worker-years in largest ILM	73.60	19.80	4.21	79.32	98.35	2,235
Employees (Compustat, 1000s)	6.06	23.33	0.00	0.66	733.00	2,519
Number employees (LinkedIn, '000s)	1.68	5.26	0.00	0.29	105.72	2,605
Number hierarchy layers	10.85	9.41	1.00	8.00	80.00	2,605
Hierarchization (German)	0.62	0.07	0.00	0.63	0.75	2,605
Hierarchization (ONet)	0.66	0.07	0.00	0.67	0.80	2,605
Hierarchization (this paper)	0.76	0.21	0.00	0.82	0.96	2,514
	Panel B: Firm-year					
	Mean	SD	Min	Median	Max	N
Number of internal labor markets	281.87	474.17	1.00	121.00	5224.00	17,652
Size of largest ILM (roles)	222.02	457.65	1.00	72.00	6045.00	17,652
Number of unique roles	365.78	697.27	2.00	142.00	9364.00	17,652
% role in largest ILM	55.55	14.98	4.00	57.69	100.00	17,652
% worker-years in largest ILM	75.53	18.70	2.05	80.74	98.35	15,262
Employees (Compustat, 1000s)	7.09	24.52	0.00	0.76	733.00	16,149
Number employees (LinkedIn, '000s)	1.92	5.97	0.00	0.30	116.34	17,652
Number hierarchy layers	11.67	10.03	1.00	9.00	80.00	17,652
Hierarchization (German)	0.62	0.07	0.00	0.63	0.75	17,652
Hierarchization (ONet)	0.66	0.07	0.00	0.67	0.80	17,652
Hierarchization (this paper)	0.76	0.21	0.00	0.83	0.96	15,731

Table 4: Common titles by hierarchy layer

Notes: The table reports the most common titles across all firms and years by the top and bottom layers across all firms.

Bottom layer	**Top layer**
Sales Associate	Technical Support Engineer
Customer Service Representative	Shift Manager
Scientist	Software Engineer
Software Engineer	General Manager
Neuroscience Sales Specialist	Supervisor
Assistant Manager	Store Manager
Server	General Manager
Cashier	Office Manager
Store Manager	Electrical Engineer
Retail Wireless Consultant	Operations Manager

Table 5: Regressions of log # organizational layers on firm characteristics: firm size proxies

Notes: The table reports regression estimations where the dependent variable is the log of the number of hierarchical layers. The unit of observation is a firm-year. The regressions include various measures of firm size, all lagged to the previous fiscal year. "Log business segments" (column 4) is the count of historical business segments as reported in Compustat. Column 5 includes the "Log (# US subsidiaries)" which is the log of the count (1 if missing) of subsidiaries from the WRDS subsidiaries data. Column 6 includes the total number of individuals reported in Execucomp (logged). The final column includes all control variables in one regression (thus, the sample is smaller). Year FE are year fixed effects, Industry FE? are 4-digit SIC fixed effects, and Incorp. State FE are fixed effects for the firms state of incorporation. Standard errors clustered at the company. *** denotes p-value < .01, ** denotes p-value < .05, and * denotes p-value < .1.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8))
Log assets	0.29***							0.037*
	(0.0047)							(0.020)
Log employees (CS, 1000s)		0.35***						0.21***
		(0.0046)						(0.032)
Log revenue			0.25***					0.080***
			(0.0052)					(0.025)
Log business segments				0.27***				-0.031**
				(0.034)				(0.013)
Log (# US subsidiaries)					0.21***			-0.0059
					(0.012)			(0.0047)
Log # execs in Execucomp						0.43***		-0.0068
						(0.058)		(0.030)
Log board size							1.25***	0.11***
							(0.057)	(0.033)
Observations	13648	12590	12465	13516	13669	5465	11344	5111
# firms	2299	2246	2145	2276	2303	785	2072	751
Mean dep. var.	2.15	2.19	2.25	2.15	2.15	2.76	2.24	2.76
R^2	0.78	0.86	0.75	0.38	0.43	0.35	0.52	0.86
Within-R^2	0.66	0.77	0.63	0.034	0.11	0.046	0.25	0.80
Year FE?	Y	Y	Y	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y	Y	Y	Y
Incorp. State FE?	Y	Y	Y	Y	Y	Y	Y	Y

Table 6: Differences in managers

Notes: The table reports regression estimations where the dependent variable is the fraction of managers or supervisory workers. The unit of observation is a firm-year. The odd columns consider the fraction of workers across all employees ("All"). The even columns consider the dependent variable that averages the fraction of managers or supervisors across a firm's hierarchy layers ("Layer"). Year FE are year fixed effects, Industry FE are 4-digit SIC fixed effects, and Asset decile FE are fixed effects for asset-based deciles. Standard errors clustered at the company. *** denotes p-value $< .01$, ** denotes p-value $< .05$, and * denotes p-value $< .1$.

	% managers		% supervisors	
	All	Layer	All	Layer
	(1)	(2)	(3)	(4)
Log # layers	0.018***	0.019***	0.0021	0.0026
	(0.0062)	(0.0060)	(0.0023)	(0.0023)
Observations	15438	15438	15438	15438
# firms	2519	2519	2519	2519
Mean dep. var.	0.22	0.21	0.018	0.020
R^2	0.33	0.30	0.32	0.30
Within-R^2	0.032	0.035	0.0087	0.012
Asset decile?	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y

Table 7: Differences in employee education by number of hierarchical layers

Notes: The table reports regression estimations where the dependent variable measures the fraction of workers (in LinkedIn) with a certain education level. The unit of observation is a firm-year. The first four columns consider the fraction of workers across all employees. Year FE are year fixed effects, Industry FE? are 4-digit SIC fixed effects, and Asset decile FE are fixed effects for asset-based deciles. Standard errors clustered at the company. *** denotes p-value < .01, ** denotes p-value < .05, and * denotes p-value < .1.

	All employees			
	BA	MA	MBA	Adv.
	(1)	(2)	(3)	(4)
Log # layers t-1	0.023**	0.0065	-0.0086**	-0.011**
	(0.0100)	(0.0054)	(0.0042)	(0.0047)
Observations	14927	14927	14927	14927
# firms	2511	2511	2511	2511
Mean dep. var.	0.57	0.13	0.075	0.060
R^2	0.017	0.011	0.0052	0.0068
Asset decile?	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y

Table 8: Differences in new employees, turnover, and tenure

Notes: The table reports regression estimations of measures of new employees and tenure of employees with the previous year (log) of number of hierarchical layers. Under "Employee", column 1 reports the fraction of employees on LinkedIn in the current year. Column 2 considers the fraction of existing employees on LinkedIn that are promoted in the year. The columns "Tenure (log years)" report the log of average tenure for employees in the year. Column 3 considers all employees in the year, including new employees. Column 4 only considers employees who did not join in the current year. The final three columns have dependent variables related to the CEO. The dependent variable in column 5 equals one if the firm had a new CEO who worked there before starting the CEO position (the sample is all CEO turnover events). The dependent variable in column 6 is the log of total compensation (salary and bonus) for the CEO. The final column's dependent variable is the ratio of CEO pay to the median worker. Year FE are year fixed effects, Industry FE? are 4-digit SIC fixed effects, and Asset decile FE are fixed effects for asset-based deciles. Standard errors clustered at the company. *** denotes p-value $< .01$, ** denotes p-value $< .05$, and * denotes p-value $< .1$.

	Employees		Tenure (log years)		Internal	CEO	CEO
	% new	% promote	Non-new	All	CEO	pay	pay ratio
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log # layers t-1	-0.051***	0.042***	0.058**	0.098***	0.38***	0.14**	20.2**
	(0.0063)	(0.0068)	(0.024)	(0.027)	(0.059)	(0.069)	(9.70)
Observations	9540	9536	9536	9540	5708	565	4558
# firms	1432	1432	1432	1432	816	419	1167
Mean dep. var.	0.14	0.18	1.72	1.62	8.60	0.55	132.9
R^2	0.045	0.025	0.032	0.040	0.38	0.013	0.070
Asset decile?	Y	Y	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y	Y	Y

Table 9: Regressions of firm productivity and hierarchical layers

Notes: The table reports regression estimations of firm-year variables. Panel A considers various measures of firm productivity: Return on Assets (column 1), Gross margin (column 2), Net profit margin (column 3), and revenue per employee (column 4). Panel B reports regressions with different cost measures (scaled by current-year assets). All variables are winsorized at the 5% level. Asset decile FE are fixed effects for firm assets deciles in the previous year, Year FE are year fixed effects, and Industry FE are 4-digit SIC fixed effects. Standard errors clustered at the company. *** denotes p-value < .01, ** denotes p-value < .05, and * denotes p-value < .1.

	Productivity				Costs	
	ROA	Gross margin	Net margin	Rev/emp.	SG&A	Wages
	(1)	(2)	(3)	(4)	(5)	(6)
Log # layers t-1	0.0243***	0.0978***	0.519***	-88.31***	0.0771***	0.00214
	(0.00833)	(0.0332)	(0.0969)	(24.73)	(0.0110)	(0.0183)
Observations	13061	13470	13470	14021	12967	1028
# firms	2320	2349	2349	2438	2312	204
Mean dep. var.	-0.0549	0.196	-0.984	464.9	0.373	0.245
R^2	0.589	0.416	0.423	0.568	0.536	0.537
Within-R^2	0.287	0.0421	0.130	0.0862	0.336	0.122
Asset 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 10: Acquisition activity

Notes: The table reports regression estimations of firm acquisition activity and its relationship with lagged hierarchical layers. The dependent variable in column 1 is an indicator of whether the firm has an acquisition (of any type) as reported by SDC. Column 2 only considers acquisitions with reported valuations, which tend to be public targets or larger acquisitions. The third column reports Poisson estimates for the total number of acquisitions in the year. Column 4's dependent variable is the log of the sum of all reported valuations in acquisitions (excluded if no valuation is reported). The final column presents a dependent variable that equals one if the acquirer's primary SIC in SDC does not match any of the SIC codes of the target. Asset decile FE are fixed effects for firm assets deciles in the previous year, Year FE are year fixed effects, and Industry FE are 4-digit SIC fixed effects. Standard errors clustered at the company. *** denotes p-value $< .01$, ** denotes p-value $< .05$, and * denotes p-value $< .1$.

	(1) Acq.?	(2) Acq. w/ value?	(3) # Acq. (all)	(4) $\log(\sum V)$	(5) Diversify?
Log # layers t-1	0.018**	0.014**	0.39***	0.49***	0.030
	(0.0074)	(0.0056)	(0.13)	(0.14)	(0.042)
Observations	14678	14678	12877	1024	1711
# firms	2499	2499	2179	532	685
Mean dep. var.	0.12	0.073	0.22	4.69	0.62
R^2	0.020	0.0096		0.43	0.0059
$Pseudo - R^2$			0.16		
Model	OLS	OLS	Poisson	OLS	OLS
Asset decile FE?	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y

Table 11: Firm patenting

Notes: The table reports regression estimations where the dependent variables are R&D or patent outcomes (counts and citations). Except for column 1, all specifications are Poisson. The right-hand side variable is the log of the number of hierarchical layers measured in the prior year. The unit of observation is a firm-year. The dependent variable in Column 1 is R&D scaled by firm assets. Columns 2-4 presents Poission estimates where the dependent variable is the count of the number of patents applied for in the current year (including 0). Column 5-7's dependent variable are the average number of citations received by the end of the sample, averaged across all patents for the firm in a given year. Asset decile FE are fixed effects for firm assets deciles in the previous year, Year FE are year fixed effects, and Industry FE are 4-digit SIC fixed effects. Standard errors clustered at the company. *** denotes p-value $< .01$, ** denotes p-value $< .05$, and * denotes p-value $< .1$.

	(1)	(2)	(3)	(4)	(5)
	R&D	# patents		# cites	
Log # layers t-1	0.013*	2.20***	2.16***	0.11	0.13
	(0.0078)	(0.45)	(0.50)	(0.26)	(0.29)
Leverage t-1		0.80		-0.41	
		(0.53)		(0.29)	
Cash/Assets t-1		1.18**		1.44***	
		(0.51)		(0.30)	
Q t-1			0.18***		0.11**
			(0.058)		(0.050)
ROA t-1			-0.88		0.098
			(0.70)		(0.50)
Observations	10611	10654	8298	2525	2338
# firms	1825	2041	1782	582	558
Mean dep. var.	0.17	19.3	23.7	0.96	0.94
Pseudo-R^2		0.87	0.88	0.33	0.31
Asset decile?	Y	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y

Table 12: Stock and firm volatility

Notes: The table reports regression estimations where the dependent variable in Panel A is the square root of the sum of squared daily stock returns over the year. Panel B considers dependent variables related to operational and financial risk. Column 1 uses operating asset volatility and column 2 cash flow volatility. Column 3's dependent variable is debt to assets (leverage) and column 4 uses cash to assets. The right hand side variable is the log of the number of hierarchical layers measured in the prior year. The unit of observation is a firm-year. sset decile FE are fixed effects for firm assets deciles in the previous year, Year FE are year fixed effects, and Industry FE are 4-digit SIC fixed effects. Standard errors clustered at the company. *** denotes p-value $< .01$, ** denotes p-value $< .05$, and * denotes p-value $< .1$.

	Panel A: Stock volatility			
	(1)	(2)	(3)	(4)
Log # layers	-0.059***	-0.23**	-0.055***	-0.16***
	(0.014)	(0.099)	(0.012)	(0.055)
Observations	14324	14171	14324	14171
# firms	2398	2245	2398	2245
Mean dep. var.	0.62	0.62	0.62	0.62
R^2	0.38	0.53	0.44	0.60
Within-R^2	0.22	0.014	0.26	0.018
Asset decile?	Y	Y	Y	Y
Firm FE?	N	Y	N	Y
Year FE?	Y	Y	Y	Y
	Panel B: Operations and financial			
	σ(OA)	σ(CF)	$\frac{Debt}{Assets}$	$\frac{Cash}{Assets}$
	(1)	(2)	(3)	(4)
Log # layers t-1	-0.047**	-0.011*	0.013	-0.013
	(0.023)	(0.0057)	(0.0087)	(0.0090)
Observations	12636	12940	14279	14367
# firms	2306	2399	2494	2498
Mean dep. var.	0.56	0.061	0.26	0.31
R^2	0.087	0.24	0.053	0.11
Asset decile?	Y	Y	Y	Y
Industry FE?	Y	Y	Y	Y
Firm FE?	N	N	N	N
Year FE?	Y	Y	Y	Y

Table 13: Three-factor regression estimates: hierarchy-sorted portfolios

Notes: The table reports three-factor regression estimations using monthly returns from portfolios sorted by the (lagged) number of hierarchical layers. In each year, we sort firms by the number of layers and construct a long-short portfolio of the firms with the most (long) versus the least (short). The right-hand side variables include a constant (α) and the six factors from Ken 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. *** denotes p-value $< .01$, ** denotes p-value $< .05$, and * denotes p-value $< .1$.

	(1)	(2)	(3)	(4)	(5)	(6)
α	-0.0048	-0.0076**	-0.0095***	-0.0076**	-0.0079**	-0.010***
	(0.0045)	(0.0035)	(0.0033)	(0.0035)	(0.0035)	(0.0033)
$R^m - r^f$	-0.29**	-0.079	-0.14	-0.078	-0.048	-0.099
	(0.12)	(0.094)	(0.096)	(0.089)	(0.092)	(0.089)
SMB		-1.16***	-0.88***	-1.16***	-1.13***	-0.82***
		(0.11)	(0.13)	(0.11)	(0.12)	(0.13)
HML		-0.13	-0.19**	-0.14	-0.099	-0.12
		(0.085)	(0.087)	(0.12)	(0.092)	(0.14)
RMW			0.65***			0.68***
			(0.17)			(0.18)
CMA				0.017		-0.058
				(0.25)		(0.26)
Momentum					0.10	0.15
					(0.098)	(0.096)
Observations	108	108	108	108	108	108
R^2	0.076	0.45	0.50	0.44	0.45	0.50

Table 14: Regressions of log # organizational layers on exposure to pandemic

Notes: The table reports regression estimations where the dependent variable is the log of the number of hierarchical layers. The sample period is 2017 through 2022. The right hand side variables are measures of exposure to the Covid pandemic. The variable "Covid" is equal to one if the firm is in the SIC code 2834, "Pharmaceutical Preparations." The unit of observation is a firm-year. The last two columns narrow the control group to firms in that are not in industries directly impacted by the pandemic (travel, hotels, retail, etc.). Asset decile FE are fixed effects for firm assets deciles in the previous year, Year FE are year fixed effects, and Firm FE are firm fixed effects. Standard errors clustered at the company. *** denotes p-value < .01, ** denotes p-value < .05, and * denotes p-value < .1.

	Log # layers					
	(1)	(2)	(3)	(4)	(5)	(6)
Covid period	0.031*		0.024		0.022	
	(0.016)		(0.016)		(0.016)	
Covid X 2017		0.0048		0.0056		0.0091
		(0.023)		(0.023)		(0.023)
Covid X 2018		0.016		0.016		0.020
		(0.016)		(0.016)		(0.016)
Covid X 2020		0.023		0.021		0.021
		(0.014)		(0.014)		(0.014)
Covid X 2021		0.039**		0.031		0.031
		(0.019)		(0.019)		(0.019)
Covid X 2022		0.051**		0.045**		0.044**
		(0.021)		(0.021)		(0.021)
Observations	12093	12093	11841	11841	9514	9514
# firms	2338	2338	2318	2318	1864	1864
Mean dep. var.	2.11	2.11	2.11	2.11	2.03	2.03
R^2	0.99	0.99	0.99	0.99	0.99	0.99
Within-R^2	0.0016	0.0019	0.043	0.043	0.043	0.043
Asset decile?	N	N	Y	Y	Y	Y
Firm FE?	Y	Y	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y	Y	Y

Table 15: Firm response to AI: number of hierarchical layers

Notes: This table reports regressions of the log of the number of hierarchical layers on various measures of AI adoption. Asset decile FE are fixed effects for the previous year firm asset decile, Year FE are year fixed effects, and Firm FE are firm fixed effects. Standard errors clustered at the company. *** denotes p-value $< .01$, ** denotes p-value $< .05$, and * denotes p-value $< .1$.

	Log number of hierarchical layers			
	% posts AI			
	AI	Narrow	Category	Max %
	(1)	(2)	(3)	(4)
Share job postings AI	-0.049			
	(0.034)			
Share job postings AI (narrow)		-0.050		
		(0.031)		
Average share (skills)			-0.050	
			(0.034)	
Avg max share				-0.011
				(0.010)
Observations	3205	3205	3205	3205
# firms	869	869	869	869
Mean dep. var.	2.54	2.54	2.54	2.54
R^2	0.0020	0.0019	0.0020	0.0020
Asset decile?	Y	Y	Y	Y
Year FE?	Y	Y	Y	Y
Firm FE?	Y	Y	Y	Y
Industry FE?	N	N	N	N

Appendix

A1 Known job ladders: investment banks

We rely on the representativeness of LinkedIn with regard to the 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.[19]

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.

[19] 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 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 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 compare to the new position. Appendix Table A6 presents the transition matrix. At least 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 layer

The figure reports the average ONet job zone by hierarchy layer (see `https://www.onetonline.org/find/all`). A job zone is "group occupation[] into one of five categories based on levels of education, experience, and training necessary to perform the occupation." Each line represents different total layers in the firm (i.e., longer lines are a sample of firms with more layers).

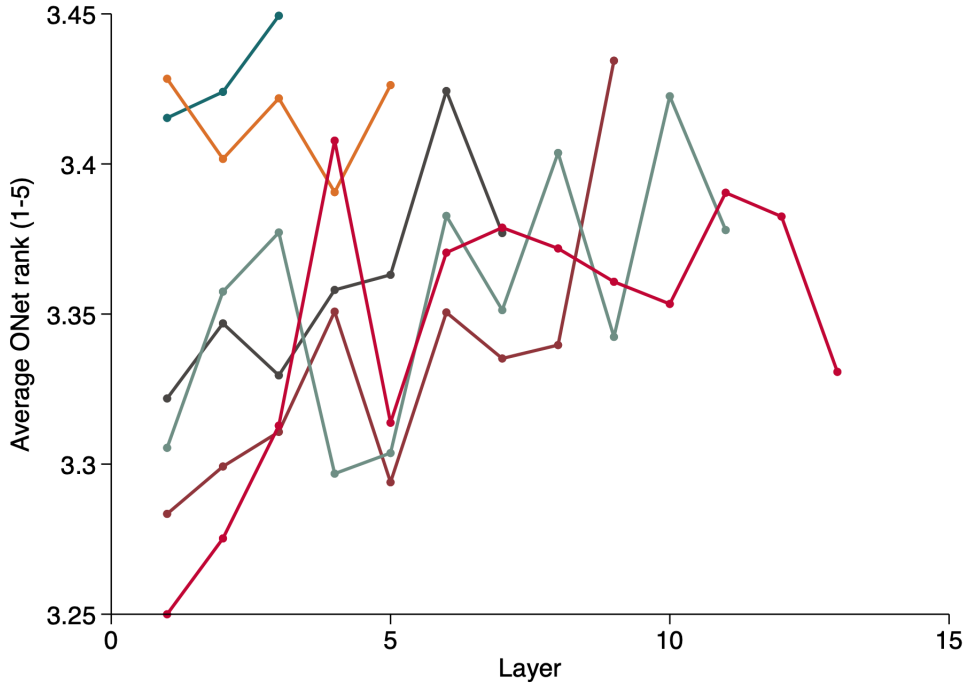


Figure A2: Constructing layers using ONet

The figure reports the average and median ONet-based layers using its job zone categorization (see `https://www.onetonline.org/find/all`). A job zone is "group occupation[] into one of five categories based on levels of education, experience, and training necessary to perform the occupation." Each line represents different total layers in the firm (i.e., longer lines are a sample of firms with more layers). We map cleaned titles to the best match job zone and calculate the number of layers. There are at most five layers.

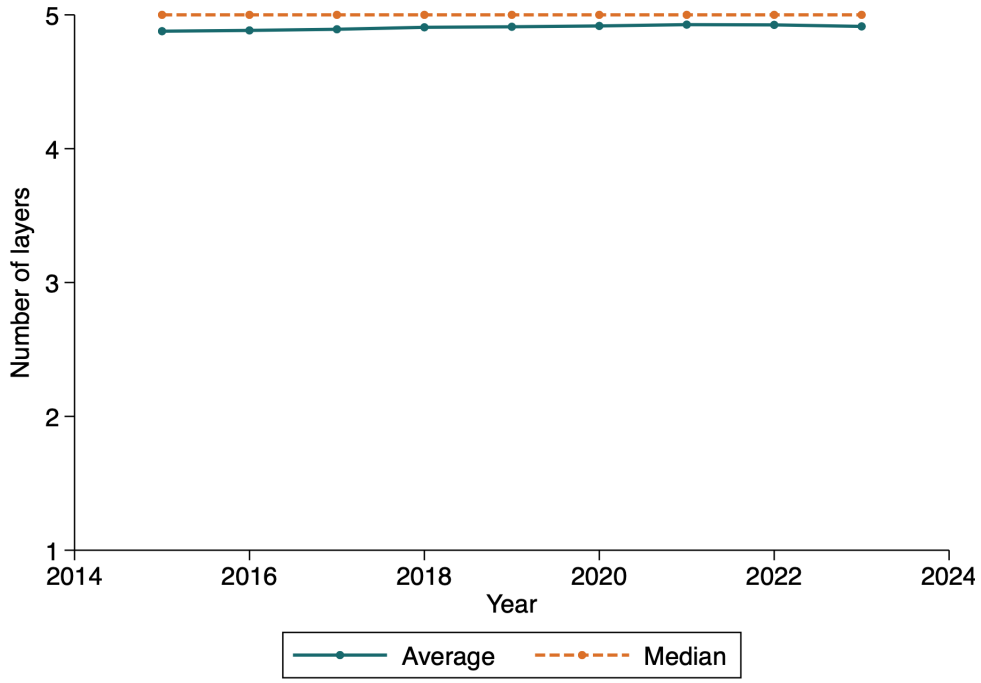


Figure A3: Span of control by hierarchy quintile and time

The figure reports the average span of control for corporate hierarchies split into quintiles. We split the hierarchy into quintiles for all firms with at least five layers. The span of control of each quintile (except the topmost) is calculated as the share of employees in the layer divided by the share in the quintile above. The "4th" is the penultimate layer. These averages are computed by quintile and year.

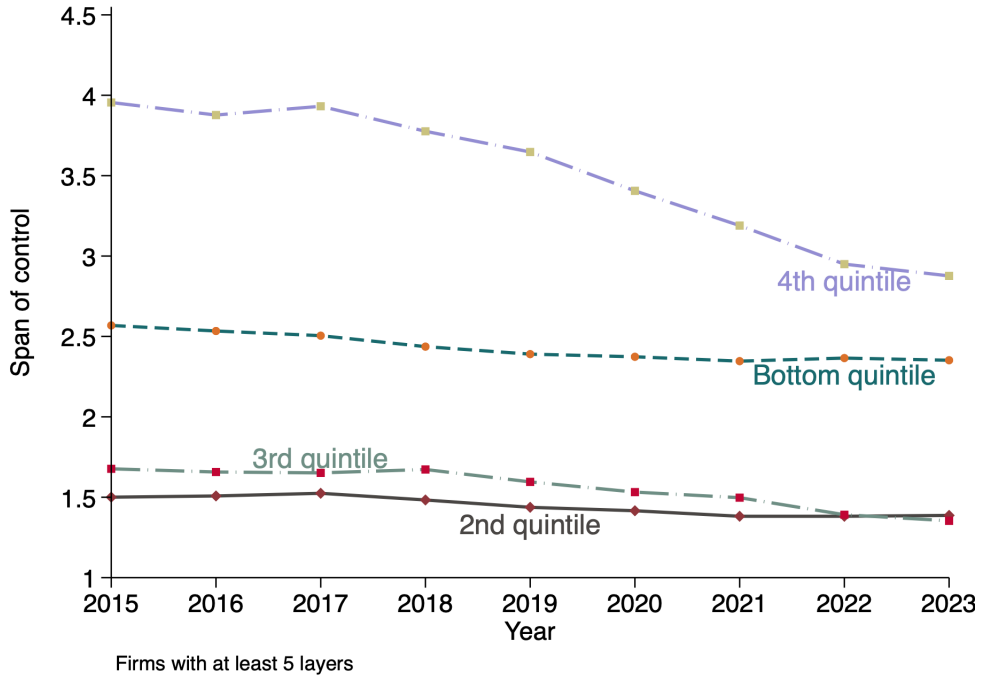


Figure A4: Promotion rates over hierarchy layers

The figure reports the fraction of all promotions in a layer (across all years) that are internal to the layer (dashed line) and to the next layer above (solid line).

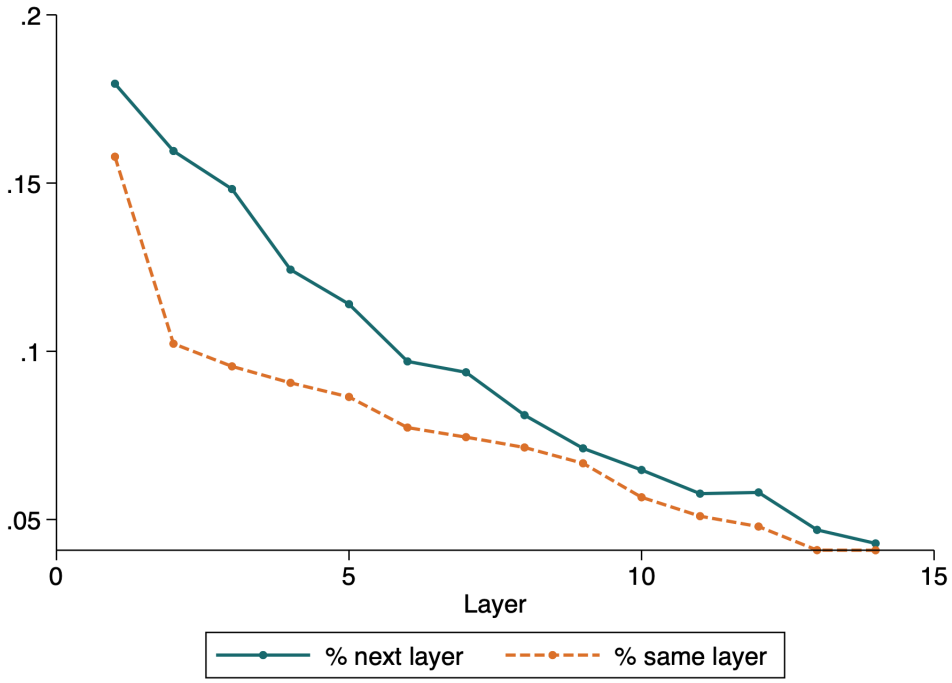


Figure A5: Years to promotion over hierarchy layers

The figure reports the average years for a worker in each layer to receive any promotion (measured using role ranks).

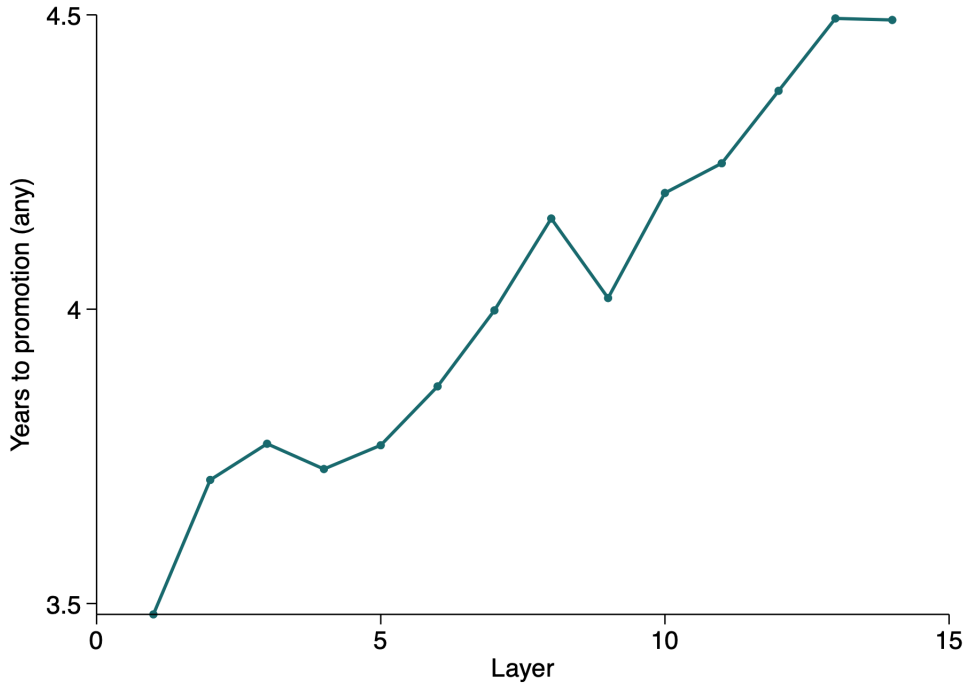


Table A1: Measurement of hierarchy in organizational studies

The table summarizes previous research that measures alternative forms of the corporate hierarchy."Depth" is "defined at the division level and represents the vertical dimension, or steepness, of the hierarchy" which is "the number of positions between the CEO and the division manager." (Guadalupe and Wulf, 2010)

Paper	**Measure**	# layers	**Country**	**Method(s)**
Adhvaryu, Bassi, Nyshadham, Tamayo, and Torres (2023)	# layers / span of control	16 (working group)	Argentina (one firm)	Firm-specific org. chart
Baker, Gibbs, and Holmstrom (1994)	Hierarchical Levels	8	US (single firm)	Job titles with transitions
Bias, Lochner, Obernberger, and Sevilir (2024)	# of 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 Levels	4	Denmark	Occupational codes
Garicano and Hubbard (2016)	Associate-Partner Ratio in law firms	2	US	1992 Census of Services (title code and wages)
Guadalupe, Li, and Wulf (2014)	# CEO reports and # managers	N/A (span)	US	Direct observation (Hewitt Associates, HR consulting)
Guadalupe and Wulf (2010)	Depth / span of control	N/A (span)	US	Direct observation (Hewitt Associates, HR consulting)
Gumpert, Steimer, and Antoni (2022)	# of Hierarchical Layers / Share of Managerial Occupations	4	Germany	Occupational codes
Huitfeldt, Kostøl, Nimczik, and Weber (2023)	Number of Hierarchy Levels	Max 66, mean 7	Norway	Job titles with transitions
Liberti and Mian (2009)	Approval Level	5	Argentina	Direct observation of reporting
Rajan and Wulf (2006)	CEO Span of Control / Depth	N/A	US (300 firms)	Direct observation (Hewitt Associates, HR consulting)
Skrastins and Vig (2018)	# managerial layers	8	India (single bank)	Bank reports
Tag, Astebro, and Thompson (2016)	Number of Layers / Span of Control	4	Sweden	Occupational codes

Table A2: Top 10 Most Common Titles and Frequency Percentages for Each Quartile

The table reports the 10 most common strings in titles for hierarchical layer quartiles (for firms with at least 4 layers).

Quartile 1		Quartile 2		Quartile 3		Quartile 4	
Title	Frequency	Title	Frequency	Title	Frequency	Title	Frequency
manager	9.13	manager	14.46	manager	16.19	manager	13.76
engineer	7.34	engineer	5.19	senior	4.92	senior	5.14
senior	5.17	senior	4.97	engineer	4.27	engineer	5.06
sales	4.95	sales	3.44	director	4.24	director	4.67
specialist	3.18	director	3.27	sales	2.87	president	3.68
analyst	3.06	analyst	2.54	project	2.47	vice	3.65
associate	2.77	project	2.36	sr	2.19	software	3.54
account	2.03	specialist	2.26	operations	2.09	sales	2.95
sr	2.02	sr	2.24	program	1.99	program	2.10
representative	1.98	lead	2.16	product	1.89	project	2.05

Table A3: Transition matrix between layer quintiles

Notes: This table reports the transition probability from roles by layer quintile. A cell reports the fraction of t to $t+1$ jobs transition to all other quintiles (including the same one in the diagnol) in the next year.

	Layer quintile $t+1$				
	1.0	2.0	3.0	4.0	5.0
1	**90.9**%	3.4%	2.4%	2.1%	1.2%
2	6.8%	**84.8**%	3.5%	3.1%	1.8%
3	6.5%	3.8%	**84.2**%	3.5%	2.1%
4	5.7%	3.4%	2.9%	**85.9**%	2.1%
5	5.3%	3.2%	2.8%	2.7%	**86.1**%

Table A4: Most layers by year and industry

Notes: The table reports the firms in 2023 with the most hierarchical layers (Panel A), among IT firms (Panel B), and drug/chemical firms (Panel C).

	Panel A: All firms	
	# hierarchy layers	Employees ('000s)
Intl Business Machines Corp	80	282.2
Microsoft Corp	76	221
Apple Inc	70	161
General Motors Co	65	163
Ford Motor Co	62	177
Intel Corp	61	124.8
Accenture Plc	60	733
Lockheed Martin Corp	58	122
Boeing Co	57	171
Chevron Corp	56	45.6
Pfizer Inc	56	88
	Panel B: IT firms	
	# hierarchy layers	Employees ('000s)
Intl Business Machines Corp	80	282.2
Microsoft Corp	76	221
Cognizant Tech Solutions	49	347.7
Sap Se	45	107.602
Jacobs Solutions Inc	42	60
Leidos Holdings Inc	40	47
Automatic Data Processing	36	63
Fiserv Inc	36	42
Intuit Inc	35	18.2
Thomson-Reuters Corp	35	25.6
Caci Intl Inc -Cl A	33	23
	Panel C: Drug/chemical firms	
	# hierarchy layers	Employees ('000s)
Pfizer Inc	56	88
Merck & Co	51	72
Gsk Plc	50	70.212
Bristol-Myers Squibb Co	48	34.1
Sanofi	48	86.088
Abbvie Inc	45	50
American Vanguard Corp	44	.845
Dow Inc	43	35.9
Amgen Inc	40	26.7
Dupont De Nemours Inc	38	24
Sherwin-Williams Co	35	64.088

Table A5: Least layers in 2023 and by industry

Notes: The table reports the firms in 2023 with the least hierarchical layers (Panel A), among IT firms (Panel B), and drug/chemical firms (Panel C).

	Panel A: All firms	
	# hierarchy layers	Employees ('000s)
Gold Resource Corp	1	.488
Innoviva Inc	1	.112
Iteos Therapeutics Inc	1	.157
Aquabounty Technologies Inc	2	.104
Decisionpoint Systems Inc	2	.21
Intellinetics Inc	2	.182
Interlink Electronics	2	.107
National Energy Services Reu	2	6.231
Pro-Dex Inc/Co	2	.146
Surgepays Inc	2	.262
	Panel B: IT firms	
	# hierarchy layers	Employees ('000s)
Decisionpoint Systems Inc	2	.21
Ftai Aviation Ltd	2	.17
Intellinetics Inc	2	.182
Glimpse Group Inc(The)	3	.186
Omniq Corp	3	.182
Research Solutions Inc	3	.145
Riot Platforms Inc	3	.534
Staffing 360 Solutions Inc	3	.15
Arena Group Hldngs Inc (The)	4	.448
Simulations Plus Inc	4	.197
	Panel C: Drug/chemical firms	
	# hierarchy layers	Employees ('000s)
Innoviva Inc	1	.112
Iteos Therapeutics Inc	1	.157
Acelyrin Inc	2	.13
Aquabounty Technologies Inc	2	.104
Fortress Biotech Inc	3	.186
Geron Corp	3	.141
Keros Therapeutics Inc	3	.136
Monte Rosa Therapeutics Inc	3	.133
Phathom Pharma Inc	3	.452
Vbi Vaccines Inc	3	.138

Table A6: Transition matrix between between firms

Notes: This table reports the transition probability from roles by layer quintile between firms. A cell reports the fraction of jobs transition between firms from (the left) to quintiles in the new firm (top).

	Layer quintile new firm				
	1.0	2.0	3.0	4.0	5.0
1	**53.5**%	17.6%	12.5%	10.8%	5.6%
2	38.3%	**22.7**%	16.4%	14.3%	8.3%
3	36.5%	21.2%	**17.2**%	16.4%	8.7%
4	31.9%	19.6%	17.2%	**20.4**%	10.9%
5	29.5%	19.6%	16.3%	20.5%	**14**%