

Huitfeldt, Ingrid; Kostøl, Andreas Ravndal; Nimczik, Jan Sebastian; Weber, Andrea

Working Paper

Internal labor markets: A worker flow approach

Discussion Papers, No. 961

Provided in Cooperation with:

Research Department, Statistics Norway, Oslo

Suggested Citation: Huitfeldt, Ingrid; Kostøl, Andreas Ravndal; Nimczik, Jan Sebastian; Weber, Andrea (2021) : Internal labor markets: A worker flow approach, Discussion Papers, No. 961, Statistics Norway, Research Department, Oslo

This Version is available at:

<https://hdl.handle.net/10419/250128>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Internal labor markets: A worker flow approach

Ingrid Huitfeldt, Andreas R. Kostøl, Jan Nimczik, Andrea Weber

TALL

SOM FORTELLER

DISCUSSION PAPERS

961

*Ingrid Huitfeldt, Andreas R. Kostøl,
Jan Nimczik, Andrea Weber*

Internal labor markets: A worker flow approach

Abstract:

This paper develops a new method to study how workers' career and wage profiles are shaped by internal labor markets (ILM) and job hierarchies in firms. Our paper tackles the conceptual challenge of organizing jobs within firms into hierarchy levels by proposing a data-driven ranking method based on observed worker flows between occupations within firms. We apply our method to linked employer-employee data from Norway that records fine-grained occupational codes and tracks contract changes within firms. Our findings confirm existing evidence that is primarily based on case studies for single firms. We expand on this by documenting substantial heterogeneity in the structure and hierarchy of ILMs across a broad range of large firms. Our findings on wage and promotion dynamics in ILMs are consistent with models of careers in organizations

Keywords: Internal Labor Markets; Organization of Labor; Wage Setting.

JEL classification: J31, J62, M5

Acknowledgements: This project received financial support from the Norwegian Research Council through grant 227115 and 275123. We would like to thank our discussants Thomas Lemieux, Kathryn Shaw and Michael Waldman, as well as Robert Gibbons, Patrick Kline and Daniel Silverman for very constructive feedback and participants in several seminars and workshops for many useful suggestions.

Address: Ingrid Huitfeldt, Statistics Norway, Research Department. E-mail: ish@ssb.no

Discussion Papers

comprise research papers intended for international journals or books. A preprint of a Discussion Paper may be longer and more elaborate than a standard journal article, as it may include intermediate calculations and background material etc.

The Discussion Papers series presents results from ongoing research projects and other research and analysis by SSB staff. The views and conclusions in this document are those of the authors

Sammendrag

Denne artikkelen studerer hvordan lønnsvekst påvirkes av jobbytter innad i foretak i Norge. Vi utvikler en ny metode for å estimere foretaks karrierestiger uten direkte kjennskap til organisasjonskart ved bruk av relative strømmer mellom stillingstitler i foretak og administrative registerdata. Vi anvender metoden vår på alle store private foretak i Norge. Resultatene våre kan oppsummeres med tre brede konklusjoner:

Vi dokumenterer stor heterogenitet i organisasjonsstrukturer mellom foretak. De fleste foretakene har flere interne arbeidsmarkeder i tillegg til enkeltstående yrker som ikke er tilknyttet det interne arbeidsmarkedet. Antall interne arbeidsmarkeder og enkeltstående yrker øker med størrelsen på foretaket. De fleste arbeidere og jobbstillinger tilhører foretakets største interne arbeidsmarked, men strukturen på det interne arbeidsmarkedet varierer på tvers av foretak. Mens den lengste karrierestigen har opptil 66 trinn, har omtrent halvparten av foretakene i vårt utvalg tre til fire trinn. Antall trinn øker når foretakets størrelse og antall involverte yrker øker.

Vi identifiserer mønstre som er konsistent med økonomisk teori for interne arbeidsmarkeder (f.eks. Doeringer & Piore, 1971) og karrierestiger (f.eks. Garicano & Rossi-Hansberg, 2006). Ifølge disse teoriene entrer arbeideren foretaket nederst på karrierestigen og jobber seg oppover mot mer komplekse jobber over tid. Dette viser også våre empiriske resultater. Vi viser at karrierestigene typisk har en pyramidestruktur, der det er relativt mange sysselsatte på de lavere trinnene på stigen, men færre mot toppen. Vi finner også at andelen lederstillinger øker over karrierestigen, som indikerer at kompleksiteten i arbeidsoppgaver øker. I tråd med modeller om inngangsporter til foretak, viser vi at andel ansettelser utenifra er høyest nederst på karrierestigen. Andelen ansettelser utenifra synker mot toppen av stigen, der de fleste ansettelser skjer internt. I tråd med teorier om karrierer innead i foretak, finner vi at omtrent en tredjedel av nyansatte får forfremmelse i løpet av de fem første årene, og at ansiennitet øker langs foretaks karrierestige.

Til sist finner vi en sterk sammenheng mellom foretaks karrierestige og individuelle lønninger. Gjennomsnittlig log lønn øker tilnærmet lineært i karrierestigen. Denne sammenhengen holder også etter å ha kontrollert for alder, utdanning og ansiennitet. Vi dekomponerer denne sammenhengen videre ved å estimere personfaste effekter i en AKM (Abowd et al., 1999), og finner at personer med høyere faste effekter har høyeresannsynlighet å bli forfremmet.

1 Introduction

While a large empirical literature studies how worker mobility across firms and industries affects wages (e.g., [Krueger & Summers, 1988](#) and [Abowd *et al.*, 1999](#)), there is much less evidence on how worker careers and wages are shaped by the internal labor markets of firms. Early theoretical work characterizes internal labor markets as collections of jobs in firms (see, e.g., [Doeringer & Piore, 1971](#)), where promotion along internal job ladders follows fixed rules set by the firm. These rules can be used to induce effort ([Lazear & Rosen, 1981](#)) and to allocate talent within the firm (e.g., [Sattinger, 1975](#) and [Gibbons & Waldman, 1999b](#)). The lack of empirical evidence is unfortunate as life-cycle labor market outcomes of workers are shaped by factors from both internal labor markets and external labor markets ([Topel & Ward, 1992](#)).

One explanation for the scarcity of evidence is that classifying internal labor markets and job ladders within firms is challenging from both a conceptual and practical standpoint. Except for management positions, the main purpose of occupational codes or job titles is to describe tasks and they are not designed to give a clear representation of the hierarchical order within the firm. As a result, the reconstruction of firm-specific job ladders requires a substantial amount of knowledge about the organizational charts that map the hierarchy structures in companies. The second, more practical problem concerns data availability. Most large-scale linked employer-employee registers do not record detailed occupations or occupational changes within firms. The tedious process of collecting data and evaluating job titles has limited the existing evidence to case studies from particular firms. While [Baker *et al.* \(1994\)](#) provided the first evidence on job ladders using changes in job titles within one particular firm, evidence from a broader population of firms is still missing.

In this paper, we contribute to the conceptual issue by proposing a two-step method of measuring the internal career structure of firms. Our method is based on observed worker flows between fine-grained occupational codes or job titles within firms. We combine recent advancements in panel data estimators and clustering techniques, which have been applied to worker transitions between firms, to estimate the organizational structure *within* firms. The key to our method is having access to detailed matched employer-employee data with information on changes in job titles of individual workers in firms. In the first step, we identify firm-specific networks of occupations based on observed occupational transitions within firms. Occupations that are connected by worker flows form internal labor markets (henceforth ILMs). We allow these networks to be segmented, so that a firm can consist of multiple ILMs.¹ Our method allows for measurement error in the coding of rare occupational transitions, which could potentially influence the shape of the ILMs. In particular, we apply a pruning algorithm related to

¹The counterpart to an internal labor market is a connected set of firms based on between-firm worker mobility in the terminology of [Abowd *et al.* \(1999\)](#) and [Abowd *et al.* \(2002\)](#).

the method used by [Kline *et al.* \(2020\)](#), which checks if removing a single worker breaks an ILM into further sub-markets.

In the second step, we establish a ranking of occupations within the ILM by exploiting the direction of internal network links and the flow frequencies between occupations. Our approach builds on an intuitive idea by Baker, Gibbs and Holmstrom ([1994](#), henceforth BGH): If many employees move from occupation b to occupation a within a firm but few, if any, move from a to b , the relative flows indicate that occupation a ranks higher than b . A clear path from c to b to a indicates a strong hierarchy in the internal labor market. We call such a pattern a job ladder. In contrast, an intransitive ranking with flows in multiple directions indicates flat organizations where job rotation plays a bigger role. To systematically estimate occupational ranks within each firm, we apply a ranking algorithm that computes a hierarchy score based on the fraction of upward moves along the job ladder. We then minimize the number of downward moves over all possible rankings following the Markov Chain Monte Carlo procedure in [Clauset *et al.* \(2015\)](#). Finally, we group occupations into hierarchy levels to distinguish lateral moves from promotions and demotions. We cluster occupations with similar estimated ranks into a hierarchy level using a k-means clustering algorithm, and follow the data-driven approach in [Bonhomme *et al.* \(2019\)](#) to choose the optimal number of levels. This allows us to summarize firms' organizational structure by a single statistic, the number of hierarchy levels.

We overcome the measurement challenge by leveraging rich administrative data from Norway with information on worker mobility both across firms and between job titles (fine-grained 7-digit occupations) within firms. About one-third of all job changes in our data occur within firms, and within-firm job changes are associated with larger wage increases than firm switches on average. This observation motivates our analysis of promotion dynamics in ILMs.

We implement our approach on a set of 3,611 large private sector firms in Norway. Our empirical analysis can be summarized with three broad conclusions. First, we document heterogeneity of organizational structures across firms. In particular, we find that the vast majority of Norwegian firms have multiple ILMs and a few single occupations not connected to the main ILM. The number of ILMs and single occupations are increasing with firm size. Overall, the majority of workers and occupations in firms are employed in the largest ILM. At the same time, the structure of ILMs differs widely across firms. While the longest job ladders have up to 66 hierarchy levels, about half of our sample of firms have three to five levels. We find that the length of the hierarchy increases in firm size and the number of occupations employed by the firm.

Second, the ILMs we identify in the data are broadly consistent with theories of internal labor markets (e.g., [Doeringer & Piore, 1971](#)) and of hierarchies in the organization of labor (e.g., [Garicano & Rossi-Hansberg, 2006](#)). According to these theories, workers enter the firm at the bottom of the hierarchy and then move up the career ladder towards more complex jobs over time. We find several pieces of evidence showing that our estimated job ladders are consistent with these models. Hierarchies have a pyramidal structure, where employment is concentrated at the lower rungs of the ladder, and employment shares decline toward the top. We also find the share of management positions increases over the hierarchy, indicating that task complexity increases. In line with the idea of ports of entry, we document that the share of external hires is highest at the lowest hierarchy level. This share declines toward the top of the hierarchy, where most workers are hired internally from lower levels of the hierarchy. Consistent with careers in organizations, about one-third of recently hired workers are promoted during the first five

years, and average tenure increases with levels of the hierarchy in the firm.

Third, we document a strong link between the internal hierarchy and individual wages. We find that average log wages increase almost linearly with the hierarchy level. This association holds for the average firm and within firms, and after flexibly controlling for age, education, and tenure. We further decompose this relationship by estimating person-fixed effects from AKM wage regressions and find that higher ability individuals are more likely to be promoted to higher levels of the hierarchy.

Our paper is closely related to research on organizational structures and career progressions within firms. In the original theory of [Doeringer & Piore \(1971\)](#), internal labor markets are characterized by ports of entry and exit as the only points of interaction between internal and external markets. In contrast to a spot market where workers are paid their marginal productivity, workers have careers within firms and receive wages attached to the job characteristics and are not subject to influence from the outside market.² [Baker *et al.* \(1994\)](#) offers the first empirical assessment of these concepts using personnel data from a specific firm. To the best of our knowledge, we are the first to apply this idea to a broader set of firms.

Evidence on the hierarchical structure of jobs and wages from personnel records of individual firms has further triggered a wide variety of theoretical approaches (see, e.g., [Gibbons, 1998](#) for an overview). In the seminal work of [Lazear & Rosen \(1981\)](#) and [Waldman \(1984\)](#), firms offer wage premiums to promoted workers to elicit effort, incentivize human capital investment, and prevent employees from being poached by competitors. Our evidence supports these theories but is also consistent with assignment models where a firm assigns more talented individuals to positions higher in the hierarchy (e.g., [Gibbons & Waldman, 1999b](#)).³ While distinguishing between competing theories is beyond the scope of this paper, we develop a method that permits a fresh take on longstanding questions, such as the assessment of the relationship between the firm's hierarchy structure and wage contracts (see, e.g., and [Chiappori & Salanie](#) and [Lazear & Shaw, 2009](#) for surveys).⁴

A small literature develops alternative measures of internal labor markets and studies labor market impacts of ILMs beyond single firms. In an early study, [McCue \(1996\)](#) uses survey data and self-reported promotions by either workers or employers, and [Lazear & Oyer \(2004\)](#) and [Van der Klaauw & Da Silva 2011](#) define promotions from broad occupational groups. More recent research studies how internal connections between larger organizational units such as business groups or establishments offer workers partial insurance against external economic shocks (e.g., [Cestone *et al.*, 2019](#) and [Giroud & Mueller, 2019](#)). Our approach provides a measure of ILMs using fine-grained occupational codes. To assess the importance of highly disaggregated occupational codes, we repeat our analysis on different levels of aggregation. We conclude that fine-grained codes perform somewhat better in explaining promotion dynamics within particular firms. At the same time, we show that many elements of ILMs and hierarchies

²[Gibbons & Waldman \(1999b\)](#) extends the theory by providing an integrated theory of job assignment, human capital attainment and learning.

³Consistent with predictions from tournament theory, where wage spreads must rise with the number of workers to compensate for the increased competition for higher-ranked jobs, [Gabaix & Landier \(2008\)](#) show using cross-country data that the size of firms explains the bulk of differences in compensation. [Eriksson \(1999\)](#), [Bognanno \(2001\)](#), [Garicano & Hubbard \(2007\)](#) and [DeVaro & Kauhanen \(2016\)](#) provide further empirical support for the predictions of these models, and [Lazear *et al.* \(2015\)](#) use company-based data and estimate that CEOs are paid many times more, but are only 1.75 times as productive as the average worker, suggesting that incentives explain the bulk of the variation in compensation in the particular firm.

⁴Our paper also relates to research on industry wage differentials (see, e.g., [Katz *et al.*, 1999](#) for an overview, and [Krueger & Summers, 1988](#) and [Gibbons & Katz, 1992](#) for prominent examples). Our evidence highlights differences in the role of ILMs and in the potential to climb job ladders as a possible explanation to why similarly skilled individuals are paid differently across industries.

are well captured using more aggregated occupational categories.

The paper proceeds as follows. Section 2 describes our data sources and the institutional background. Section 3 introduces our methodological approach to identifying internal labor markets and job ladders. Section 4 discusses the properties of estimated ILMs and job ladders in the sample of large Norwegian firms. Section 5 discusses the robustness of our estimation methods to alternative assumptions about measurement error in the data and alternative clustering approaches. Section 6 concludes.

2 Data and Institutional Background

This section describes the administrative matched employer-employee data that we use, explains our key variables, and provides institutional background on labor markets in Norway.

2.1 Data

We use administrative matched employer-employee data from Norway. The data can be linked by unique and anonymized identifiers for every labor force participant, firm, and establishment. The Norwegian employer-employee register includes virtually all employment contracts from 2006 to 2014, except for contracts with fewer than four hours of work per week or below 10,000NOK (roughly 1,100USD) per annum. The contracts are reported by the employer to the authorities at the end of the year.⁵ Each reported worker contract includes information on the exact dates of alterations to the terms of the contract, the corresponding wage, industry and occupational codes, geographic location of the workplace, and tenure. The data thus allow us to observe transitions within organizations that cross different establishments as well as transitions across occupations, e.g., a person working as a systems engineer in one plant who becomes an operations manager at another plant, and moves on to a central position at the headquarters. Since our data cover every change of contract - including a change in the occupational code - we have a very detailed and reliable measure of job changes. We proceed by constructing time series of earnings for each worker, and by tracking all cases where the worker switches occupation, establishment, or firm.

2.1.1 Sample Selection

Our empirical analysis focuses on private sector firms, which can consist of multiple establishments (or plants). Since our main interest is in the organization of jobs in firms that offer career possibilities to workers, we restrict the sample to larger firms with at least 30 employees at some point during the period from 2006-2014 and at least 10 internal movers over this sample period. We further restrict our sample to firms with at least 15 external hires over the period to reduce bias that arises from limited mobility in the estimation of the AKM wage model.⁶ This leaves us with a sample of 3,611 firms that employ about 167 workers on average in each year.

⁵We exclude data prior to 2006 as the occupational codes were incomplete in some firms. From 2006 to 2014 earnings are reported per spell, where every contract change is recorded. The process of income reporting has changed since January 2015 and is based on complete monthly payments. We exclude observations after 2014 to avoid mechanical changes in the wage structure. While we use data from 2006 to 2014 to estimate the organizational structure of firms, all our empirical analyses are based on the period from 2007 to 2014. The reason is that we need one additional year of data prior to the start to be able to say where workers come from if they switch their job or firm.

⁶See e.g. Andrews *et al.* (2008), Kline *et al.* (2020), and Lamadon *et al.* (2019) and the discussions therein.

The sample of workers includes every full-time employed male and female aged 20 to 61. This restriction is customary in the literature and avoids issues related to work hours and labor force participation. We organize worker observations in an annual panel focusing on one job per worker and year. If a worker has multiple employers over the year, we choose the employer in the last month observed, and if the person has multiple employers in the last month, we select the employer with the highest total earnings in the particular year.

2.1.2 Variables

From our data source, we extract all transitions of workers between occupational codes, both within and across firms. We exclude transitions of workers who have a period of more than six months of non-employment between two positions or firms. We keep track of worker earnings and other characteristics. In the following, we describe our main variables.

Occupations. The key to our approach is fine-grained occupational codes that can be used to describe the variety of jobs or positions at a given firm in detail. We use 7-digit occupational codes based on the international standard classification of occupations by the International Labor Organization (ISCO). There are about 6,000 different occupations in the Norwegian version, where some job descriptions have been adjusted to meet Norwegian standards and occupational licensing rules. Given the detailed and fine classification of occupations, we use the terms occupation, job (title), or position interchangeably. For some jobs, the occupational descriptions include information about the rank of the occupation in the hierarchy, e.g. assistants, mid-level managers, top-level management, or members of the executive board. For the majority of jobs, however, it is not possible to classify the occupational code into a hierarchy structure based on the occupational description. In Section 3, we therefore introduce a data-driven method to classify occupations based on the flow of workers across 7-digit occupations within firms. We will compare results based on 7-digit occupations with occupational task definitions at higher levels of aggregation in Section 5.

Wages. Our measure of wage is the natural logarithm of the average monthly earnings of a worker in a firm. In our empirical application, we primarily use log wages. We also use residualized log wages when estimating individual and firm fixed effects. We residualize wages by regressing the log wage on a flexible specification of calendar year indicators that capture common year effects and include individual characteristics, such as dummies for each year of schooling, whether a person is married, a dummy for each number of children below age 18, gender, and each age category.

Worker and Firm Heterogeneity. Following the seminal work of Abowd *et al.* (1999, henceforth AKM), we decompose log wages into additive fixed effects that represent unobserved worker and employer heterogeneity. Let w_{ijt} denote the log residualized wage of individual i in year t and in firm j . The AKM wage model with firm fixed effects is described by

$$w_{it} = \alpha_i + \phi_{j(i,t)} + r_{it}, \quad (1)$$

where α_i is a time-invariant person effect for worker i , and $\phi_{j(i,t)}$ is the permanent firm fixed effect of firm j that employs i in year t . The time-varying residual component r_{it} is assumed to be uncorrelated

with worker mobility. We estimate this model on our sample of 3,611 larger firms. Due to the mobility restrictions, all firms in this sample are part of the major component of connections and concerns about limited mobility bias should be mitigated. In our analysis we use the AKM estimates as additional characteristics that capture unobserved heterogeneity. We interpret worker effects as a proxy for unobserved ability and firm effects as a proxy for the firm wage premium that the firm pays to all its workers.

Demographics. To capture complete information on workers’ geographic location and other socio-economic characteristics we link the matched employer-employee data with longitudinal administrative registers provided by Statistics Norway. These administrative data sources cover every Norwegian resident from 1967 to 2014 and contain individual demographic information like gender, age, zip code, and education.

Table 1: Firm Characteristics

	Mean	st.dev	p25	median	p75
Average number of workers per year	167.2	488.2	40	68.6	138.1
Number of occupations within firm (all years)	48.3	44.2	23	36	59
Number of occupations within firm (average year)	27.7	25.7	13	20	33
Number of internal moves (all years)	103.4	790.9	16	28	62
Number of external hires (all years)	246.0	719.0	55	97	209

Notes: This table reports firm-level characteristics for the sample of 3,611 private sector firms. See Section 2.1.1 for a detailed definition of the sample.

2.1.3 Descriptive Statistics

Table 1 shows characteristics of our sample of firms. Due to our sample restrictions, firms are fairly large and the average firm employs 167 workers per year. The firm size distribution is right skewed with a few very large firms in the tail of the distribution while the firm at the median only employs about 69 worker per year. Workers are employed in 48 different occupations in the average firm over the full observation period, and in 36 occupations in the median firm, but not all occupations are filled in every year. In a single year, the average firm employs workers in about 28 occupations (the median is 20). Over the 8-year period, we observe a total of 103 internal moves and 246 external hires in the average firm, but only 28 internal moves and 97 external hires in the median firm.

Table 2 provides some descriptive statistics on individual worker transitions in our data. We split the sample into stayers who remain in the same firm and occupation from one year to the next, and into movers. Among movers, we distinguish between those who switch between firms and internal movers who switch occupations within a firm. Not surprisingly, the majority of workers are not moving from one year to the other. While the literature in labor economics mainly focuses on external job transitions, our data show that internal transitions are important as well: among all moves, roughly 30% occur within firms. Moreover, internal movers experience a higher average wage growth than those who do not move and those who stay. Internal transitions are also associated with a lower share of negative wage changes than external transitions. Internal movers are positively selected as indicated by a higher average education level.

Consistent with the previous literature, we observe that workers who move across firms are younger (see e.g., [Neal, 1999](#)). These facts motivate the remainder of the paper, where we attempt to unpack the internal labor markets of firms.

Table 2: Descriptive Statistics: Movers vs. Stayers

	Movers								
	A. Stayers			B. Internal			C. External		
	mean	st.dev	median	mean	st.dev	median	mean	st.dev	median
In wage	10.6	0.47	10.6	10.7	0.49	10.6	10.3	0.55	10.3
Wage growth	0.074	0.24	0.054	0.095	0.27	0.067	0.068	0.47	0.043
Share positive growth	0.75			0.75			0.40		
Share negative growth	0.24			0.24			0.31		
Tenure in firm (months)	86.0	85.5	54	78.2	83.7	47	14.2	33.7	6
Age	41.5	10.9	42	40.0	11.1	40	34.4	10.5	32
Female	0.29			0.30			0.35		
Married	0.56			0.54			0.38		
6-9 years education	0.047			0.040			0.048		
10-13 years education	0.62			0.58			0.57		
14-16 years education	0.33			0.38			0.38		
AKM Person fixed effect	10.3	0.32	10.3	10.3	0.32	10.3	10.2	0.37	10.2
Worker-years		3,249,331			373,424			888,254	

Notes: This table documents the characteristics of movers and stayers using 7-digit occupations. The sample is described in Section 2.1.1. Observations are worker-year. Wage is at nominal levels.

2.2 The Norwegian Labor Market

The Norwegian labor market is characterized by a combination of institutional regulation and flexibility. Hiring and firing practices follow European labor law. Firms can hire employees on either fixed-term or permanent contracts, where a permanent contract typically entails a probationary trial period of six months, during which the employee can be dismissed on the grounds of the employee's lack of suitability for the work or lack of proficiency or reliability following a 14 day notice. Fixed-term hiring has stricter regulations, and an employee can only be temporarily hired if the work is also temporary, or if the employee is a temporary replacement hire, a trainee, or a participant in an active labor market program.

Union membership in Norway is relatively high compared to other countries in the OECD and the U.S., but has fallen from 58 to 53 percent from 1992 to 2013.⁷ Still, virtually all private sector jobs are covered by collective bargaining agreements, and wages and working hours are typically set in accordance with collective agreements between unions and employer associations.⁸ Tariff wages at the industry-level are first set centrally, after which wages are supplemented by local adjustments, or wage drift, which is bargained over at the firm level. The two-tier framework is considered a key reason for the highly compressed wage structure in Norway, with comparably low inter-industry wage differentials (see [Barth et al., 2014](#)).

⁷OECD Statistics [Trade Union Statistics](#), Accessed June 14th, 2020

⁸In some occupations, such as apprenticeships and for teachers, promotion is based on tenure in an occupation. These thresholds may be negotiated between unions and employer organizations.

3 A Worker Flow Approach to Internal Labor Markets

This section describes our two-step method for identifying internal labor markets and job ladders.

3.1 Internal Labor Markets

The first step of our method uses flows of workers between occupations to identify the boundaries of internal labor markets. Consider an intuitive example firm in the manufacturing sector, where the majority of production workers are part of a large ILM, including production line management positions and top executives. Within the ILM, workers switch positions, e.g., through promotions or job rotation. In a separate ILM, the firm employs workers for logistical purposes, like security personnel or receptionists. There are, however, no workers flows from one ILM to the other one – the two ILMs are entirely unconnected. This basic idea is captured in the concept of connected components from graph theory.

An ILM comprises occupations within the firm that are connected by some path established through realized worker flows. Two occupations are connected if a worker transitions between those two occupations, and a connected component is a set of occupations linked by at least one worker moving between them.⁹ We interpret these components as ILMs where hiring in one component is independent of hiring in the other components in the firm. This step thereby separates one ILM from another or from single occupations that exclusively hire workers from the external labor market.¹⁰

The first step of our method is closely related to the approach used to examine the internal structure of a particular firm in BGH’s seminal contribution. In contrast to that paper, our method is entirely data-driven and can be applied to examine the full workforce of a firm (not only management positions) and the full population of firms. Our method thus deviates from ILMs that are defined by observable categories such as business units (e.g., [Cestone et al., 2019](#)), establishments (e.g., [Giroud & Mueller, 2019](#)), or broad occupational categories ([Lazear & Oyer, 2004](#)).¹¹ In fact, we observe a large fraction of job transitions across detailed occupational categories and establishments, indicating that aggregating jobs into broader (e.g., 1-digit) occupational categories might not capture all relevant ILMs. We investigate this issue further in Section 5.1.¹²

Handling Measurement Error. A potential concern is that the data might contain measurement error in occupational coding. Although firms are obliged to report their employees’ occupation, they have no strategic incentives to do so correctly. Since we are mainly interested in transitions between occupations *within* firms, the issue appears less problematic as we can reasonably assume that the misclassification of occupations is constant within firms. Occasionally, however, due to turnover in HR staff or due to errors in data entry, we may observe implausible transitions between seemingly unrelated occupational codes.

⁹A potential limitation of our study is that the detailed occupational data are available only for eight years. To the extent that transitions between different jobs in the internal market are less frequent than eight years, we cannot capture connections between occupations with very low turnover.

¹⁰From the observed matrix of within-firm transitions, connected components can be easily identified using a simple breadth-first search algorithm implemented in standard statistical packages.

¹¹Our approach nests these other structures. If the mobility of workers is in fact limited to an establishment or a broad occupational category, that will be endogenously determined in our model.

¹²In the Norwegian data, 70 percent of all within-firm occupational switches are also transitions across broader 4-digit occupations and 49 percent correspond to occupational switches across 1-digit occupations. This is in line with recent evidence from the US ([Schubert et al., 2020](#)) where 86 percent of all 6-digit occupation switches (across firms) are also 2-digit switches. Moreover, 13 percent of all within-firm job switches in our data involve moving to a new establishment.

To address this concern, we employ a data-driven cleaning procedure that aims to separate plausible from implausible links. Our pruning algorithm iteratively removes workers from the network and assesses whether the removal breaks the network apart. Removing a worker breaks the network apart if the link established by this worker accounts for less than 10 percent of all transitions into or out of the jobs between which the worker switches and, hence, is an exceptional move.

The pruning algorithm is based on the concept of bi-connected components in graph theory and akin to a procedure previously used in a related context by [Kline *et al.* \(2020\)](#) (see Appendix B for a detailed description of the algorithm). In contrast to the leave-one-out components employed in [Kline *et al.* \(2020\)](#), we extend the algorithm to allow for rare but true transitions. Intuitively, if the observed network contains one transition from a senior manager position to the CEO position, this transition should not be capped to split the network. If, however, the network contains many accountants and many financial analysts who become financial managers and only one secretary who switches to financial manager, the algorithm breaks the link between the secretary and financial manager position. We call this the leave-X-percent-out procedure and choose 10 percent as our threshold. This proportional method has the additional advantage that it is scale-invariant, i.e., it yields the same classification of internal labor markets in two firms with the same structure but a different number of workers and transitions.¹³

3.2 Job Ladders

In the second step, we identify job ladders within internal labor markets by tracing the hierarchical structure of job flows. We rank all occupations within an ILM using relative flows between them. The basic intuition is simple: more observed transitions from occupation b to occupation a than vice-versa imply that occupation a is ranked higher than b . Using this logic, we overcome the fundamental challenge of not observing job hierarchies in linked employer-employee data.¹⁴

The idea of using worker flows to elicit job hierarchies is not new in personnel economics. However, the tedious process of collecting data and evaluating the flows between job titles has limited the existing evidence to case studies from particular firms. To extend the hand-curated approach of BGH beyond one firm, we introduce a minimum violations ranking approach. This algorithm orders occupations in an internal labor market such that there are as few transitions downwards (i.e., reverse) in the hierarchy as possible.

3.2.1 Minimum Violation Ranking

Specifically, the ranking algorithm orders occupations within each ILM based on relative flows between them. While the idea that a ranks higher than b if more workers move from b to a than in the opposite direction is simple, the actual estimation is more complicated because the observed relations between occupations are not transitive. We therefore use an algorithm that ranks occupations within each internal labor market such that the number of transitions towards lower steps on the occupational job ladder is

¹³We thank our discussant Thomas Lemieux to point out this idea.

¹⁴Several data sets and standard occupational codes include rough classifications of positions into levels. These rough levels can typically explain a relatively large portion of wage variation (see e.g., [Caliendo *et al.* \(2015a\)](#); [Bayer & Kuhn \(2019\)](#); [Lazear & Oyer \(2004\)](#)). Our algorithm, however, has several advantages over such predefined categories. First, it is firm-specific and therefore able to capture even subtle forms of firm heterogeneity. Second, it is much more detailed and allows us to examine promotions, demotions, and the underlying incentives on a much finer scale. Finally, if all job transitions follow such predefined categories, our algorithm will trace out exactly those categories as the relevant hierarchical structure.


minimized (**minimum violation ranking**). The ranking algorithm is based on the fraction of links that are upward moves along the job ladder, i.e. the ranking of the target occupation is higher than that of the source. We maximize this fraction over all possible rankings following the procedure in [Clauset *et al.* \(2015\)](#). Starting from an initial ranking where occupations are ranked according to the number of outbound transitions, we converge to the optimal ranking by repeatedly swapping ranks of a randomly chosen pair of occupations and accepting swaps where the associated new ranking has the same or a higher number of upward moves (or, equivalently, a lower number of ranking “violations”). Since there are potentially several equally plausible rankings with the maximum possible fraction of upward moves, we sample optimal rankings from the set of permutations with the maximal fraction of upward moves. Our results are then averaged rankings over the sampled sets while uncertainty around the estimated ranking can be measured by the distribution of ranks across optimal rankings. The resulting ranking represents ILM-specific hierarchies of jobs.

Figure 1 illustrates the logic of our algorithm. Suppose the ILM consists of three occupations—green, blue, and red—that are connected as shown in the top panel of the figure. Two possible rankings share the same minimum number of one violation (see the two middle panels). In both cases, the green triangle is the highest-ranked occupation. The red square and the blue circle, however, are each ranked second in one case and ranked lowest in the other case. In order to find a consensus ranking, the MCMC algorithm will therefore—after having converged to the minimum number of violations—collect many (random) samples from the set of possible rankings with the lowest number of ranking violations and subsequently average the ranks from all samples. The lower panel of Figure 1 shows that in the consensus ranking, the blue circle and the red square both receive rank 1.5 as we expect them both to be ranked last (i.e., rank 1) and second in half of the samples. Running the algorithm several times also provides us with a measure of uncertainty around the estimated rankings. Due to the stochastic nature of the algorithm, the consensus ranking will vary across runs. We compute the standard deviation of ranks across these runs as our measure of uncertainty.

3.2.2 Job Levels

After estimating the minimum violation ranking, we group occupations into hierarchy levels. To do so, we cluster occupations with similar ranks into the same level of hierarchy using a k-means clustering algorithm (see, e.g., [Bonhomme *et al.* \(2019\)](#)). In contrast to the clustering procedure in [Bonhomme *et al.* \(2019\)](#), we cluster occupations based on a single dimension: the estimated rank. This enables us to employ a stable dynamic programming algorithm that guarantees optimality and reproducibility ([Wang & Song, 2011](#)). Given a number of clusters K , the k-means algorithm assigns each occupation 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. In particular, following the suggestion in [Bonhomme *et al.* \(2019\)](#) we choose K such that the value of the k-means objective function is at least as low as the average standard deviation of the estimated ranks. Details are provided in Appendix B.

Figure 1: Two potential minimum violation rankings



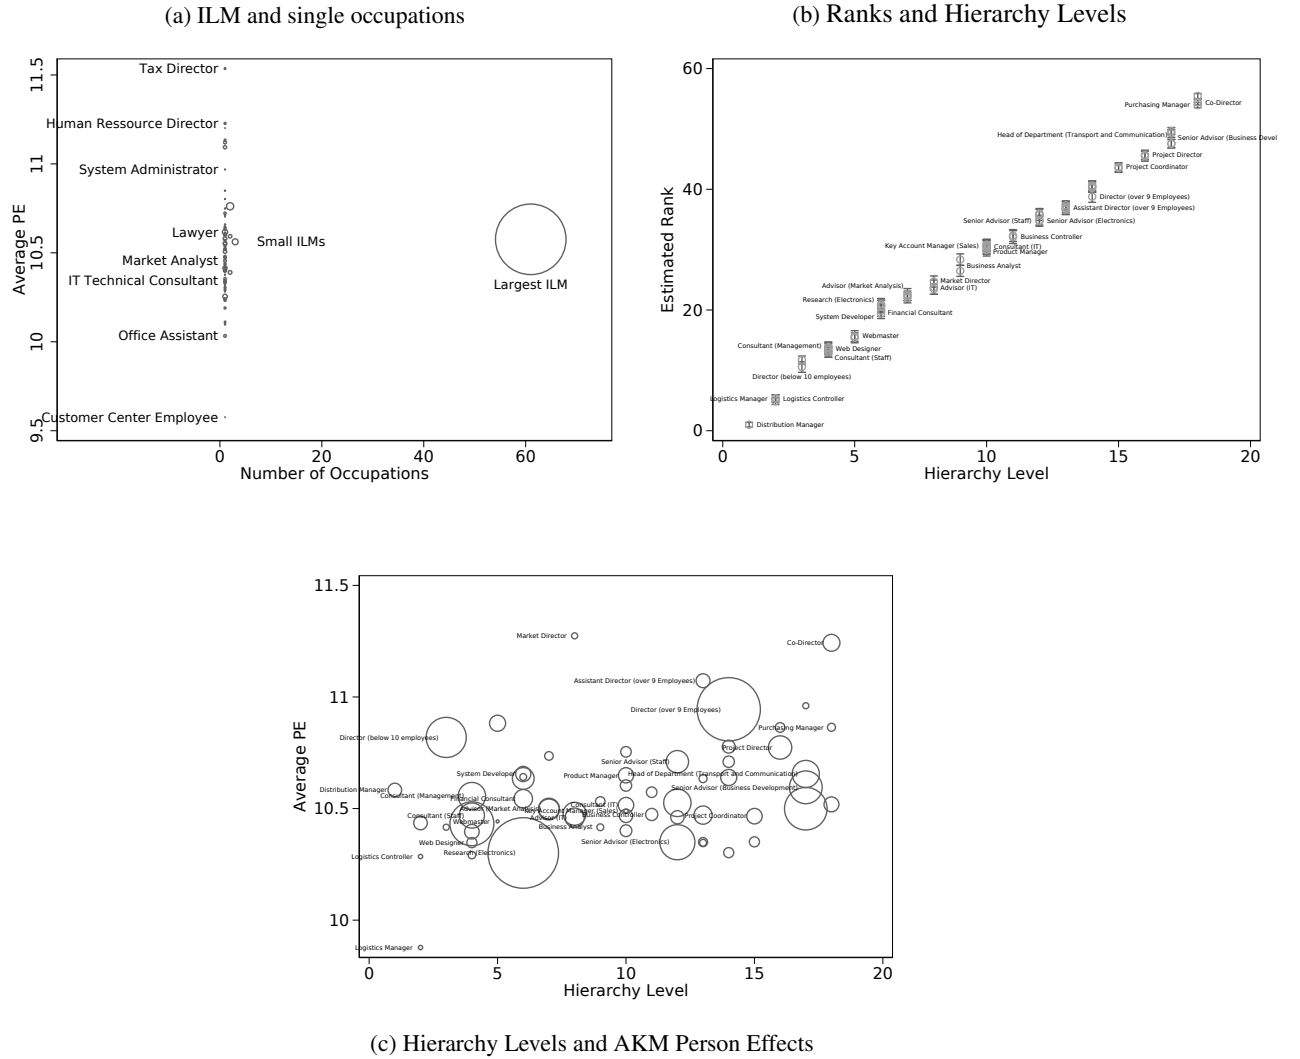
Notes: These figures illustrate the logic of our ranking algorithm. Suppose the component consists of three occupations—green, blue, and red—that are connected as shown in the top panel of the figure. The middle two figures illustrate two possible rankings that share the same minimum number of one violation. In both cases, the green triangle is the highest-ranked occupation. The red square and the blue circle, however, are each ranked second in one case and ranked lowest in the other case. The lower panel shows that in the consensus ranking the blue circle and the red square both receive rank 1.5 as we expect them both to be ranked last and second in half of the samples.

3.3 Illustrative Example

We illustrate our method based on an example firm to which we apply the algorithms described above. The firm is in the manufacturing industry, has four plants, and is relatively large with 4,229 worker-year observations and 153 occupations over the full sample period. After applying our pruning procedure, the firm has 83 single occupations that are not connected to any internal labor market.¹⁵ These single occupations, however, contain less than eight percent of all worker-year observations, whereas the largest ILM contains 73 percent of all observations. This largest ILM consists of 61 occupations that are connected by worker flows. Finally, there are three small ILMs that consist of two to three occupations.

¹⁵Either there does not exist any observed link to an internal labor market, or their connections are rare exceptions that are removed by our cleaning procedure.

Figure 2: Internal Labor Market and Job Ladders in a sample firm



Notes: Panel A illustrates the internal labor market structure of an example firm. Panel B shows the estimated ranks of the occupations in the ILM and how they are assigned to levels of hierarchy. Panel C shows the relation of the hierarchy levels and AKM person effects. The larger the circle, the more individual workers are employed in the ILM.

Internal Labor Market. Panel A of Figure 2 depicts the structure of the internal labor markets in the sample firm. The single occupations range from low-skilled occupations such as customer center employees and office assistants to high-skilled positions such as tax or human resource director. Workers in these occupations are hired externally, and when workers leave these positions they move to somewhere outside the sample firm. These are occupations with skills presumably tied to the occupation rather than to the core business of the firm. In terms of the individual wage component, measured by the average AKM person fixed effect in the occupation, there are single occupations both on a higher and lower level than the average in the ILM. The size of the circles in the figure is proportional to the number of workers and emphasizes that the single occupations and the three ILMs that contain two or three occupations are small compared to the largest ILM.

Ranks and Hierarchy Levels. Panel B of Figure 2 zooms into the largest ILM of the sample firm and shows how the estimated ranks are grouped into 18 hierarchy levels using our k-means clustering algorithm. The vertical bars indicate 95 percent confidence intervals for the estimated rank of a given oc-

cupation. We note that the clustering is based on the mean estimated rank only and the number of clusters is determined by the overall uncertainty in estimated ranks. Nevertheless, most levels of hierarchy are clearly statistically distinct from each other while positions within the same level have mostly overlapping confidence intervals. The classification of occupations into hierarchy levels appears to reasonably reflect the hierarchical structure of the job titles. We find several director and senior advisor positions towards the top of the hierarchy while the lower end of the hierarchy is populated by consultant, advisor, and research positions as well as managers of small units.¹⁶

Job Ladders. Panel C of Figure 2 shows how the job ladder structure of our sample firm relates to the person wage fixed effect from our AKM decomposition. The figure shows a positive relation between the hierarchy level of each occupation and the average AKM person effect in that occupation.

4 Main Findings

This section applies our method to our sample of firms and presents the main evidence on the structure of internal labor markets and job ladders.

4.1 Internal Labor Markets

Table 3 reports summary statistics on the organizational structure in the 3,611 Norwegian firms in our sample. In total, we identify about 180,000 components in the data, 97,993 or 54 percent of which include more than one occupation while 82,045 are singletons with only one occupation. Per firm, we identify on average 4.4 internal labor markets (ILMs) with multiple occupations and 22.7 singleton components.¹⁷

¹⁶Note that BGH restricts the top management positions to be at the highest levels. We could condition on the additional information from the occupational titles to restrict the output. Note also that not all high-powered incentives (i.e., stock options) are included in the wages.

¹⁷As we are interested in internal mobility, we refer to network components with two or more occupations as internal labor markets and to components with only one occupation as single occupations.

Table 3: Firms, Internal Labor Markets and Single Occupations

Panel A. Organizational Characteristics	Mean	st.dev	p25	median	p75
Number of components in firm	27.1	27.3	12	20	33
Number of ILMs	4.42	3.16	2	4	6
Number of singleton components in firm	22.7	25.9	9	16	28
Panel B. Largest ILM Characteristics	Mean	st.dev	p25	median	p75
Number of internal moves in largest ILM	83.4	302.3	11	21	52
Number of occupations in largest ILM	17.0	21.5	6	10	19
Share worker-years in largest ILM	0.69	0.24	0.55	0.75	0.89
Share occupations in largest ILM	0.36	0.19	0.21	0.34	0.50
Number of hierarchy levels	6.39	5.77	4	5	7
Panel C. Other components	Mean	st.dev	p25	median	p75
Share worker-years in singleton-components	0.18	0.15	0.067	0.14	0.24
Share workers-years in second largest component	0.09	0.11	0.02	0.050	0.12


Notes: This table reports firm-level characteristics for the sample of 3,611 private sector firms. The sample is described in Section 2.1.1. Singleton components are occupations not connected to other occupations within the firm.

Looking at the internal labor markets at the firm level, we see that most firms have one large ILM and several smaller ones. This corresponds to the evidence from the illustrative example firm in the previous section, where the largest ILM employs workers in the core business process, while smaller ILMs include occupations that are hired externally and not linked to other jobs in the firm. Focusing on the largest ILM in each firm, we see that out of the 103 internal moves in the average firm (see Table 1), 83 occur in the largest ILM, and 17 out of the 35 non-singleton occupations are employed in the largest ILM. In terms of worker-year observations 70 percent are employed in the largest ILM. The second largest ILM, in contrast, contains only around nine percent of all worker-year observations, while in the average firm, 18 percent of workers are observed in singleton components.

Figure 3 relates the number of internal labor markets and single occupations per firm to firm size and to the estimated AKM firm fixed effect. The graphs paint a consistent picture. The number of ILMs per firm increases with firm size (measured by the number of worker-year observations in Panel A) and with the firm pay premium (represented by the firm fixed effect from the AKM decomposition in Panel B). While a firm at the bottom of the size or pay premium distribution has roughly three ILMs, a firm at the top of the distribution has about five ILMs. The number of single occupations is also increasing with firm size and firm fixed effects. This suggests that larger and higher-paying firms increasingly rely on additional occupations that are hired predominantly on the external labor market and are not part of the ILM.

As the vast majority of workers and occupations are part of the core internal labor market, even in large firms, we concentrate on each firm's largest ILM in the remainder of this paper.

Figure 3: Components by Firm Size and Firm Fixed Effect



Notes: Panel A illustrates the number of ILMs per firm, and the number of single components (with only one occupation) per firm by firm size decile. Panel B illustrates the number of ILMs per firm, and the number of single components (with only one occupation) per firm by AKM firm fixed effect decile. The sample is described in Section 2.1.1.

4.2 Job Ladders

We now turn to the hierarchical structure of internal labor markets. As described above, we focus on the largest ILM of the firm, defined by the highest number of occupations, and restrict the sample to ILMs with at least two occupations. This corresponds to a set of 3,607 ILMs.¹⁸ Overall, we document substantial variation in the hierarchical structure of ILMs. The average ILM has 6.4 hierarchy levels as shown in Table 3. However, as illustrated by the histogram of the hierarchy level distribution in Figure 13b, the distribution is highly skewed. About half, or 56%, of our sample of firms has three, four, or five hierarchy levels in their largest ILM, while the longest job ladders have up to 66 hierarchy levels. In the following, we classify ILM's or firm types by the number of hierarchy levels (i.e., the number of rungs of the internal job ladder) in the largest ILM.

Appendix Table A1 shows summary statistics for a set of firm and workforce characteristics separately for different hierarchy levels. Notably, as firm size increases and the number of occupations expands, hierarchies get longer. The probability of promotion and demotion increase with the length of the hierarchy. The probability of being promoted in a single year is around eight percent for workers in firms with longer hierarchies. This means it takes on average about 12.5 years to be promoted to a higher level. Relative to promotions, demotions are observed less frequently; about half a percent of the workers experiences a demotion in a given year.¹⁹

4.3 Internal Careers

To get a better sense of what the hierarchies imply for careers in firms, we begin by describing some basic facts. We first count the employment in terms of worker-year observations, and plot the size of each hierarchy level for the different firm types in Figure 4a.²⁰ This picture confirms that larger firms

¹⁸In this section, we refer interchangeably to the ILM as the largest internal labor market or firm.

¹⁹The low frequency of promotions and the existence of demotions, suggest that the job ladders we identify do not correspond to automatic movements along a pay scale combined with an occupational upgrade, where all workers would typically move more regularly.

²⁰We group firms with longer ladders into one category with nine or more hierarchy levels, for expositional clarity. The patterns hold for longer job ladders too.

have longer hierarchies as the longer lines (i.e., firms with longer ladders) lie on a higher curve. We also see that the number of workers is decreasing in the hierarchy level for a given type of firms. This means that hierarchies have a pyramid shape – consistent with the finding in BGH . We consistently find this pyramid structure across all firm types. In terms of number of occupations, Figure 4b shows that the average number of jobs across levels is hump-shaped, again consistently across firm types. At the entry-level, there are fewer occupations and specialization tends to increase towards the middle of the hierarchy, before it declines at the top hierarchy levels.²¹

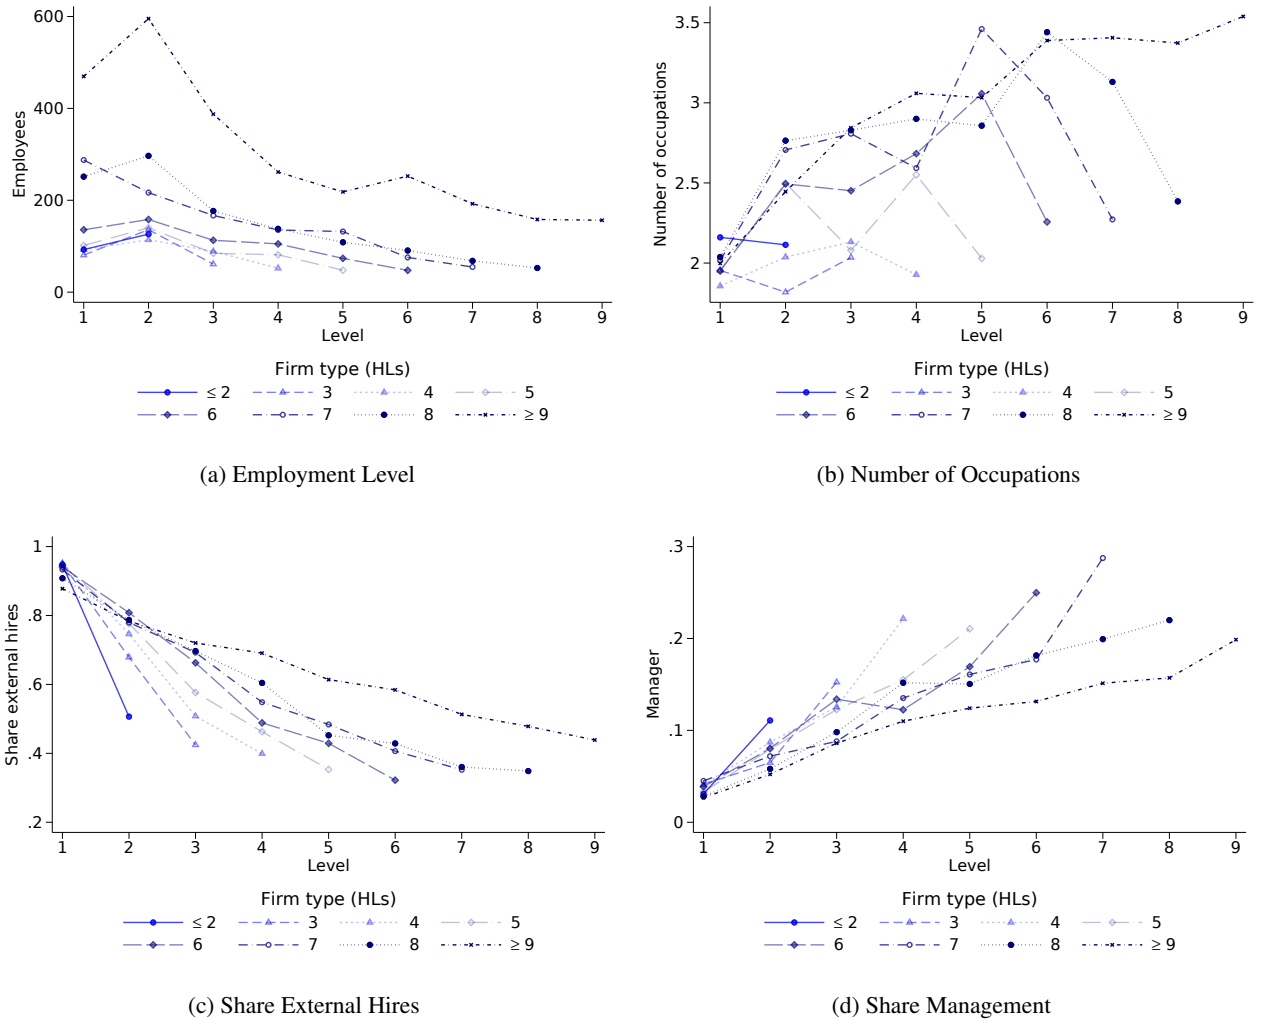
In Doeringer & Piore (1971), workers enter an internal labor market at entry positions and then climb up the job ladder. This strict form of an ILM is typically found in bureaucratic organizations (see, e.g., Bertrand *et al.*, 2020). For private sector firms, the empirical literature on job ladders generally documents less strict ladders with a declining share of external hires over hierarchy levels (e.g., Baker *et al.*, 1994 and Lazear & Oyer, 2004). Figure 4c confirms this finding for the Norwegian labor market. The share of external hires is above 90 percent in all firms at the lowest hierarchy level and it declines to around 40 percent at top hierarchy levels.²² Interestingly, the figure shows that the share of external hires approaches 40 percent at the top hierarchy level for all firm types, irrespective of the length of the job ladder.

Finally, Figure 4d shows that the share of workers in management positions increases with the hierarchy level. This occupational category includes mid-level and top-level management, for example, managers in a retail shop as well as the CEO and the executive board. Reassuringly, the management share is close to zero at the lowest hierarchy levels in all firm types. Moving further up the hierarchy, we see that the share increases monotonically. This pattern is in line with the intuition that task complexity and responsibilities increase as workers move up the job ladder. Interestingly, both the low shares of management at the lowest levels and the increasing pattern of management shares with level is similar across firm types. But there is an uptick in the management share at the top hierarchy levels across firm types.

²¹This pattern is also consistent with the organizational chart in BGH, who classified job titles into levels by relying on information about moves between job titles. BGH initially focused on fourteen titles that each represented at least 0.5 percent of employee-years. They then added the code for Chairman-CEO, and two other titles observed in moves from their original titles to Chairman-CEO to fill in the job ladder to the top of the organization. The lowest level was identified by the hiring patterns: It consisted of job titles exclusively filled externally, who then later moved into job titles at higher levels. The next levels were determined by manually minimizing the rank reversal: Most moves other than stays or exits from the lowest level went to six other titles. These six job titles were only trivially filled by workers coming from other titles where external hires were much less important. This procedure was continued until all job titles were assigned a level.

²²The remaining hires at the lowest hierarchy level are demotions or lateral moves within the hierarchy level.

Figure 4: Empirical Job Ladders



Notes: This figure illustrates the key characteristics of our empirical job ladders by firm types. HL in the figure refers to the number of hierarchy levels. We group ILMs with two or fewer occupations, and nine or more hierarchy levels into one category and plot the averages for all ILMs in this group. We consider the largest ILM per firm. The sample is described in Section 2.1.1. Number of employees are worker-year observations.

Individual Careers.

The patterns shown in Figure 4 provide a strong indication that the estimated hierarchy structures capture common properties of internal labor markets. According to [Doeringer & Piore \(1971\)](#), a key concept to ILMs is that workers have careers within their firm. We proceed by examining individual workers' career paths in firms with different numbers of hierarchy levels. To track the career within a firm, we construct a sample of workers who enter one of our firms in 2007 or 2008 at any level of the hierarchy. We then follow the promotion, demotion, and wage dynamics of these workers in the same firm over the following five years.


The panels of Figure 5 document the average career paths of the new entrants. Panel 5a examines how the duration of the employment spell varies with the firm type. We observe quite a lot of turnover, especially in the first two to three years, and on average only 20 percent of entrants remain in the same firm after five years. The figure also shows that exit probabilities vary by firm type with a clear ranking. Firms with longer job ladders potentially offer better prospects of promotions and are able to keep workers longer. By contrast, firms with shorter hierarchies experience the highest level of churn.

The next panels decompose the career dynamics of workers who are moving between positions in the same firm into three categories. Panel 5b shows a clear ranking in the cumulative number of promotions across firm types. While there are few promotions in firms with short job ladders, up to 35 percent of workers who remain in a firm with at least nine levels for five years experience a promotion. The graph also indicates the probability of being promoted over time is almost linear for each firm type. Panel 5c, shows the corresponding graph for demotions. Across all firm types, demotions are much rarer than promotions and happen less than one fifth as often to workers still in the firm after five years.²³ Finally, Panel 5d shows the cumulative number of lateral moves (i.e., moves within the same hierarchy level). Similar to demotions, lateral moves are relatively rare, and tend to happen slightly more often in larger firms with longer job ladders.


The evidence in Panels (a) to (d) indicates that there is heterogeneity in the speed of promotions across firm types. At the same time, career dynamics could differ substantially across workers which means that the patterns could be driven by dynamic selection. To shed some light on this question, we explore heterogeneity in the number of promotions across firm and worker types in Figure 6. For simplicity, we collapse the firm types into two groups with short job ladders (two to five hierarchy levels) and with long job ladders (six or more hierarchy levels). Additionally, we separate between workers with an AKM person effect below and above the median. The figure shows an interesting pattern: Conditional on the quantile of the AKM person effect, there are still more promotions in firms with longer ladders compared to those with short ladders. And, strikingly, workers with higher AKM person effects climb the job ladder substantially faster regardless of the firm type. We view this pattern as evidence for a sorting mechanism that assigns higher ability types to the upper part of the hierarchy. Such sorting is consistent with tournament models (e.g., Lazear & Rosen, 1981) and assignment models (e.g., Gibbons & Waldman, 1999a), and will be further explored in the next section.

²³This analysis offers a different statistic of the performance of our ranking algorithm. The relationship of promotions to demotions confirms the performance of our ranking method, which had the aim of minimizing downward career moves.

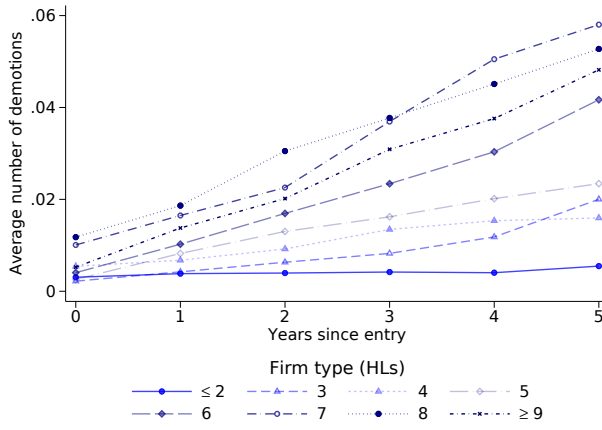
Figure 5: Careers




(a) Share of workers remaining in firm



(b) Number of internal promotions




(c) Number of internal demotions



(d) Number of lateral job changes

Notes: Figure 6a illustrates the attrition of workers since entering the firm. Figure 6b illustrates the total number of moves in the same firm. A move is when a worker switches job title. Figure 6c illustrates the total number of promotions in the same firm. A promotion is when a worker moves to a higher level in the firm. The sample is restricted to workers who join the firm from 2007 to 2009 (see further details of the sample restrictions in Table A1).

Figure 6: Person Effects and Number of internal promotions



Notes: Figure 6 illustrates the total number of promotions in the same firm by the level of person fixed effects. The baseline sample is described in Section 2.1.1, and is further restricted to workers who join the firm from 2007 to 2009 to have a balanced panel of five years.

4.4 Internal Labor Markets and Wages.

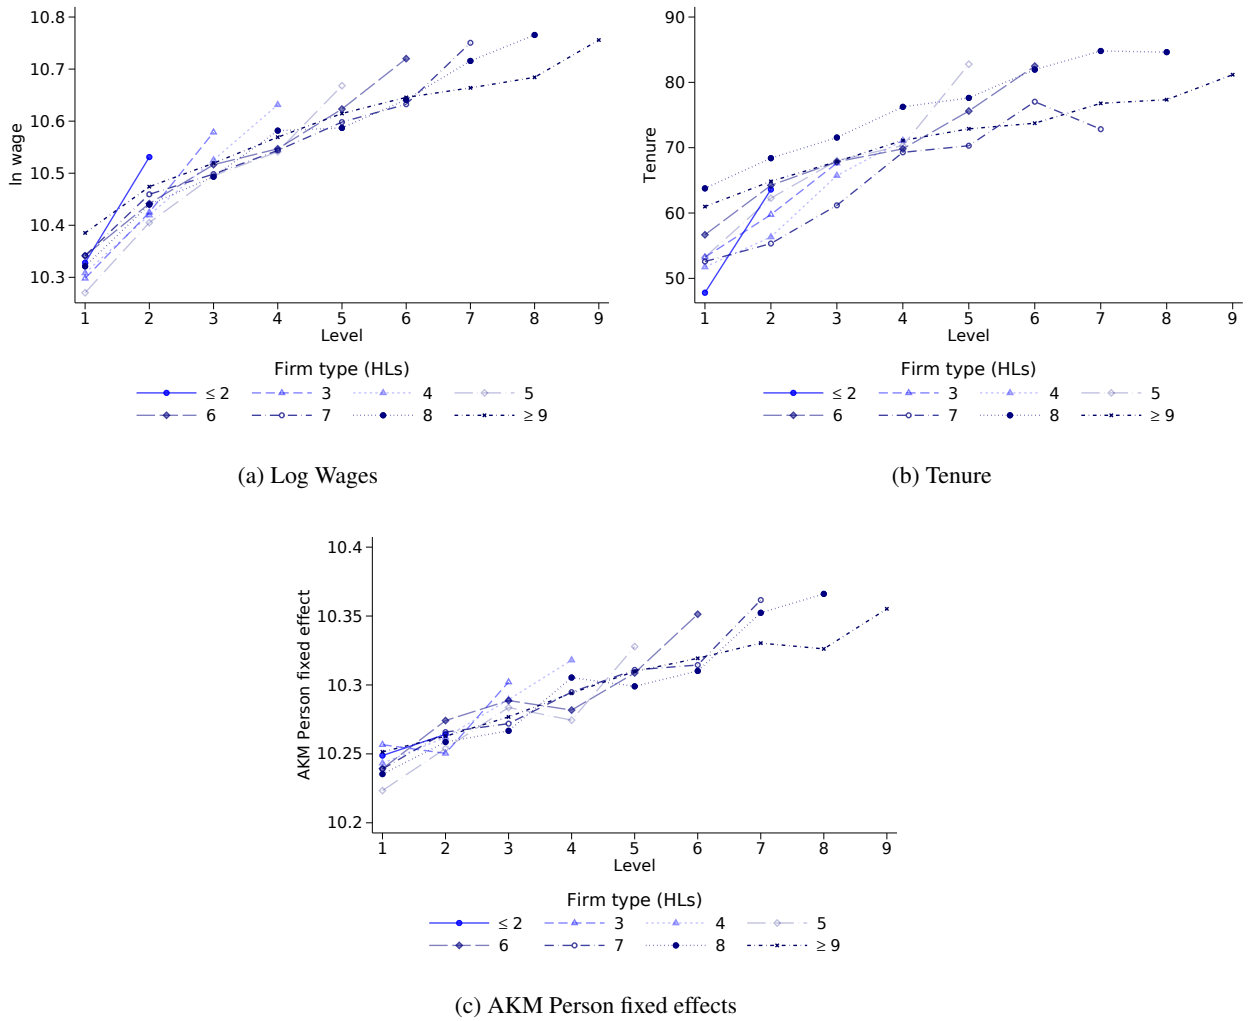
In models of internal labor markets, wage policies are tightly connected to the job hierarchy. Firms can use administrative rules to assign wages to job titles, or use the internal hierarchy structures to design promotion systems as incentive or allocation mechanisms. In the rank-order tournament model of Lazear & Rosen (1981), promotions are used to elicit effort from employees. A key prediction from this model is that the wage spreads between hierarchy levels increases with the level. This increasing premium reflects higher competition from other workers in the same and other rungs of the ladder, thereby compensating for the lower probability of winning the tournament.

We confirm that the estimated job ladders in Norwegian firms reflect the predictions from this theory. Figure 7a shows that average log wages are increasing by hierarchy levels. The upward pattern appears somewhat steeper for firms with shorter hierarchies. But importantly, the profile for each firm type is almost linear in mean log wages and there is an additional wage boost at the final rung of the career ladder for each firm type. Firms can also increase retention by offering more attractive careers and promotions (see, e.g., Gibbons & Waldman, 1999a). This incentive structure implies that higher ability types are more likely assigned to the upper part of the hierarchy, or they climb upwards faster (see the evidence in the previous section).

To assess these implications, we begin by investigating the association between firm tenure and levels of the job hierarchy. In Figure 7b, we plot average worker tenure over hierarchy levels for different firm types. Consistent with evidence on promotions and declining shares of external hires, the figure shows that tenure is increasing at higher levels of the hierarchy. The average employee at the lowest rung of 8-level firms has about 65 months of tenure, whereas the employees at the top rung have about 20 months longer tenure. Next, we use worker fixed effect from AKM wage decompositions as a proxy for ability. Figure 7c shows average person fixed effects from the AKM wage decomposition by hierarchy level. The fixed effects can be seen as measures for worker characteristics that are transferable across firms. Similar

to wages, person fixed effects follow parallel and upward sloping patterns in hierarchy levels across all firm types. In entry-level jobs, the average person effect is substantially lower than at higher hierarchy levels. This suggests that higher ability types move up the job ladder more quickly and this mechanism generates a sorting pattern.

Figure 7: Wage and Promotions




Notes: This figure illustrates the key characteristics for our empirical job ladders by firm types. HL in the figure refers to the number of hierarchy levels. We group ILMs with nine or more hierarchy levels into one category and plot the averages for all ILMs in this group. We consider the largest ILM per firm. The sample is described in Section 2.1.1.

Worker Careers and Wages.

To further assess the relation between the incentives posed by longer job ladders and the incentives posed by pay premiums associated with a firm, we return to the analysis of the careers of workers who entered their employer in 2007 or 2008. Figure 8a documents wage dynamics for those workers who entered in 2007 and 2008 and stayed at their employer for (at least) five years. Over the first five years at the firm, the data show monotonically rising wage profiles. This holds true for any firm type. Firms with longer ladders tend to pay higher starting wages on average while (log) wage profiles increase in parallel. Panel

8b collapses firm types into firms with short ladders (two to five hierarchy levels) and firms with long ladders (six or more hierarchy levels). Additionally, we split firms into terciles according to the AKM firm fixed effect. The data show that both the length of the job ladder and the AKM firm pay premium shape the wage setting process. Within a given tercile of the AKM firm effect, we still see that firms with longer job ladders pay higher wages. At the same time, there is a clear ranking that confirms that - given the length of the job ladder - firms with higher AKM firm effects pay higher wages. While we note that we have no (quasi-)experimental variation that allows us to distinguish the causal impact of job ladders on wage setting from unobserved firm heterogeneity, we view this evidence as suggestive that both mechanisms are operative.

Figure 8: Careers



Notes: Panel a displays wage profiles of workers who entered in 2007 or 2008 and stayed for at least five years separately by the length of the job ladder in the firm. Panel b additionally separates by terciles of the AKM firm fixed effect. The baseline sample is described in Section 2.1.1, and is further restricted to workers who join the firm from 2007 to 2009 to have a balanced panel of five years.

Individual Wage Regressions.

In BGH, hierarchy levels explain about 70 percent of the cross-sectional variance of wages within the firm (and substantially more than typical Mincer regressions). This finding is echoed in a recent paper by [Bayer & Kuhn \(2019\)](#) who use five rough levels of job hierarchies reported in German data. We examine whether this holds on average in our sample and explore the distribution across different firms.

Table 4: Effects of Human Capital and Hierarchical Level on Current Salary.

	(1) ln wage	(2) ln wage	(3) ln wage	(4) ln wage
Tenure		0.00183*** (0.0000138)	0.00198*** (0.0000138)	0.00195*** (0.0000136)
Tenure ²		-0.00000432*** (4.20e-08)	-0.00000453*** (4.21e-08)	-0.00000448*** (4.17e-08)
Education:				
10-13 years education			-0.0260*** (0.00157)	-0.0274*** (0.00152)
14-16 years education			0.0986*** (0.00166)	0.0861*** (0.00161)
Level:				
Level 2				0.0443*** (0.00107)
Level 3				0.0911*** (0.00122)
Level 4				0.117*** (0.00140)
Level 5				0.154*** (0.00158)
Level 6				0.165*** (0.00174)
Level 7				0.184*** (0.00197)
Level 8				0.188*** (0.00225)
...				⋮
Level 16+				0.242*** (0.00197)
Observations	4,253,859	4,253,859	4,253,859	4,253,859
R ²	0.537	0.545	0.554	0.565
Adjusted R ²	0.521	0.529	0.539	0.550
Gender and age FE	Yes	Yes	Yes	Yes
Firm x year FE	Yes	Yes	Yes	Yes

Standard errors in parentheses are robust and clustered at the individual level * p<0.05, ** p<0.01, *** p<0.001.


Notes: The baseline sample is described in section 2.1.1). Column (1) displays the output from a regression of log wages on age dummies, tenure and tenure squared. Column (2) displays the output from a regression of log wages on age dummies, tenure, tenure squared and education dummies (years of schooling). Reference level of years of schooling is 0-9. Column (3) displays the output from a regression of log wages on age dummies, tenure, tenure squared, education dummies and dummies for each level in the hierarchy. The reference level in the organizational structure is 1.

Table 4 reports regression results for a pooled regression of log wages on human capital measures and hierarchy level indicators. The first column reports the explained variance of the benchmark model with fully interacted firm and year fixed effects with an adjusted R^2 of 0.52. The second and third columns show that controlling for a quadratic polynomial in tenure and including dummies for years of schooling add about 1.8 percentage points to the explained variance. The fourth column shows the full specification with dummies for 16 hierarchy levels. Comparing the adjusted R²s reveals that the hierarchy structure explains almost as much of the variance in log wages as the tenure and human capital variables together.

With an adjusted R2 in column 4 of 0.55, there still appears to be a substantial individual component to wage setting.

The coefficients on the hierarchy level dummies, also plotted in Figure 9, confirm that promotions give an extra push to wage growth beyond the wage increases due to seniority. Within the same firm and holding individual characteristics fixed, the average wage in the top hierarchy level is about 20 percentage points higher than the wage at entry level. The increase in the (log) promotion bonus is almost linear at least from level 4 onwards. This again confirms the hypothesis of increasing spreads between hierarchy level from tournament theory. It also implies that the gap in wage growth between workers in firms with longer hierarchies and to workers in firms with shorter hierarchies, leads to increasing income inequality over the life cycle.²⁴

Figure 9: Mincer Regression.



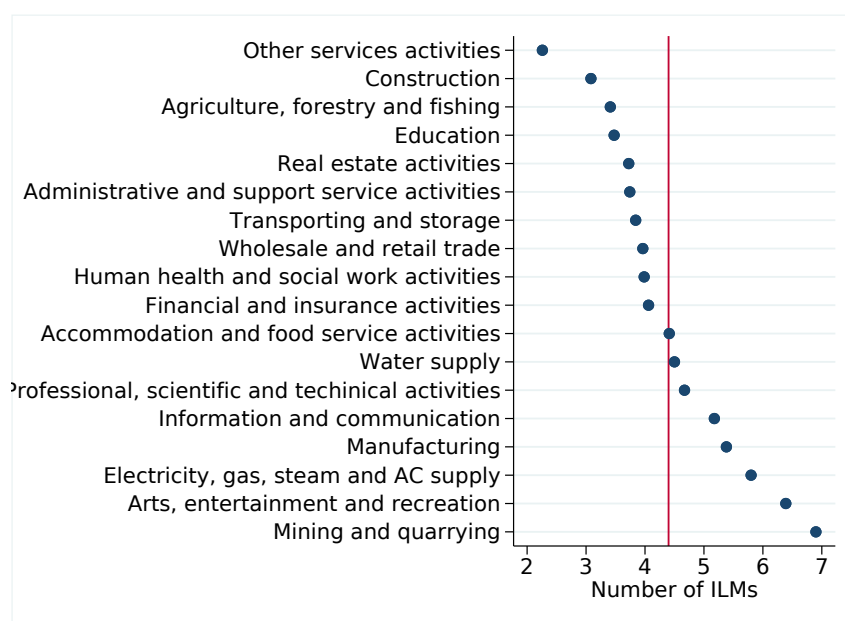
Notes: This figure presents the regression coefficients from Table 4, Column 4. Regressions include fully interacted firm and year dummies and dummies for gender and each age. The baseline sample is described in Section 2.1.1. Estimate for level 16 includes levels 16-66. Vertical bars represent 95 percent confidence intervals. Standard errors are clustered at the individual level.

4.5 Heterogeneity by Industry

The organizational structure in firms is likely to reflect competition between workers for promotions (e.g., Lazear & Rosen, 1981), differences in shirking or monitoring costs (e.g., Shapiro & Stiglitz, 1984), and production processes across industries (e.g., Lazear, 1989). To explore the heterogeneity in the organizational structure across industries, we proceed in three steps. We begin by documenting differences in the number of internal labor markets by industry in Figure 10. The vertical axis displays two digit industries and the horizontal axis shows the number of ILMs in each industry. The vertical line corresponds to the full sample average and the dots represent the average number of ILMs for each industry, which ranges from two to seven ILMs. We find the largest number of ILMs in manufacturing, electricity and gas supply, arts and entertainment, and in the mining and quarrying industry, which is dominated by large and high-paying firms. The lowest numbers of ILMs are found in firms in other service activities, construction, agriculture, and education.

²⁴We have also estimated a specification with an additional dummy variable indicating the top level of the firm's hierarchy structure, which captures the extra wage boost at the top hierarchy levels for each firm type in Figure 8a. In this specification the log wage premium for the top hierarchy level is 0.056 (SE 0.002).

Figure 10: Average Number of Internal Labor Markets Across Industries

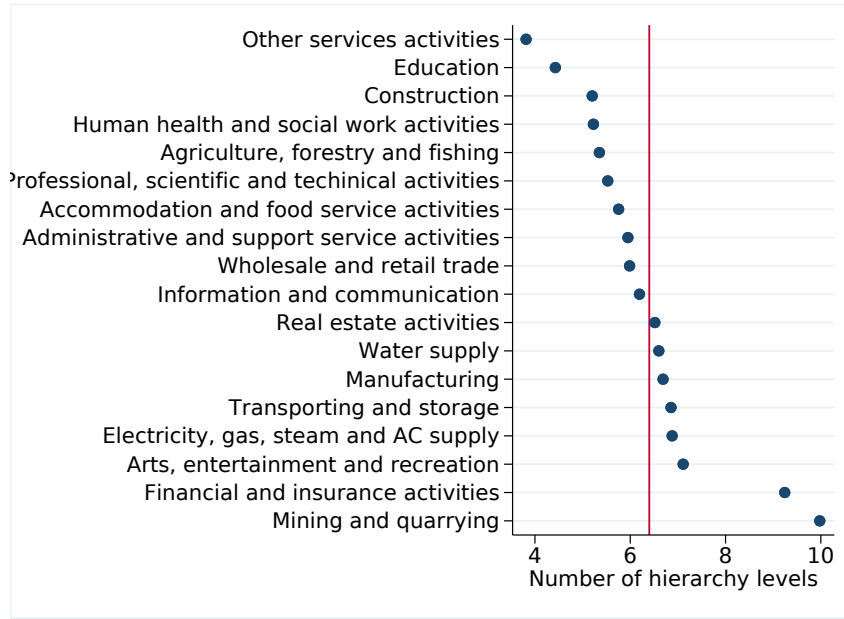


Notes: The figures show the average number of internal labor markets by industries for our sample of firms (see further details of the sample restrictions in Section 2.1.1).

In the second step, we measure the degree of segregation within firms by computing the share of workers in the largest ILM. An equal distribution of workers across components implies a high degree of segregation, whereas a concentration of workers in one ILM implies that the majority of occupations are connected by worker transitions. Appendix Figure A1 shows significant variation in the concentration of workers in the largest ILM. We find industries with more segregated organizational structures, such as electricity and gas supply, arts and entertainment, accommodation and food services, and information and communication, have fewer than 60% of all worker-year observations in the largest ILM. In these industries, the employment share in single occupations is 20%, roughly equal to the share in smaller ILMs. At the other end of the distribution, we find construction, real estate, finance and insurance, and other services have more integrated structures. In these industries, more than three quarters of all workers are employed in the firm's largest ILM.

Turning to the hierarchy structure, we focus on job ladders in the largest ILM of each firm. Figure 11 shows that the average number of hierarchy levels also varies substantially across industries. Firms with long job ladders are more prevalent in mining and quarrying, with about 10 levels on average, and financial and insurance services, with about nine levels. On the other end of the scale, we find service industries, such as education, health services, and the construction industry, where firms have on average four to five hierarchy levels.

Figure 11: Average Hierarchy Levels Across Industries



Notes: The figures show the average number of hierarchy levels by industries for our sample of firms (see further details of the sample restrictions in Section 2.1.1).

Overall, this analysis shows that organizational structures are industry-specific and do not only vary by firm size or pay levels. It allows us to identify industries with a more integrated organizational structure, where the majority of workers have access to the ILM, and where job ladders are long. This is the case for firms in financial and insurance activities, which tend to offer better internal career opportunities to their workers on average than, for example, firms in the construction industry.

5 Robustness

This section examines the robustness of our method to different levels of occupational aggregation, different assumptions about the noise in the data, and different clustering methods. We assess whether the inferred hierarchy structure is stable over time.

5.1 Sensitivity to Occupational Aggregation Levels

Our baseline specification uses 7-digit occupational codes, which is the most disaggregated categorization of job titles. In many alternative data sources, however, occupations are only available at higher levels of aggregation. This can limit the ability to capture some mobility between occupations. To examine the gain from a higher level of disaggregation in the occupational definition and the applicability of our approach to other occupational categories, we re-estimate internal labor markets and job ladders in our sample of firms based on 4-digit and 1-digit occupational codes.

Whereas the 7-digit occupational codes include 6,092 different occupations in our sample of firms, there are only nine different 1-digit occupations.²⁵ Figure 12 shows firm level averages of the number of occupations employed and the number of internal moves observed, separately for different levels of

²⁵The 10th occupational category covers military occupations, which is excluded from our sample of private sector firms.

aggregation from one to seven digits of the occupational code. The average number of occupations is increasing in the level of disaggregation: the average firm employs workers in six different 1-digit occupations but in 48 different 7-digit occupations. Notably, we see that even the most aggregated occupational codes capture substantial internal mobility within firms – falling from 103 for 7-digit occupations to 51 per firm in the case of 1-digit occupations. Hence, it appears that even with a high level of aggregation, we can still capture some important mobility of workers within the firm.²⁶

Appendix Table A2 summarizes the main characteristics of internal labor markets and job ladders comparing the estimates based on 7-digit, 4-digit, and 1-digit occupational codes, and Appendix Figure A2 plots the histograms of the hierarchy levels using 4- and 1-digit versions. In terms of ILMs, we see that the number of ILMs per firm falls from 4.4 on average in the 7-digit case to 2.2 in the 4-digit and to one in the 1-digit specifications.²⁷ The largest decline is in singleton occupations which drops from 23 to 10 to only one across the three different specifications. When focussing on the largest ILM per firm, we see that more aggregated occupational codes still capture some elements of the job hierarchy, where the number of hierarchy levels at the median firm drops from five to four and three in the case of the most aggregated occupational code. In terms of promotion dynamics, we see the promotion share of all internal moves is 81 percent in the most disaggregated version, and drops to 77 and 73 percent in the two other specifications. Looking at wage increases associated with the levels, Appendix Figure A3 shows a positive relationship up to the sixth level, after which the two more aggregated versions exhibit a somewhat less stable relationship than the baseline specification.

These comparisons show that ultimately, the optimal aggregation level of occupational codes depends on the research question. If one is interested in the change of the organizational structure in specific firms such as outsourcing of occupations or the relationship between different internal labor markets, a detailed level of disaggregation will be essential to capture the difference between linked and unlinked occupations or internal markets. On the other hand, if one aims to compare hierarchy structures across firms, a more unified occupational definition based on broader classification might be more informative (Kauhanen & Nappari, 2012).


5.2 Sensitivity to Eliminating Potentially Misclassified Occupational Transitions

In our main analysis, we account for potential measurement error in occupational coding by using a pruning algorithm that removes misclassified occupational links. We believe that this is a conceptually important step that helps to distinguish workers that are part of the ILM from those in single occupations that are hired only externally. In this section, however, we show that our findings are robust to omitting the cleaning step and using the unadjusted connected components where we have not removed any links from the networks. Table 5 reports summary statistics for the unadjusted components. The vast majority of firms are still characterized by a single large ILM (the share of worker-year observations in the largest ILM slightly increases from 69 to 71 percent). In comparison to our baseline in Table 1, there are significantly fewer single occupations in firms.

²⁶The reason for this seems to be that firms mainly apply highly disaggregated codes internally, while higher aggregation levels are more comparable across firms. The average 7-digit occupation is used in about 29 firms and only in five firms for the median, while the average 1-digit occupation is used in the majority of firms.

²⁷The share of worker-year observations corresponding to an internal move is nine percent in the 7-digit specification; it drops to six percent and four percent for the higher levels of aggregation, respectively. The share of internal hires observed in each specification declines by half overall, from 32 percent in the 7-digit case to 23 and 16 percent in the other two specifications.

Figure 12: Occupations and mobility by digits of occupational coding



Notes: This figure illustrates the number of internal moves and number of unique occupations by digits in the occupational code.

Table 5: Internal Labor Markets and Single Occupations

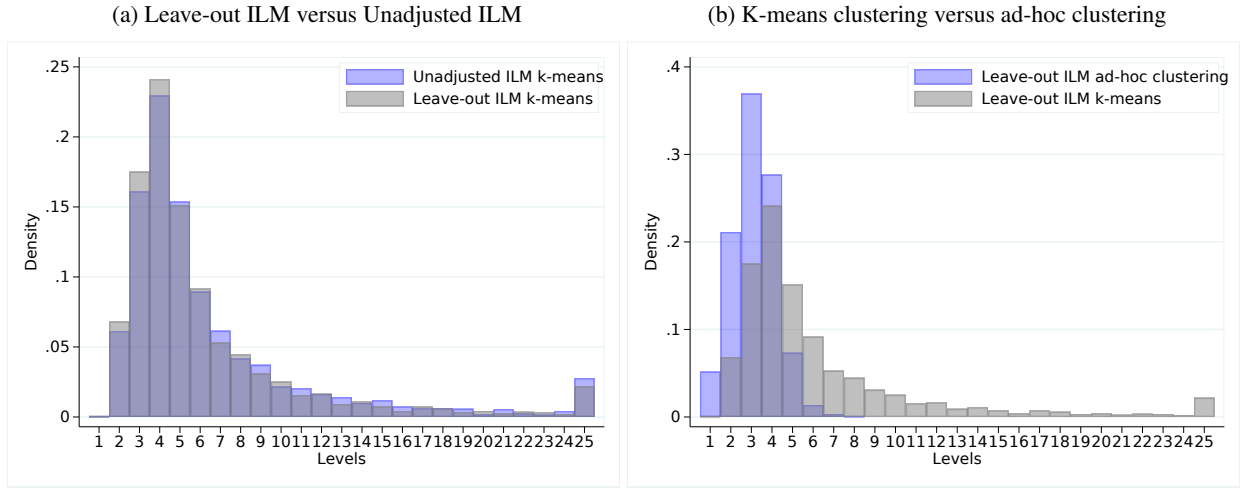
	Mean	sd	p25	median	p75
Number of components in firm	25.3	25.6	11	19	31
Number of ILMs	4.19	2.99	2	4	6
Number of singleton components in firm	21.1	24.4	8	15	26
Number of occupations in largest ILM	19.3	26.4	7	11	21
Number of internal moves in largest ILM	84.7	303.3	11	22	54
Share worker-years in largest ILM	0.71	0.24	0.56	0.78	0.90
Share occupations in largest component	0.40	0.21	0.23	0.38	0.54
Share worker-years in singleton-component (7)	0.17	0.15	0.060	0.13	0.24

Notes: This figure documents the characteristics of the internal labor markets when using the unadjusted ILMs. The sample includes 3,611 private sector firms (see further details of the sample restrictions in Section 2.1.1).

We now turn to the question of whether including these single occupations in the main ILM affects the structure of hierarchy and our empirical results. We re-estimate the hierarchy structure using the unadjusted ILMs that ignore potential misclassification of occupational links. The histogram in Figure 13a shows the worker-year weighted distribution of the number of hierarchy levels in the largest ILM of each firm. Differences that arise from omitting the cleaning step appear negligible.

In Appendix Table A3, we compare our baseline specification (Panel A) to the unadjusted version (Panel B) by regressing number of employees, number of occupations, the share of external hires, an indicator for management position, log wages and AKM fixed effects on a linear and a squared term of the hierarchy level. This exercise shows that our main empirical findings are robust to omitting the cleaning step. Importantly, the relationship between wages and AKM person fixed effects and hierarchy levels remains virtually unchanged.

Figure 13: The Distribution of Job Ladders



Notes: This figure shows the worker-year weighted histogram of the number of hierarchy levels per firm. In Figure 13a, grey bars represent our baseline method where we apply the pruning algorithm, and the blue bars represent the unadjusted components. In Figure 13b, grey bars represent our baseline method, and the blue bars represent the ad-hoc clustering method. The sample is described in Section 2.1.1. Levels are right censored at 25.

5.3 Other Clustering Methods

In our baseline specification, we group occupations to levels of hierarchy using a k-means algorithm based on the estimated hierarchy score. A potential drawback of this approach is that it takes into account the overall level of uncertainty in the estimation of ranks to determine the number of hierarchy levels, but does not take into account the uncertainty around each individual rank. In this section, we show that our results are robust to an alternative clustering method that is based on the pairwise statistical difference between adjacent ranks. In particular, occupations that have statistically indistinguishable estimated ranks are subsumed into the same level of the hierarchy. To do so, we order occupations according to their rank and compute rank differences between adjacent pairs. Starting from the highest rank, occupations are grouped together if the rank difference is lower than the estimated standard deviation of the higher ranked occupation. We iteratively apply the grouping procedure until all adjacent occupation pairs are evaluated. Figure 13b shows that this ad-hoc clustering method leads to a coarser classification of hierarchy levels compared to our baseline. The maximum number of hierarchy levels is seven instead of 66. Panel C of Appendix Table A3 shows that the regression results remain similar, and the qualitative conclusion remains unchanged.²⁸

5.4 Stability over Time

Finally, we turn to the stability of the organizational structure over time. Throughout the paper, we have treated the number of occupations and hierarchy levels within a firm as fixed characteristics. However, while the job ladder is identified using observed worker transitions between occupations in any two years of the observation period, this does not necessarily imply that the firm employs workers in the same set of occupations in every year. We proceed in two steps: First, we examine which occupations are discontinued at the firm level by tracking occupations that are present within a firm in a given year

²⁸The magnitude of coefficients increases as expected since the number of hierarchy levels is substantially reduced.

but not present in the following year(s). Second, we examine the stability of the hierarchical structure by tracking changes in the number of actually populated hierarchy levels over time.

5.4.1 Discontinued Occupations

Extensive discussions revolve around whether certain business functions should be provided within the firm or externally hired. [Goldschmidt & Schmieder \(2017\)](#) document the impact of outsourcing on the wage structure in Germany. They concentrate on the outsourcing of specific low-skilled occupations, such as cleaners, food preparers, or security personnel. While a thorough analysis of outsourcing events is beyond the scope of this paper, our method allows us to predict which occupations are discontinued at the firm level more broadly in a data-driven way. Specifically, we show that single occupations not connected to an internal labor market within a firm are much more likely to be discontinued than occupations that are part of an ILM. To do so, we construct an indicator that equals one if an occupation is present at a firm in year t but not in year $t + 1$. On average, 14% of all occupations that are present in t are absent in $t + 1$. We then regress this indicator on a range of occupation and firm characteristics. Appendix Table A4 presents regression results with and without occupation and firm fixed effects. Overall, we find the most important predictor for being discontinued is whether the occupation is a singleton occupation (i.e., with no connections to the firm's ILMs). The probability of being discontinued increases by 4.7 to 6.9 percentage points for singletons compared to occupations within an ILM. Importantly, the difference between single occupations and occupations within ILMs is not driven by compositional effects, as can be seen by comparing the coefficient across columns 1 to 3. Further, the regression lends support to polarization of occupations within the firm. We see that occupations with low average person effects are more likely to be discontinued than occupations with high average person effects, and low-wage occupations are more likely to be discontinued than high-wage occupations.

5.4.2 Stability of Hierarchies

We now assess whether hierarchy structures within internal labor markets are stable over time. Firms can add levels to their job ladders over time, e.g., by opening a new production line, or they can reduce the number of levels, e.g., due to outsourcing. Our analysis of the firm's structural stability focuses on the largest ILM in 3,607 firms with at least two hierarchy levels.

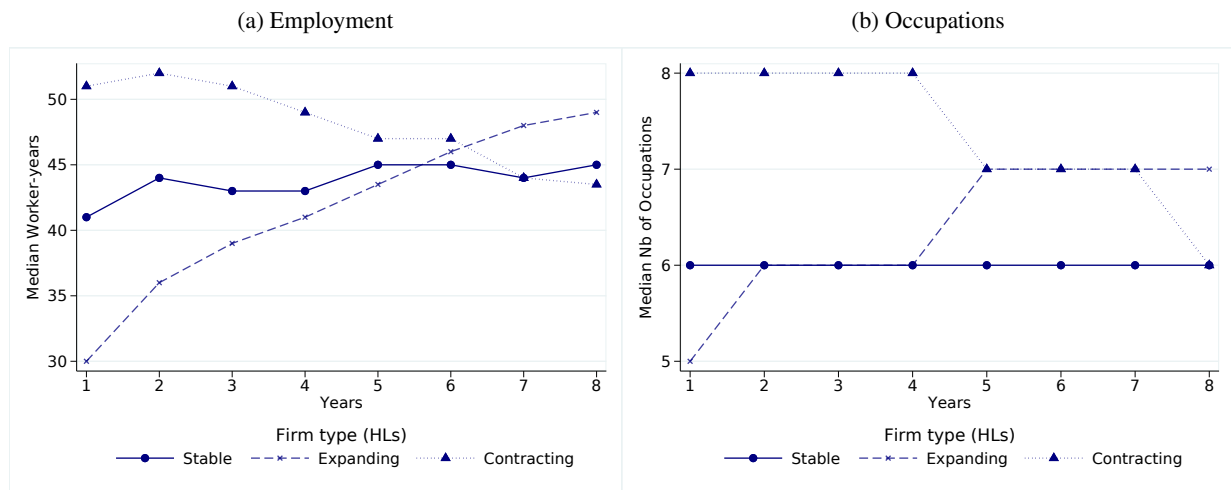
In our sample, 80 percent of ILMs with more than one hierarchy level change the number of levels over time. Among those that change, roughly two thirds expand by adding one or more levels, and the other third removes levels. Table 6 shows annual employment growth rates for the different types of firms. Most remarkably, employment growth differs between expanding and contracting firms. While 70 percent of all firms increase their employment over the eight year observation period, 79 percent of expanding firms and only 48 percent of contracting firms grow. The median firm grows by four percent annually.

Table 6: Firm growth form 2007 to 2014

	Median Number of Workers in first year	Share of Growing Firms	Median Annual Growth Rate	Number of Firms
Stable Hierarchies	41	0.66	0.03	727
Expanding Firms	30	0.79	0.06	1,947
Contracting Firms	51	0.48	-0.003	933
All Firms	37	0.68	0.04	3,607

Notes: This table shows average annual employment growth rates of firms with more than one hierarchy level between 2007 and 2014. The sample includes 3,607 firms at least two hierarchy levels. The sample is described in Section 2.1.1.

Figure 14: Organizational Structures



Notes: The figures show the development of employment, number of occupations, and log wages over time in firms with different types of structures. The sample includes 892 firms with stable structures, 396 expanding firms, 218 contracting firms, and 516 firms that are both expanding and contracting (see further details of the sample restrictions in Section 2.1.1).

Figure 14 shows time patterns in the median number of workers and median number of occupations employed in the three types of firms. The graphs confirm that changes in the hierarchy structure of the largest ILM in the firm are associated with corresponding changes in employment and occupations. Firms with stable hierarchies employ roughly the same number of workers in the same number of occupations each year. In contrast, firms with expanding hierarchies add a substantial fraction of their workforce over time; at the median they steadily grow from 30 to 50 workers over eight years and they add two occupations to the ILM. Contracting firms reduce their workforce and the number of occupations, see Figures 14a and 14b. Our finding that in medium-sized firms, strong employment growth is associated with an expansion of the hierarchy structure, is consistent with [Caliendo et al. \(2015a\)](#) who show that French manufacturing firms grow by adding additional levels to their hierarchies (see also [Caliendo et al., 2015b](#) and [Friedrich, 2015](#)). This finding suggests it may be necessary to adapt the algorithm to changes in the firm structure, and allow the hierarchical structure to vary before and after structural changes in organizations.

6 Conclusion

This paper developed a method to study how promotion dynamics and wages are shaped by the structure of internal job hierarchies in firms. It addressed the conceptual challenge of ordering different job titles in firms into levels of a hierarchy by using observed flows between jobs within firms. In the first step of our approach, we identified occupational networks from observed transitions between positions within the firm. These networks were allowed to be segmented so that a firm can consist of multiple internal labor markets. In the second step, we extended the hand-curated approach of [Baker *et al.* \(1994\)](#). We combined a data-driven approach to rank occupations using flow frequencies between positions and clustering techniques to form levels of the job hierarchy.

We applied our method on eight years of linked employer-employee data from Norway – which is particularly well-suited for our purpose, as it reports highly disaggregated occupational codes and tracks contract changes within firms. We documented a wide variation in the structure of internal labor markets and job hierarchies across about 3,600 firms. The evidence provided broad support to theories of careers in organizations, job assignment models, and a wage-setting process linked to promotions systems. We then verified that the method is robust to alternative assumptions about noise in the data and coarser classifications of occupations (e.g., four-digit codes available in many other data sets).

Our analysis scraped the surface of how internal labor markets shape wage-setting in firms. At the same time, we hope our method can help advance our understanding of wage policies and how firms affect wage distributions (e.g., [Card *et al.*, 2013](#); [Song *et al.*, 2018](#); [Lamadon *et al.*, 2019](#)). Having an estimate of firms' organizational structure permits a fresh take on questions about drivers of dispersion in firm quality, sorting of high-wage workers to high-wage firms, and opens further avenues for future research. For example, it enables a comprehensive assessment of how internal labor markets of firms affect the decision to outsource, related to [Goldschmidt & Schmieder \(2017\)](#) who documented impacts of outsourcing certain occupations, such as cleaners or security personnel, on the wage structure in Germany. Our framework extends to studies of external mobility, such as whether workers are willing to take wage cuts to access internal labor markets with more promising wage prospects, and the role of career opportunities in shaping inter-industry wage differentials.

References

- ABOWD, JOHN M., KRAMARZ, FRANCIS, & MARGOLIS, DAVID N. 1999. High Wage Workers and High Wage Firms. *Econometrica*, **67**(2), 251–333.
- ABOWD, JOHN M., CREECY, ROBERT H., & KRAMARZ, FRANCIS. 2002. *Computing Person and Firm Fixed Effects Using Linked Longitudinal Employer-Employee Data*. Working Paper TP-2002-06. US Census Bureau: Longitudinal Employer Household Dynamics.
- ANDREWS, M. J., GILL, L., SCHANK, T., & UPWARD, R. 2008. High wage workers and low wage firms: negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **171**(3), 673–697.
- BAKER, GEORGE, GIBBS, MICHAEL, & HOLMSTROM, BENGT. 1994. The internal economics of the firm: Evidence from personnel data. *The Quarterly Journal of Economics*, **109**(4), 881–919.
- BARTH, ERLING, MOENE, KARL O, & WILLUMSEN, FREDRIK. 2014. The Scandinavian model—an interpretation. *Journal of Public Economics*, **117**, 60–72.
- BAYER, CHRISTIAN, & KUHN, MORITZ. 2019. Which Ladder to Climb? Decomposing Life Cycle Wage Dynamics.
- BERTRAND, MARIANNE, BURGESS, ROBIN, CHAWLA, ARUNISH, & XU, GUO. 2020. The glittering prizes: Career incentives and bureaucrat performance. *The Review of Economic Studies*, **87**(2), 626–655.
- BOGNANNO, M.L. 2001. Corporate Tournaments. *Journal of Labor Economics*, **19**, 290–315.
- BONHOMME, STEPHANE, LAMADON, THIBAUT, & MANRESA, ELENA. 2019. *Discretizing Unobserved Heterogeneity*. Working Paper.
- CALIENDO, LORENZO, MONTE, FERDINANDO, & ROSSI-HANSBERG, ESTEBAN. 2015a. The anatomy of French production hierarchies. *Journal of Political Economy*, **123**(4), 809–852.
- CALIENDO, LORENZO, MION, GIORDANO, OPROMOLLA, LUCA DAVID, & ROSSI-HANSBERG, ESTEBAN. 2015b. *Productivity and organization in Portuguese firms*. Tech. rept. National Bureau of Economic Research.
- CARD, DAVID, HEINING, JOERG, & KLINE, PATRICK. 2013. Workplace Heterogeneity and the Rise of West German Wage Inequality. *Quarterly Journal of Economics*, **128**(3), 967–1015.
- CESTONE, GIACINTA, FUMAGALLI, CHIARA, KRAMARZ, FRANCIS, & PICA, GIOVANNI. 2019. *Insurance Between Firms: The Role of Internal Labor Markets*. Working Paper 11336. CEPR.
- CHIAPPORI, PIERRE ANDRE, & SALANIE, BERNARD. 2003. Testing Contract Theory: A Survey of Some Recent Work. *Advances in Economics and Econometrics*, **1**.
- CLAUSET, AARON, ARBESMAN, SAMUEL, & LARREMORE, DANIEL B. 2015. Systematic inequality and hierarchy in faculty hiring networks. *Science Advances*, **1**(1).
- DEVARO, JED, & KAUFMANEN, ANTTI. 2016. An Opposing Responses Test of Classic Versus Market-Based Promotion Tournaments. *Journal of Labor Economics*, **34**(3), 747–779.
- DOERINGER, PETER, & PIORE, MICHAEL. 1971. *Internal Labor Markets and Manpower Analysis*. Lexington, MA: D. C. Heath and Company (Armonk, NY: London: M. E. Sharpe, Inc., reprinted 1985).
- ERIKSSON, TOR. 1999. Executive Compensation and Tournament Theory: Empirical Tests on Danish Data. *Journal of Labor Economics*, **17**(2), 262–280.
- FRIEDRICH, BENJAMIN. 2015. *Trade shocks, firm hierarchies and wage inequality*. University of Aarhus, Department of Economics.
- GABAIX, XAVIER, & LANDIER, AUGUSTIN. 2008. Why has CEO pay increased so much? *The Quarterly Journal of Economics*, **123**(1), 49–100.

- GARICANO, LUIS, & HUBBARD, THOMAS N. 2007. Managerial leverage is limited by the extent of the market: Hierarchies, specialization, and the utilization of lawyers' human capital. *The Journal of Law and Economics*, **50**(1), 1–43.
- GARICANO, LUIS, & ROSSI-HANSBERG, ESTEBAN. 2006. Organization and inequality in a knowledge economy. *The Quarterly Journal of Economics*, **121**(4), 1383–1435.
- GIBBONS, ROBERT. 1998. Incentives in organizations. *Journal of economic perspectives*, **12**(4), 115–132.
- GIBBONS, ROBERT, & KATZ, LAWRENCE. 1992. Does unmeasured ability explain inter-industry wage differentials? *The Review of Economic Studies*, **59**(3), 515–535.
- GIBBONS, ROBERT, & WALDMAN, MICHAEL. 1999a. Careers in organizations: Theory and evidence. *Handbook of labor economics*, **3**, 2373–2437.
- GIBBONS, ROBERT, & WALDMAN, MICHAEL. 1999b. A theory of wage and promotion dynamics inside firms. *The Quarterly Journal of Economics*, **114**(4), 1321–1358.
- GIROUD, XAVIER, & MUELLER, HOLGER M. 2019. Firms? Internal Networks and Local Economic Shocks. *The American Economic Review*, forthcoming.
- GOLDSCHMIDT, DEBORAH, & SCHMIEDER, JOHANNES F. 2017. The rise of domestic outsourcing and the evolution of the German wage structure. *The Quarterly Journal of Economics*, **132**(3), 1165–1217.
- KATZ, LAWRENCE F, *et al.* 1999. Changes in the wage structure and earnings inequality. *Pages 1463–1555 of: Handbook of labor economics*, vol. 3. Elsevier.
- KAUHANEN, ANTTI, & NAPPARI, SAMI. 2012. Career and Wage Dynamics: Evidence from Linked Employer-Employee Data. *Research in Labor Economics*, **36**, 35–76.
- KLINE, PATRICK, SAGGIO, RAFFAELE, & SØLVSTEN, MIKKEL. 2020. Leave-out estimation of variance components. *Econometrica*, **88**(5), 1859–1898.
- KRUEGER, ALAN B., & SUMMERS, LAWRENCE H. 1988. Efficiency Wages and the Inter-Industry Wage Structure. *Econometrica*, **56**(2), 259–293.
- LAMADON, THIBAUT, MOGSTAD, MAGNE, & SETZLER, BRADLEY. 2019. *Imperfect Competition, Compensating Differentials, and Rent Sharing in the U.S. Labor Market*. Working Paper 25954. NBER.
- LAZEAR, EDWARD P. 1989. Pay equality and industrial politics. *Journal of political economy*, **97**(3), 561–580.
- LAZEAR, EDWARD P., & OYER, PAUL. 2004. The Structure of Wages and Internal Mobility. *The American Economic Review*, **94**(2), 212–216.
- LAZEAR, EDWARD P., & ROSEN, SHERWIN. 1981. Rank-Order Tournaments as Optimum Labor Contracts. *Journal of Political Economy*, **89**(5), 841–864.
- LAZEAR, EDWARD P, & SHAW, KATHRYN L. 2009. Wage structure, raises and mobility: An introduction to international comparisons of the structure of wages within and across firms. *Pages 1–57 of: The structure of wages: An international comparison*. University of Chicago Press.
- LAZEAR, EDWARD P, SHAW, KATHRYN L, & STANTON, CHRISTOPHER T. 2015. The value of bosses. *Journal of Labor Economics*, **33**(4), 823–861.
- MCCUE, KRISTIN. 1996. Promotions and wage growth. *Journal of Labor Economics*, **14**(2), 175–209.
- NEAL, DEREK. 1999. The Complexity of Job Mobility among Young Men. *Journal of Labor Economics*, **17**(2), 237–261.
- SATTINGER, MICHAEL. 1975. Comparative advantage and the distributions of earnings and abilities. *Econometrica: Journal of the Econometric Society*, 455–468.

- SCHUBERT, GREGOR, STANSBURY, ANNA, & TASKA, BLEDI. 2020. *Monopsony and Outside Options*. Working Paper.
- SHAPIRO, CARL, & STIGLITZ, JOSEPH E. 1984. Equilibrium Unemployment as a Worker Discipline Device. *American Economic Review*, **74**(3), 433–444.
- SONG, JAE, PRICE, DAVID J, GUVENEN, FATIH, BLOOM, NICHOLAS, & VON WACHTER, TILL. 2018. Firming up inequality. *The Quarterly journal of economics*, **134**(1), 1–50.
- TOPEL, ROBERT H., & WARD, MICHAEL P. 1992. Job Mobility and the Careers of Young Men. *Quarterly Journal of Economics*, **107**(2), 439–479.
- VAN DER KLAUW, BAS, & DA SILVA, ANTÓNIO DIAS. 2011. Wage dynamics and promotions inside and between firms. *Journal of Population Economics*, **24**(4), 1513–1548.
- WALDMAN, MICHAEL. 1984. Job Assignment, Signalling, and Efficiency. *RAND Journal of Economics*, **8**, 255–287.
- WANG, H, & SONG, M. 2011. Ckmeans.1d.dp: Optimal k-means clustering in one dimension by dynamic programming. *The R Journal*, **3**(2), 29–33.

A Appendix: Additional Tables and Figures

Table A1: Descriptive Statistics

	Hierarchy levels								
	1-2	3	4	5	6	7	8	9+	All
In wage	10.4	10.4	10.4	10.4	10.5	10.5	10.5	10.6	10.5
Wage growth	0.079	0.075	0.074	0.072	0.072	0.072	0.076	0.077	0.075
Tenure	57.9	59.3	59.7	64.1	66.6	62.2	71.1	71.9	63.6
Age	37.9	37.9	38.1	38.3	39.0	38.7	39.1	40.1	38.6
Female	0.27	0.29	0.27	0.29	0.27	0.28	0.26	0.27	0.28
Married	0.47	0.46	0.46	0.47	0.49	0.48	0.49	0.53	0.48
6-9 years education	0.044	0.042	0.044	0.042	0.043	0.045	0.044	0.047	0.044
10-13 years education	0.65	0.68	0.67	0.66	0.65	0.63	0.66	0.61	0.65
14-16 years education	0.30	0.28	0.29	0.29	0.31	0.33	0.30	0.34	0.30
AKM Firm fixed effect	0.19	0.17	0.18	0.18	0.20	0.20	0.22	0.25	0.20
AKM Person fixed effect	10.2	10.3	10.2	10.2	10.3	10.3	10.3	10.3	10.3
Fraction stayers	0.68	0.70	0.69	0.71	0.70	0.70	0.71	0.71	0.70
Fraction hired internally	0.082	0.076	0.081	0.082	0.090	0.092	0.097	0.11	0.087
Fraction hired externally	0.24	0.22	0.23	0.21	0.21	0.21	0.19	0.18	0.21
Fraction exiting to other internal job	0.096	0.085	0.092	0.091	0.096	0.100	0.096	0.10	0.094
Fraction exiting to other external job	0.22	0.21	0.20	0.19	0.19	0.18	0.18	0.17	0.19
Occupation change within ILM	0.076	0.072	0.077	0.077	0.086	0.087	0.091	0.099	0.082
Fraction promoted	0.067	0.063	0.069	0.070	0.076	0.075	0.078	0.083	0.072
Fraction demoted	0.0030	0.0031	0.0037	0.0040	0.0058	0.0079	0.0095	0.012	0.0058
Fraction lateral moves	0.0064	0.0054	0.0042	0.0038	0.0043	0.0045	0.0038	0.0034	0.0044
Number of occupations in ILM	4.23	5.75	7.88	11.1	14.8	18.7	22.2	50.6	17.0
Average number of workers per year	35.9	39.7	50.0	62.3	85.3	140.8	155.4	456.6	132.7
Number of worker-years	217.0	274.9	343.4	452.1	626.4	1056.4	1175.2	3501.8	990.8
Number of Firms	247	632	870	546	331	191	161	629	3607
Annual growth in total employees	0.16	0.13	0.17	0.15	0.12	0.23	0.072	0.43	0.20

Notes: The table displays characteristics by firm types. The baseline sample is described in Section 2.1.1. The sample consist of 3,607 firms where the largest component has more than one occupation. Firm types are defined by the number of levels. 8+ levels refer to firms with 8 – 66 levels. Component level variables refer to the largest component in the firm.

Table A2: Comparison 7, 4, 1 digit occupations

	7 digit				4 digit				1 digit			
	Mean	Median	p25	p75	Mean	Median	p25	p75	Mean	median	p25	p75
Number of firms	3,611				3,611				3,611			
Number of occupations	6,092				524				10			
Number of ILMs	15,948				7,765				3,795			
Number of singletons	82,045				37,062				4,703			
Occupations												
Firms per occupation	28.6	5	2	21	166	19.5	1	120	2069	2360	507	3295
by Firm												
Number of occupations	48.3	36	23	59	24.1	21	14	31	6.3	6	5	8
Number occ in avg. year	27.7	20.3	13	33	16.3	13.7	9	20	5.4	5.4	4.4	6.5
Number occ in all years	10.8	7	3	14	8.7	7	3	11	4.3	4	3	6
Share occ in all years	0.24	0.21	0.11	0.33	0.36	0.35	0.21	0.5	0.68	0.71	0.5	0.86
Number of internal moves	103	28	16	62	72	19	11	43	51	13	7	32
Share intern movers	0.072	0.063	0.041	0.093	0.049	0.04	0.023	0.063	0.036	0.028	0.016	0.46
Share internal hires	0.26	0.24	0.16	0.35	0.2	0.17	0.1	0.26	0.15	0.13	0.07	0.21
Share promotions in internal movers	0.822	0.85	0.74	0.93	0.77	0.81	0.68	0.91	0.73	0.74	0.62	0.88
Number of ILMs	4.4	4	2	6	2.2	2	1	3	1.05	1	1	1
Number of singletons	22.7	16	9	28	10.3	8	5	13	1.3	1	0	2
by Largest ILM												
Number of hierarchy levels	6.4	5	4	7	4.8	4	3	6	3.2	3	3	4
Number of occupations	17	10	6	19	11.1	8	5	14	4.9	5	4	6
Share occupations	0.36	0.34	0.21	0.5	0.46	0.44	0.3	0.61	0.77	0.8	0.66	1
Share workers	0.68	0.75	0.54	0.86	0.77	0.86	0.7	0.94	0.92	0.99	0.94	1
Share movers	0.088	0.076	0.049	0.112	0.06	0.048	0.027	0.076	0.038	0.03	0.017	0.048
Share internal hires	0.32	0.3	0.19	0.42	0.23	0.2	0.11	0.32	0.16	0.14	0.08	0.22
Share promotions in internal movers	0.81	0.84	0.73	0.94	0.77	0.8	0.67	0.93	0.73	0.74	0.62	0.88

Notes: This table shows summary statistics for 7 digit, 4 digit and 1 digit occupations. The sample is described in Section 2.1.1.

Table A3: Summary of Hierarchy Regressions

Panel A	Baseline					
	(1) Employees	(2) Number of occupations	(3) Share external hires	(4) Manager	(5) Log Wages	(6) AKM Person Effect
Hierarchy Level	-6.424*** (0.667)	0.0606*** (0.00407)	-0.0386*** (0.00209)	0.0207*** (0.00125)	0.0338*** (0.00185)	0.0136*** (0.000864)
Squared Hierarchy Level	0.0842*** (0.0123)	-0.00108*** (0.000111)	0.000601*** (0.0000626)	-0.000284*** (0.0000335)	-0.000505*** (0.0000548)	-0.000193*** (0.0000230)
Observations	146570	146570	102855	146570	146570	146570
Adjusted R^2	0.211	0.168	0.170	0.227	0.560	0.406
Dependent mean	24.38	2.024	0.664	0.133	10.57	10.31
Panel B	Naïve clustering					
	(1) Employees	(2) Number of occupations	(3) Share external hires	(4) Manager	(5) Log Wages	(6) AKM Person Effect
Hierarchy Level	-1.242*** (0.255)	0.0356*** (0.00554)	-0.0355*** (0.00278)	0.0168*** (0.00123)	0.0287*** (0.00233)	0.0110*** (0.000906)
Squared Hierarchy Level	0.0183*** (0.00415)	-0.000665*** (0.000167)	0.000395*** (0.0000722)	-0.000156*** (0.0000245)	-0.000305*** (0.0000631)	-0.000110*** (0.0000219)
Observations	154475	154475	107330	154475	154475	154475
Adjusted R^2	0.164	0.163	0.178	0.215	0.546	0.397
Dependent mean	23.72	2.103	0.677	0.129	10.57	10.30
Panel C	Ad hoc clustering					
	(1) Employees	(2) Number of occupations	(3) Share external hires	(4) Manager	(5) Log Wages	(6) AKM Person Effect
Hierarchy Level	19.03*** (5.339)	1.496*** (0.278)	-0.276*** (0.00667)	0.0571*** (0.0103)	0.163*** (0.00934)	0.0307*** (0.00586)
Squared Hierarchy Level	-5.447*** (1.010)	-0.257*** (0.0503)	0.0253*** (0.00128)	0.00288 (0.00214)	-0.00937*** (0.00185)	0.00185 (0.00117)
Observations	71250	71250	51875	71250	71250	71250
Adjusted R^2	0.497	0.490	0.227	0.379	0.648	0.520
Dependent mean	50.16	4.164	0.717	0.110	10.50	10.28

Standard errors in parentheses are robust and clustered at the individual level * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Notes: This table reports results from regressing a number of career outcomes on hierarchy level using the three clustering methods discussed in Section 5. Standard errors in parenthesis. The sample is described in section 2.1.1.

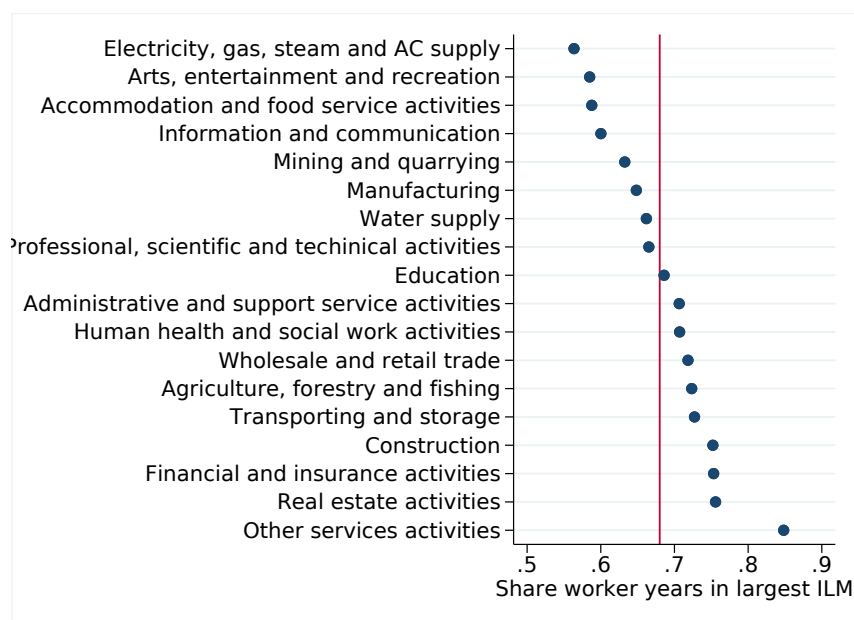
Table A4: Discontinued Occupations

	(1)	(2)	(3)
Indicator for singleton occupation	0.051*** (0.003)	0.069*** (0.002)	0.047*** (0.003)
Average wage in firm	-0.039*** (0.003)	-0.037*** (0.003)	-0.025*** (0.004)
Firm Effect Quintiles			
2	-0.047*** (0.002)	-0.035*** (0.002)	-0.048*** (0.002)
3	-0.043*** (0.002)	-0.029*** (0.002)	-0.049*** (0.002)
4	-0.024*** (0.003)	-0.010*** (0.002)	-0.036*** (0.003)
5	0.007*** (0.003)	0.014*** (0.003)	-0.012*** (0.003)
Firm FE	no	yes	no
Occupation FE	no	no	yes
Adjusted R2	0.01	0.08	0.03
Observations	624,230	624,230	623,591

Standard errors in parentheses are robust and clustered at the firm level * p<0.05, ** p<0.01, *** p<0.001.


Notes: This table shows regression coefficients for regressions on the occupation by firm level. Dependent variable is an indicator that equals one if an occupation is present in a firm in year t but not present in year t+1. Additional regressors are year fixed effects, AKM firm effect, and the share of female employees in the occupation by firm cell. The sample is described in Section 2.1.1.

Figure A1: Share of Workers in Largest IML




Notes: The figures show the average number of internal labor markets by industries for our sample of firms (see further details of the sample restrictions in Table A1).

Figure A2: Hierarchy Levels By Levels of Occupational Aggregation



Notes: The figures report hierarchy levels by levels of occupational aggregation. Panel a shows the result using four digits and panel b shows the results using one digit of the occupational code. The sample is described in Section 2.1.1

Figure A3: Wage Regressions By Levels of Occupational Aggregation



Notes: The figures report regression coefficients of levels on log wages. Regressions include fully interacted firm and year dummies and dummies for gender and each age. Panel a shows the result using four digits and panel b shows the results using one digit of the occupational code. The sample is described in Section 2.1.1. Vertical bars represent 95 percent confidence intervals. Standard errors are clustered at the individual level.

B Details on the Algorithms

B.1 Internal Labor Markets

We identify ILMs as connected sets in the network of occupations within firms. Let $G = (U, V, E)$ denote a firm-specific bi-partite graph where U denotes the set of workers in the firm, V denotes the set of occupations, and E denotes the set of links between workers and occupations. A connected component is a subgraph of G in which any two occupations are connected to each other by some path of worker transitions, and which is connected to no additional occupations in G . Computing connected components

is a prerequisite for several applications in Economics, e.g., the wage decomposition in AKM. Functions that identify components in a network typically rely on a breadth-first search algorithm and are available in standard statistical packages such as *igraph* in R or *Boost Graph Library* in MATLAB.

In order to account for potential measurement error in occupational coding, we propose a leave-X-percent-out procedure that prunes the data. This pruning procedure requires that the bi-partite network remains connected when any one worker is removed. This boils down to finding workers that constitute cut vertices or *articulation points* in the bi-partite network. In contrast to the standard procedure of identifying these cut points, we cut the network only if removing a single worker destroys less than X percent of the vertices that are attached to an occupation.

Algorithm 1: Pruning Network

Result: Leave-X-percent-out Component Structure

compute degree d_v for each occupation in $v \in V$;

construct G' where each link that enters occupation $v \in V$ is duplicated $100/(d_v * X)$ times ;

construct G'_1 from G' by deleting all workers that are articulation points in G' ;

remove duplicated links and return G_1 ;

B.2 Minimum Violations Ranking

We rank occupations within each ILM using a minimum violations ranking (Clauset *et al.*, 2015). In contrast to the first step of our algorithm, the direction of the worker transition that links two occupations is important for this step. Let $H = (V, T)$ denote a network of occupations within a firm V that are connected by directed and weighted worker transitions, T . Our goal is to define a ranking of all occupations $v \in V$ such that the number of "violations" is minimized. A violation is a worker transition (u, v) where the rank of the origin occupation u is higher than the rank of the target occupation v . Complex networks like job to job transitions within firms often exhibit multiple minimum violation rankings, in which several distinct orderings of occupations produce the same smallest number of violations. We apply the sampling procedure proposed in Clauset *et al.* (2015) to find a consensus ranking in the case of several possible optimal rankings. The algorithm starts from an initial ranking that sorts occupations according to their out-degree. At each step in the algorithm, a randomly chosen pair of occupations is chosen and a ranking is proposed in which their ranks are swapped. We accept each proposal that leads to a lower or equal number of violations. Once the algorithm has converged on the minimum number of violations, this might still lead to different accepted rankings (since we are also accepting neutral proposals). We then sample from the set of rankings with the minimum number of violations and average the ranks for each occupation to get a consensus ranking. We repeat the procedure R times to compute the standard deviation of mean ranks across repetitions as our measure of uncertainty.

Algorithm 2: Minimum Violations Ranking Algorithm

Result: Average mean ranking over R repetitions, std. dev. of mean rankings

```
for  $i = 1$  to  $R$  do
    sort occupations according to out-degree (decreasing);
    compute number of violating links  $S$  ;
    set  $t = 0$  ;
    while  $t < T$  do
         $t = t + 1$  ;
        switch ranks of two randomly chosen occupations ;
        compute number of violating links  $S'$  ;
        if  $S' \leq S$  then
            if  $S' < S$  then
                 $S = S'$  ;
                delete previously stored rankings ;
                 $t = 1$  ;
            end
            store updated ranking ;
        end
    end
    store number of violations and average ranking over samples ;
end
compute summary stats over all repetitions: mean ranking, std. ranking ;
```

B.3 Clustering Ranks to Hierarchy Levels

In order to group occupations into hierarchy levels, we apply a one-dimensional kmeans clustering algorithm based on the estimated average rank r_v for each occupation $v \in V$. In particular, we partition the N occupations within a firm into K groups, corresponding to group indicators $\hat{k}_v \in 1, \dots, K$ in order to minimize the objective function

$$(\hat{h}, \hat{k}_1, \dots, \hat{k}_N) = \arg \min_{h, k_1, \dots, k_N} \sum_{v=1}^N (r_v - h(k_v))^2 \quad (\text{B1})$$

where $\hat{h}(k)$ is the mean of r_v in group $\hat{k}_i = k$ and K is a pre-determined number of hierarchy levels. Wang & Song (2011) provide an optimal dynamic programming algorithm to minimize (B1) in this one-dimensional setting that leads to guaranteed optimality and is reproducible since it does not rely on computational approximations such as kmeans in higher dimensions.

Following Bonhomme *et al.* (2019), we choose the number of hierarchy levels, K , such that the within-group variation in ranks is of the same magnitude as the overall noise level in the estimation of the ranks. Specifically, let $\hat{Q}(K) = \frac{1}{N} \sum_{v=1}^N (r_v - \hat{h}(k_v))^2$ denote the value of the kmeans objective function corresponding to K groups. Then, we choose

$$\hat{K} = \min_{K \geq 1} \left\{ K : \hat{Q}(K) \leq \frac{1}{N-1} \sum_{v=1}^N \text{Var}(r_v) \right\} \quad (\text{B2})$$

where $Var(r_v)$ is the variance of the estimated rank obtained from the minimum violations ranking algorithm.

to K groups. Then, we choose

$$\hat{K} = \min_{K \geq 1} \left\{ K : \hat{Q}(K) \leq \frac{1}{N-1} \sum_{v=1}^N \text{Var}(r_v) \right\} \quad (\text{B2})$$

where $\text{Var}(r_v)$ is the variance of the estimated rank obtained from the minimum violations ranking algorithm.