

Job levels and Wages*

Christian Bayer[†]

Moritz Kuhn[‡]

September 8, 2026

Abstract

Job levels summarize the complexity, autonomy, and responsibility involved in executing tasks. They are conceptually distinct from occupations and education. Our finding that job levels can be constructed from task execution data demonstrates their economic content. We present a novel theory of employer job design in which jobs with different levels emerge endogenously. Using matched employer-employee data, we demonstrate the theory-consistent mediating role of job levels in life-cycle wage growth, gender wage gaps, and returns to education and seniority. Our theory interprets some wage differences as arising from differences in human capital utilization and offers a new perspective on observed wage differences.

Keywords: job levels, wage structure, career ladder

JEL Codes: D33, E24, J31.

*We thank Simon Heiler for outstanding research assistance and Jim Albrecht, Thomas Dohmen, Jan Eeckhout, Simon Jäger, Gueorgui Kambourov, Fatih Karahan, Greg Kaplan, Francis Kramarz, Iouri Manovskii, Simon Mongey, Elena Pastorino, Esteban Rossi-Hansberg, Todd Schoellman, Ludo Visschers, Michael Waldman, and conference and seminar participants at various places for comments and suggestions. The authors gratefully acknowledge support through the project ADEMU, “A Dynamic Economic and Monetary Union,” funded by the European Union’s Horizon 2020 Program under grant agreement No. 649396 and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2126/1– 390838866, as well as under CRC-TR-224 EPoS (Projects A03 and C05), and KFG-50 CASFI (Bayer). Bayer would like to thank the EUI for hosting him while conducting part of the research that led to this paper. Kuhn would like to thank the Federal Reserve Bank of Minneapolis for hosting him while conducting part of the research that led to this paper.

[†]Universität Bonn, ECONtribute CEPR, and IZA, Adenauerallee 24-42, 53113 Bonn, Germany, *email*: christian.bayer@uni-bonn.de

[‡]Universität Mannheim, CEPR, and IZA, L7 3-5, 68161 Mannheim, Germany, *email*: mokuhn@uni-mannheim.de

1 Introduction

The question of what determines a worker’s wage has been a long-standing topic of economic research. We offer a fresh perspective on the topic by combining the analysis of largely unexplored data with novel economic theory. Based on the task-based view of wage determination (Autor et al., 2003), our theory rationalizes the existence of distinct jobs within the production process and delivers a wage structure in which pay is tied to jobs and their content. We expand the existing view that builds on the task content derived from occupations and introduce the concept of *job levels*. First, we show that job levels, in the data, measure the complexity, autonomy, and responsibility associated with task execution and provide a new theory in which job levels are the solution to an information problem in which workers’ skills can only be set-identified, i.e. known up to a range. Job levels allow employers to design jobs that group tasks of similar complexity to match the skills of workers identified in groups. In model and data, higher job levels are more autonomous, more productive, and better paid. Second, we demonstrate that the concept of job levels provides a unified interpretation of several important wage facts, which are the result of workers progressing through job levels at different rates.

The concept of job levels has a long history in labor market statistics, dating back to the 1950s. Job levels often form the basis of union collective bargaining agreements and corporate job-based compensation schemes. For example, the U.S. Bureau of Labor Statistics has reported wages by job level in its White Collar Pay Survey since at least 1959. Similarly, the German Statistical Office has reported them since at least 1957. Additionally, many unions rely on job-leveling schemes for collective bargaining agreements, and firms hire job-leveling consultancies—even outside of collective bargaining. This strongly suggests the conceptual importance of job levels. However, the limited availability of job-level data in commonly used datasets has prevented a systematic economic analysis of the relationship between job levels and wages. Our paper addresses this issue by using high-quality, administrative, matched employer-employee data to demonstrate the significance of job levels in explaining observed wage differences.

What is the *level* of a job? Occupations describe the tasks that workers perform and have been widely studied in economic research (e.g. Kambourov and Manovskii, 2009a; Autor et al., 2003). Job levels provide an additional distinction in task execution, comparable across and within occupations in terms of complexity, autonomy, and responsibility (CAR). The complexity dimension defines the minimum skills required for the job, consistent with the task-based approach of matching workers’ skills to jobs and their associated tasks (Acemoglu and Autor, 2011). Autonomy and responsibility are closely related to the organization of the production process because they describe how work is organized. Thus, our findings confirm Acemoglu and Autor’s conjecture that task execution is ultimately related to “(...) *the allocation of authority within the organization (...) and the nature of the responsibility system*” (p. 84). For an example of differences in job levels, consider two bakers in the same five-digit occupation (SOC 51-3011). One baker follows recipes and rules for mixing and baking dough. The other baker, in addition, develops new recipes. According to the U.S. O*Net task content, both perform occupational baker tasks, but as their autonomy and responsibility differ, so do their job levels.

Our study of job levels and wages is divided into two parts. First, we introduce the empirical

concept of job levels and our new theory that rationalizes their existence. Using data from various sources, we thoroughly explain the concept of job levels, demonstrate their economic significance, and clarify how they differ from occupations and traditional measures of job tasks. We then develop a model in which job levels arise endogenously when firms can only identify workers’ skills but need to assign them to tasks of different complexities. The second part derives model predictions for wage dynamics and provides empirical evidence supporting the idea that job levels are the key mediator behind observed wage differences. Specifically, we demonstrate in the model and the data that moving across job levels during one’s working life is the main driver of life-cycle wage growth and inequality. Furthermore, we demonstrate that job-level differences and their life-cycle dynamics are the primary determinants of wage differences by gender, education, and seniority. Group differences arise when a group progresses more quickly to higher job levels. Thus, we link a large part of residual between-worker wage differences to directly observable differences in job levels, offering a new, theory-based perspective on these widely studied wage facts.

Our main data sources are the four waves of the Structure of Earnings Survey (SES), which cover the period 2006 to 2018. Each wave provides worker-level information on job levels and wages. Job levels alone account for 47% of the variation in wages, while all observable characteristics account for more than 80%, making the SES data an ideal source for studying the determinants of wages. To demonstrate the economic significance of job levels, we supplement the SES data with job requirement information from the BIBB/BAuA Employment Survey. Using the publicly available job-leveling scheme of the steel and metal workers’ collective bargaining agreement in North Rhine-Westphalia, which covers a large number of workers, we construct job-leveling factors from workers’ reported job requirements. These job leveling factors serve as building blocks for constructing job levels. We find that they account for 44% of the wage dispersion in the BIBB/BAuA data. In combination with the collective bargaining agreement, the job-leveling factors directly demonstrate the economic content of job levels by grading a job’s complexity, autonomy, and responsibility (CAR intensity). Importantly, we also demonstrate in the SES data that job levels differ from occupations and account for a significant amount of wage dispersion within and across occupations. Within an occupation, we typically find a significant proportion of workers in three of the five possible job levels. Furthermore, we demonstrate that task-based wage differences ([Autor et al., 2003](#)) are largely absorbed by average job-level differences. Thus, task-based (occupational) wage components alone account for little of the wage dispersion when job levels are considered. Job levels confirm the general idea of the task-based approach, but also provide an important refinement. Using data from the National Compensation Survey (NCS), which is conducted by the U.S. Bureau of Labor Statistics (BLS), we generalize our findings beyond the German case.¹ Our results from German and U.S. data are very similar, which leads us to conclude that the importance of job levels for macroeconomic wage differences is not peculiar to the German labor market.

Given the empirical importance of job levels in accounting for wage differences, we develop a theory to interpret these findings. Within a task-based framework, we formalize the concept of

¹The NCS provides data on job levels, and it has already been documented that similarly striking results on the explanatory power of job levels for wages apply to the U.S. labor market ([Pierce, 1999](#)).

job levels. Tasks are ordered by difficulty, and workers differ in skills. Therefore, employers want to allocate tasks to workers according to their comparative advantage. However, when screening workers, employers can only determine the range of a worker’s skills, not their exact level. This gives rise to a job design problem, for which the job-level structure is an optimal solution. Under an optimal job design, higher job levels entail higher minimum skill requirements, reflecting greater complexity. They encompass a broader set of tasks, necessitating increased autonomy, and they have a higher marginal product for appropriately qualified workers. A key insight from the theory is that any measured productivity difference stems from two factors: inherent (fixed) skill differences and how these skills are utilized in a job. Thus, our theory offers different interpretations of observed between-group wage differences. While the measured wage differences between educational groups are closely related to the skill differences of these groups, the gender wage gap is the result of the different utilization of the skills of equally skilled men and women. In fact, in a second step, we use the empirical and theoretical results from the first part to decompose wage dynamics over the life cycle by applying synthetic panel methods, a standard tool from the macroeconomic toolkit, to the repeated cross sections of the SES data (Deaton, 1985; Verbeek, 2008). We construct cohort-level panel data and estimate the coefficients of interest based on a cohort-level wage regression. Using the estimated coefficients, we construct the worker-level wage components arising from observable individual and job characteristics, plus an employer component. Using this decomposition, we document the contribution of each component to wage growth and increasing wage dispersion over the life cycle. When we decompose the contribution of job characteristics into job- and occupational-level components, we find that the former account for most of the job component. For life-cycle wage dynamics, we find that career progression—i.e. transitions across job levels as workers age—accounts for 50% of wage growth and nearly all of the increase in wage dispersion. Thus, we find a key mediating role for job levels in life-cycle wage dynamics. We contrast our empirical results on the important role of job-level wage dynamics to a calibrated version of the newly developed model that we embed into a frictional labor market. Career opportunities in the model arise when workers change employers or when workers on higher job levels leave the current employer, opening a path to internal promotion. We calibrate the baseline model for men to the average mobility rates between employment and unemployment and to differences in unemployment rate between skill groups. The parsimoniously calibrated model closely matches the job-level wage increase over the life cycle and the increase in wage dispersion from differences in career progression.

A central prediction of the model is the important role of career progression within the same employer. We compare this prediction with descriptive evidence on career dynamics based on data from the German Socio-Economic Panel (SOEP). We document life-cycle profiles of promotion and demotion rates on the career ladder and demonstrate that our model closely matches the empirical promotion pattern. Additionally, we discuss how labor market mobility between employers, occupations, and through non-employment is associated with steps up and down the career ladder. Although we find that employer and occupational mobility are associated with career progression, we also find that most career ladder moves occur while staying with the same employer and within the same occupation, as predicted by our theory. While such career ladder dynamics have been the topic of single-employer case studies (Baker et al., 1994), we are the first to document their importance at the macroeconomic level.

The SOEP data allow us to directly compare the career dynamics of men and women. In the model, we calibrate women’s life-cycle labor market experiences by matching the observed increase in part-time employment and transition to marginal employment during childbearing age. By interpreting these transitions as career slowdowns (part-time work) and breaks (transitioning to marginal employment), the model closely matches the life-cycle profile of the gender wage gap. Through the lens of the model, the gender wage gap becomes a gender promotion gap, which is consistent with empirical evidence from the SOEP data. In the model, men and women are equally skilled and productive in the same jobs. The underutilization of women’s skills at lower job levels results in a gender wage gap and lower measured productivity when job-level differences are not considered. By contrast, when studying the return to education, we find that it is closely related to inherent skill differences. In the model, wage differences between skill groups result from more highly skilled workers climbing faster and further up the career ladder. Consistent with the mediating role of job levels, we observe a substantial return to education when job-level differences are not considered, but this return largely disappears when they are. We demonstrate a similar mediating role in the model and data for the returns to seniority. In our empirical decomposition, returns to seniority—the higher wages observed for workers who have stayed with their employer longer than their coworkers—stem from wage differences accounted for by job-level differences. Our job-level model generates empirically consistent returns to seniority in simulated data when the mediating role of job levels is not considered.

Our theory supports a model of wage determination in which workers are paid for the tasks they perform rather than directly for the skills they offer. This idea is in principle not new, nor is it original to the concept of job levels. It is, instead, as [Acemoglu and Autor \(2012\)](#) note, a key innovation of the task-based approach. [Acemoglu and Autor](#) further emphasize that the focus on task execution for wage determination is the key difference between a traditional human capital view and the task-based approach. Our results follow, corroborate, and further elaborate this view. To put our results further in perspective to these ideas, it is important to emphasize that measured job levels include more than our initial simple baker example suggests. Job levels combine several aspects of the execution of tasks (see the BLS Job Leveling Guide [U.S. Bureau of Labor Statistics, 2013](#), for the U.S. NCS data). One of these aspects is that some jobs have a (particularly) complex set of tasks, which is reflected in a minimum skill requirement that also arises as optimal solution to our job-design problem. However, this minimum skill requirement allows for situations where workers with a college education are taxi drivers as long as they have the minimum requirement of a driver’s license. Our theory explains that job levels are associated with task ranges that are complementary to worker skills. This then leads to a notion of skill / human capital utilization on the job, e.g., a taxi driver with a college education would not be using all her skills (see [Rosen, 1983](#), for consequences of potential underutilization of human capital for educational choices). Indeed, our quantitative model traces the gender wage gap back to such an underutilization of women’s skills when their careers slow down during childbearing age. The second key aspect encoded in the job level is autonomy in task execution. Our initial example of the baker provides an illustration of this aspect, as the two workers differ in how closely they must follow rules and procedures in performing their occupational tasks. Finally, the responsibility aspect of the job level captures the scope of operations affected by

the job holder’s task execution. For example, if a supervisor leads a team of a few workers, her responsibility is lower than that of a manager whose orders bind the activities of workers throughout the entire organization, even if the manager herself directly supervises few or no workers. Although job levels are independent of specific occupational tasks, i.e. baking bread or making sausages, it is also clear that, for example, managerial occupations have higher job levels by construction. Thus, we should expect job levels to correlate with occupations and other concepts derived from occupations, such as the task-based approach (Autor et al., 2003), with hierarchies within firms (Garicano and Rossi-Hansberg, 2006), or with other aspects of work organization, such as incentives and team structures (Winter, 2006). Importantly, job levels are distinct from job titles, as the latter are not tied to job tasks and their execution, but can be arbitrarily inflated, as recently documented in Cohen et al. (2026).

Finally, it is important to emphasize that we do not fully answer the question of why workers end up in the jobs they have and why some climb the career ladder while others do not (beyond education, gender/part-time, and seniority effects). In this sense, we study the consequences rather than the causes of career progression. On the theoretical side, we leave the intriguing and complicated interaction of occupations and job levels for future research.

The remainder of the paper is organized as follows: Section 2 discusses how our work relates to the existing literature. Section 3 introduces the different datasets used in the empirical analysis. Section 4 discusses job levels and their relation to wages, occupations, and the task-based approach, as well as their economic content. This section also presents our new theory of job levels. Section 5 introduces the decomposition approach used to study life-cycle wage dynamics and reports the decomposition results. The section relates the decomposition results to our theoretical model and provides a new interpretation of between-group wage differences through the lens of our theory. We conclude with Section 6. An appendix follows.

2 Related literature

Our results on task execution encoded in job levels confirm and extend the idea of the task-based approach that the tasks executed determine a job holder’s wage (Autor et al., 2003; Acemoglu and Autor, 2011, 2012). We add to this view a fundamental role for the additional distinction of *how* tasks are executed in terms of complexity, autonomy, and responsibility. Our results thus point to an important role of work organization and the implied distribution of jobs in shaping the macroeconomic wage distribution over time. In this sense, our findings align well with the established results of Katz and Murphy (1992), Krusell et al. (2000), Autor et al. (2003, 2006), and Acemoglu and Autor (2011) that wage inequality is driven by changes in the production process over time, which also helps rationalizing differences in wage dynamics among technologically similar economies (see, e.g., Krueger et al., 2010; Pham-Dao, 2019). Caroli and Van Reenen (2001) highlight explicitly the organizational structure in relation to technological progress. At the same time, we provide evidence in support of extending this view to the organizational structure of the production process, which reflects not only physical but also management techniques Bloom et al. (2016), the composition of the workforce, and labor market institutions (Acemoglu, 2002; Acemoglu and Autor, 2011). The view proposed in

our paper contributes to the macroeconomic approach that places the organizational structure of firms at the center of the analysis. [Caicedo et al. \(2019\)](#) study secular trends in the wage structure and propose a theory of vertical job differentiation as a result of specialization in the production process. [Caliendo et al. \(2015\)](#) provide empirical support for the theoretical model in [Garicano and Rossi-Hansberg \(2006\)](#). [Pastorino \(2024\)](#) proposes a model of employer learning about a worker’s ability that also emphasizes the importance of internal labor markets for wage dynamics. [Kuhn et al. \(2023\)](#) document a relationship in survey data between the coordination in the production process and the average worker pay.

By exploring the sources of wage growth and inequality throughout the life cycle, our work is directly related to the long-standing economic research agenda on the determinants of wage differences, going back at least to the seminal work of [Mincer \(1974\)](#). His work has developed into a large literature documenting a variety of life-cycle patterns of wage growth and inequality. Examples include [Deaton and Paxson \(1994\)](#); [Storesletten et al. \(2004\)](#); [Heathcote et al. \(2005\)](#); [Huggett et al. \(2006\)](#). Today, a common practice is to interpret the residuals from Mincerian wage regressions as wage risk, and a large literature is devoted to estimating stochastic processes for these residuals ([Lillard and Willis, 1978](#); [MaCurdy, 1982](#); [Carroll and Samwick, 1997](#); [Meghir and Pistaferri, 2004](#); [Guvenen, 2009](#)). More recently, [Huggett et al. \(2011\)](#), [Bagger et al. \(2014\)](#), and [Guvenen and Smith \(2014\)](#) have taken more structural approaches to exploring the drivers of life-cycle inequality. We contribute to this literature by relating diverging wages to observable steps on the career ladder and differences between employers. [Kambourov and Manovskii \(2008, 2009a,b\)](#) document the important role of occupations as determinants of wage differences in the cross section and over time. Our results complement this work, emphasizing the importance of job-level differences in addition to occupational or task differences. Differences between employers as a source of wage differences are a prominent theme in the literature that examines secular trends in wage inequality. [Card et al. \(2013\)](#), relying on the estimation approach in [Abowd et al. \(1999\)](#), find that increasing wage differentials between employers are an important contributor to rising wage inequality in Germany. [Song et al. \(2019\)](#) confirm this finding in U.S. Social Security data. [Song et al. \(2019\)](#) and [Card et al. \(2013\)](#) both argue that changes in the organizational structure of firms are the likely driver of the increase in wage differentials between firms. [Low et al. \(2010\)](#), [Hornstein et al. \(2011\)](#), and [Jung and Kuhn \(2016\)](#) are examples that explore employer differences as a source of earnings inequality in search models.

Our findings also relate to the personnel economics literature that studies internal labor markets and career dynamics following the seminal work of [Doeringer and Piore \(1985\)](#). The existing research in this strand of the literature builds on case studies of individual firms and sometimes even subgroups of workers within those firms, as in [Baker et al. \(1994\)](#). [Baker et al. \(1994\)](#) and [Dohmen et al. \(2004\)](#) find that in the absence of promotions across job levels, there is virtually no individual wage growth. [Fox \(2009\)](#) document for Sweden that promotions are a key source of life cycle earnings growth, and [Bronson and Thoursie \(2018\)](#) also document in Swedish panel data gender differences in career progression.² This strand of literature unanimously echoes the key idea formulated in [Doeringer and Piore \(1985, p. 77\)](#) that “[i]n many jobs in the economy, wages are not attached to workers but to jobs.”

²For the United States, [Guvenen et al. \(2021\)](#) document persistent gender earnings gaps at the top.

On the theoretical side, we focus on the life-cycle implications of career progression on wage dynamics as an outcome of a search process, given the structure arising from job levels. [Waldman \(2012\)](#) provides an excellent overview of theoretical career ladder models. The seminal papers are [Lazear and Rosen \(1981\)](#), which explains promotion dynamics as a result of tournaments, and [Waldman \(1984\)](#), which emphasizes the signaling role of promotions in an environment with asymmetric information about workers’ abilities. [Gibbons et al. \(2006\)](#) and [Gibbons and Waldman \(2006\)](#) extend this theory by allowing for complementarity between job levels and skills. As summarized in [Rubinstein and Weiss \(2006\)](#), the underlying assumption of these theories is that wage differences arise solely from workers’ skills, potentially amplified by job assignments that make skills differentially productive.³ In contrast, [Winter \(2004, 2006\)](#) shows that wage differences in teams may arise purely to provide optimal incentives linked to the organizational structure of the team.

3 Data

We rely on three data sources in our empirical analysis: the Structure of Earnings Survey (high-quality repeated cross sections of matched employer-employee data), BIBB/BAuA survey data (a worker survey with details on tasks and wages), and the Socio-Economic Panel (household panel data with information about jobs of household members and their incomes).

3.1 The Structure of Earnings Survey data

The Structure of Earnings Survey (*Verdienststrukturerhebung*, SES) is our main data source. We use the 2006, 2010, 2014, and 2018 waves, which include more than seven million observations of employees from over 100,000 establishments with at least ten employees across all survey years. The German Statistical Office conducts the survey, and establishments are legally required to participate. Establishments with 10 to 49 employees must report data on all employees, while those with 50 or more employees only report data for a representative sample. Data on regular earnings, overtime pay, bonuses, and regular and overtime hours paid are extracted from payroll accounting and personnel master data. This data is then transmitted through a software interface to the statistical office. Unlike German social security data, the SES reports employees’ actual (virtually uncensored) pay and hours worked. The survey also provides detailed information on workers’ education, occupation, age, tenure, and job levels. Self-employed workers are not covered. The survey contains information on approximately 3.2 million employees in 2006, 1.9 million in 2010, and 0.9 million in both 2014 and 2018. The number of employees sampled decreased over time as the sampling probability of plants decreased to reduce bureaucratic costs. In our analysis, we rescale observation weights across surveys, giving each survey equal weight. For our baseline analysis, we restrict the data to workers 25 to 55 years of age. We drop a small number of observations where earnings are censored⁴ and all observations for which the state has a major influence on the plant.⁵ We drop observations from the public administration and

³[Yamaguchi \(2012\)](#) extends this framework to capture the dynamics of endogenous accumulation of unobserved skills, where the rate of accumulation differs across different types of jobs.

⁴The censoring limit in annual gross earnings is €1M (2006) / €750K (2010+). We impose €750K throughout.

⁵We run a robustness check in which we include publicly owned/dominated plants, too; see Appendix H. For

Table 1: Summary statistics for wages and hierarchies in the SES, 2006-2018

	Wages (in 2010 €)					Pop. Share of Job Level (in %)					N. Obs
	Av.	Gini	p10	p50	p90	1	2	3	4	5	
Males											
2006	20.5	0.26	10.5	18.0	32.8	5.8	17.0	43.4	24.3	9.5	706,886
2010	20.3	0.28	9.9	17.6	33.3	7.7	17.2	41.5	22.4	11.1	581,442
2014	21.3	0.27	10.4	18.4	34.8	5.6	13.5	45.9	23.6	11.4	187,568
2018	22.0	0.27	10.8	19.0	36.4	5.7	14.1	45.2	23.3	11.7	175,441
Females											
2006	15.9	0.22	8.7	14.7	23.8	12.5	18.9	46.2	18.5	3.9	431,016
2010	15.8	0.24	8.4	14.4	24.2	13.9	17.5	45.6	18.2	4.8	353,863
2014	16.6	0.24	8.7	14.9	25.9	9.6	15.1	51.4	18.2	5.7	125,185
2018	17.7	0.24	9.5	15.8	27.7	8.2	15.0	52.2	18.5	6.1	116,332

Notes: “Wages” refers to the hourly wages in constant 2010 prices. “Av.” is the average, and “p10/50/90” are the 10th, 50th, and 90th percentiles of the wage distribution, respectively. “Pop. Share of Job Level” refers to the population share of a job level in the sample population. “N. Obs.” refers to the unweighted number of observations in the baseline sample.

mining industry and observations with missing occupation- or job-level information. For our decomposition analysis, we use plant-fixed effects, and therefore drop all observations for which our sample selection by age leaves us with fewer than ten workers at a plant. The baseline sample has 2.67 million worker-plant observations.

As wage measure, we use monthly gross earnings including overtime pay and bonuses divided by regular paid hours and paid overtime hours. As control variables, we use experience, education, sex, occupation, and the job level. We construct experience as potential experience starting at age 25. Sex is naturally coded. For education, we consider four groups: only a *secondary* education, a secondary education with additional *vocational* training, a *college* education. The fourth group, *other*, includes workers for whom education is not reported or who have other levels of education, including workers who have not completed a secondary education.⁶ For occupation coding, we mostly use two-digit 2008 ISCO codes available in the scientific use files. For some particular questions, we use more detailed occupational codes that can be accessed only on site. We rely on a crosswalk provided by the International Labour Organization (ILO) together with additional occupation codes from the German employment agency (KldB 1988) to recode occupations in the 2006 data.⁷ Table 1 reports descriptive statistics for men and women in the baseline sample (number of observations for each wave, average wages, wage inequality, and distribution of workers across job levels). There are five encoded job levels in the SES data, job levels 1 to 5, from 1 being the lowest job level to 5 being the highest job level.

a large set of observations, information on public ownership is missing. The information is available only if in a region-industry cell there are at least three firms in which the state has a major influence. “Major influence” is defined as being a government agency, as the state owning $\geq 50\%$ of shares, or as a result of strong regulation.

⁶Information in the 2014 SES shows that typically workers in “other” have not completed secondary education.

⁷International Labour Organization crosswalk, International Classification of Occupations “Correspondence ISCO-88 to ISCO-08,” www.ilo.org/ilostat-files/Documents/Correspondence_EN_ISCO_08_to_ISCO_88.xlsx.

3.2 BIBB/BAuA data

The BIBB/BAuA Employment Survey is a representative cross-sectional survey among employed persons 15 years and older in Germany. A worker is considered employed if the worker is paid for working at least 10 hours per week. The survey contains information on the occupation, education/training requirements of a job, the actual education/training of the worker, and details of the job tasks the worker executes. We use the 2012 wave, which provides us with detailed data on task execution beyond the task content of occupational data for individual workers. From the task execution data, we select the information used for job leveling. We refer to this information as job-leveling factors. The survey also collects data from workers on their monthly earnings and typical hours worked, worker demographics, industry, and employer size. We use the hours and earnings data to construct wages. Constructed wages in the BIBB/BAuA data likely contain more measurement error than wages from the SES data, which are based on employer-reported earnings and hours. This will reduce explanatory power in the regression analysis.

To align with our SES analysis, we restrict the BIBB/BAuA sample. We include workers between 25 and 55 years of age who work for private sector employers with at least 10 employees and exclude self-employed, freelance, and independent contractor workers, as well as family workers. To reduce the measurement error in hours, we further restrict the sample to workers who do not report second jobs and who report regular working time between 35 and 45 hours per week.⁸ We exclude all observations with imputed wage information.

3.3 Socio-Economic Panel (SOEP) data

The German Socio-Economic Panel (SOEP) data are the equivalent to the U.S. Panel Study of Income Dynamics (PSID) data. The SOEP data provide information on individual labor market situations together with worker demographics and income (Goebel et al., 2019). The individual-level panel data we use cover the period from 1984 to 2015. As part of these data, the SOEP collects information similar to job levels with a coding based on ideas from the sociological literature (Hoffmeyer-Zlotnik, 2003). In fact, Appendix A demonstrates that the life cycle profiles of wages by job level are broadly aligned with the SES data profiles.

We opt for the SES as our preferred data source for all questions that do not look at promotion dynamics because compared to encoded job levels in the SES data, the SOEP coding loads more heavily on education than on actual tasks. In addition, since the SES income data come directly from payroll data, it will have smaller measurement error on income than the SOEP, which also has no coworker information and is much smaller in sample size in each survey year.

To align the SOEP and SES samples, we keep workers aged 25 to 55 years who work for employers with at least 10 employees. We excluded self-employed workers, apprentices, military personnel, and public employees. All observations with missing information on educational level, industry, occupation, or number of employees at their employer are dropped. The data are collected annually, and we explain below how we measure labor market mobility and career dynamics.

⁸Appendix H shows that the results based on SES data change little when considering only full-time workers.

4 Job levels in the data and in theory

The SES data provide a unique opportunity to study the importance of job characteristics encoded in job levels for wage differences. The dataset combines rich administrative information on workers, jobs, and firms, and its wage data is of exceptional quality. In fact, the richness and precision of the data can be seen in the descriptive regression results in Table 2. Using all available information on workers, employers, and jobs, the observed information accounts for more than 81% of the observed cross-sectional variation in wages. For men in the middle of their working lives, the number is even higher, with more than 83%. The high quality of the data is a key reason for this very high degree of statistical determination.

Table 2: Importance of characteristics in explaining hourly wages

	Job levels, occupations, plant-fixed effects, worker controls	Job levels and plant-fixed effects	Job levels, industry, region, and plant size	Job levels	Plant-fixed effects
All workers, age 25-55					
(adj.) R^2	0.813	0.782	0.618	0.471	0.583
Men, age 35-45					
(adj.) R^2	0.832	0.814	0.625	0.488	0.641


Notes: Adjusted R^2 of different regressions on log wages. All regressions contain year fixed effects as additional regressors. Worker controls are education, experience, tenure, sex, and interaction dummies. The first column regression is on job-level dummies, plant fixed effects, occupation dummies, and worker controls; the second column on job-level dummies and plant fixed effects; the third column on job-level dummies, plant size dummies, regional dummies, and industry dummies; the fourth column only on job-level dummies; and the fifth column only on plant fixed effects.

4.1 Job levels and wages

Besides data quality, the other and economically more important reason for the high explanatory power is that we observe job levels. Column 2 shows that using only job levels and plant-fixed effects (neither occupation nor worker characteristics) results in a marginal decline in explanatory power, which is 78% and 81% for middle-aged men. The explanatory power remains at 62% (column 3) when describing firms by their observable characteristics (region, industry, and size) instead of using fixed effects. Column 4 reveals the striking fact about the explanatory power of job levels: five dummies for job levels alone account for 47% (49%) of the cross-sectional wage variation among all workers (men aged 35–45). The tremendously more granular plant-fixed effects (one for each of the more than 76,000 plants in our final sample) account for 58% (64%) of the wage variation in the data.⁹ The high explanatory power of job levels for wages is not

⁹Figure A1 in Appendix A shows the wage differences by job level throughout the entire life cycle in both the SES and SOEP data. Importantly, both datasets show the consistent life cycle patterns with large differences in wages across job levels.

Figure 1: Wage density across occupations by job level for different occupation codes



Notes: Density estimates for residual wages by occupation and job level. *Within all* shows residual wage density after removing the average wage, *within job levels* removes average job-level wages, and *within occupations* removes average wages by occupation. Wage observations are for occupation-job-level cells. The number of cells varies with the occupation codes applied. See text for further details. For three-digit ISCO codes, we observe 120 different occupations, and for five-digit KldB codes, we observe 1,066 occupations. We always observe 5 job levels.

particular to German data. Data from other countries that include job levels and wage data are discussed in [Strub et al. \(2008\)](#) and [Pierce \(1999\)](#). In Appendix C, we specifically look at U.S. NCS data and corroborate our findings.

Figure 1 zooms in on the relative statistical importance of job levels and occupations in accounting for wage differences. The figure compares wage differences across and within occupations. For this purpose, we aggregate wage data by job-level-occupation cells. Then we either regress these aggregated data on the five job-level dummies or the much finer-grained occupation dummies. Figure 1 shows the distribution of the regression residuals for (a) three-digit ISCO codes (120 categories) as our baseline. In (b), we restrict the underlying sample and calculate occupation-job-level wages based on the wages of 35-45-year-old men, and in (c) we form occupation-job-level cells using the finer five-digit KldB occupation codes (1,066 categories). The legend reports the variance of log wages in the raw data across all job-level-occupation cells (within all), the variance of residuals after controlling for job levels (within job levels), and the variance of residuals after controlling for occupations (within occupations). The results are striking. We find that five job-level dummies account for 36% of the wage dispersion across occupation-job-level cells, while the 120 occupation dummies account for only about 17% of this wage dispersion. Even the 1,066 five-digit occupation dummies account for less of the wage dispersion across occupation-job-level cells (31%) than five job-level dummies. The results when considering only mid-age men lie in between the two occupation specifications.

In Appendix C, we compare wage densities and standard deviations across occupation-job-level cells from U.S. NCS data to the corresponding German 2018 SES data. For both the U.S. and German data, we find that job levels have higher explanatory power than occupations. Furthermore, we report that in the U.S. data, job levels account for half of the within-occupation wage variation. We will return to the empirical differences and similarities of job levels and occupations in Section 4.3, after discussing the empirical concept of job levels next.

4.2 Job levels as an empirical concept


In general, job levels have a long history in labor market statistics. The German statistical office reports in its quarterly wage statistics data on wages by job levels going back at least to 1957, and similar long-term reporting exists for the United States by the BLS. Although the details of job-level assignments vary across schemes, they can be broadly understood as encoding the complexity, autonomy, and responsibility (CAR) inherent in task execution. Defining a “job” in terms of both occupations and job levels provides, therefore, a more comprehensive description than occupations alone. Occupations describe what tasks are performed—they specify the content of the work (e.g. teaching or programming, baking or cooking, painting or building, etc.)—while job levels allow a consistent comparison of how these tasks are executed across and within occupations. Higher job levels are associated with a broader range of tasks that require not only better skills, but also more complex, autonomous, and responsible execution. The complexity dimension reflects the minimum skill requirement needed to perform a job, but higher-skilled workers can still occupy roles with lower formal skill thresholds. [Acemoglu and Autor \(2011\)](#) discuss employers’ allocation problem of workers to tasks of different complexity. Autonomy indicates the degree to which a jobholder can diverge from preset routines and make independent decisions, while responsibility captures the extent to which these decisions affect broader operational outcomes. In short, while occupations capture the *what* of work, job levels capture the *how* in terms of task breadth and difficulty (the CAR intensity).

To be specific, the coding instructions for the five job levels in the SES data can be summarized as follows.¹⁰ At the lowest job level, minimum skill requirements are set so that task execution does not require particular training (such as an apprenticeship) and can be learned on the job in less than three months. Task execution follows clear rules and procedures, and workers do not make decisions independently but follow a clearly defined workflow. The second level also has these workflow characteristics, but task execution is somewhat more complex and requires some experience but no formal training, and can be learned in under two years. The complexity at the third level requires completed occupational training and experience and allows some discretion in the workflow. Junior clerks or salespeople would be typical examples. However, the task execution in these jobs does not include responsibility for the work of others or decisions that affect the work of others, such as strategic business decisions. These responsibilities are a key characteristic of the next two job levels. On the fourth level, task execution requires specialized training, and tasks are executed independently and with discretion over one’s own workflow. Therefore, they come with substantial decision-making power over cases, transactions, or organization of work. Jobholders have some decision-making power in regard to the work of others or their decisions affect the work of others; examples would be production supervisors, junior lawyers, or heads of administrative offices. The highest level includes primarily decision-making tasks and responsibility for others’ work, such as senior lawyers or researchers. However, a high-level job does not necessarily require lower-level workers in the production process.¹¹

¹⁰The BLS job-leveling guide describes in detail the job-leveling approach for the U.S. NCS data ([U.S. Bureau of Labor Statistics, 2013](#)).

¹¹The fact that job levels do not require subordinate hierarchies at the plant distinguishes job levels from theories of production hierarchies as in [Garicano and Rossi-Hansberg \(2006\)](#). The fact that they are linked to tasks and their execution distinguishes them from job titles that are at best vaguely related to task execution, as shown in [Cohen et al. \(2026\)](#).

Figure 2: Share of occupations with different job-level span



Notes: Share of occupations with different levels of job-level span. Job-level span is defined as the number of job levels with at least 5% (10%) of workers from a given occupation. The left panel shows two-digit ISCO codes. The right panel shows five-digit KldB codes (for 2018 SES data). Sample selection applies.

For example, all jobs in research will be classified into the two highest job levels due to their complexity, autonomy, and responsibility. Although educational requirements enter the job level description, education and job levels are distinct, as Appendix B shows.

4.3 Job levels and occupations

Which task is executed by a worker and how this task is executed are, of course, not unrelated. Thus, one should expect some relationship between job levels and occupations, especially if occupational classifications are fine-grained enough, e.g., five-digit occupation codes. Therefore, we next provide a detailed discussion of how occupations and job levels relate.


First, we quantify how much job levels vary within occupations. For this purpose, we calculate for each occupation the share of its workers on the various job levels and then count for each occupation how many job levels hold more than a threshold of 5% (alternatively 10%) of that occupation’s workforce (“job-level span”).

We report the shares of occupations by job-level span in Figure 2 for (a) two-digit and (c) five-digit occupation codes for all workers and (b) for mid-age male workers.¹² We find that most occupations span three job levels. Only if we set the threshold to 10% and use the five-digit KldB codes, we find a marginally higher share of occupations with a job-level span of two. Still, for four out of ten five-digit occupations, we find on three job levels 10% or more workers of that occupation (Figure 2(c), light gray). Thus, there is a clear distinction between job levels and occupations. Similarly, Appendix B shows that education levels span three to four job levels.

At the same time, these findings imply that an occupation does not span all job levels (equally) because not all tasks can be executed at any level of complexity, autonomy, and responsibility. In fact, there is a systematic relationship between the task content of an occupation (what one does) and the distribution of job levels across workers in that occupation (how the task is done). We follow [Dengler et al. \(2014\)](#), who build on the task-based approach ([Autor et al., 2003](#); [Acemoglu and Autor, 2012](#)) and classify occupations into five task groups based on whether

¹²Results for 5-digit codes are based on 2018 SES data alone as these are not included in the SUFs of the SES.

Figure 3: Distributions of job levels by the main task



Notes: The figure displays the distributions of job levels by the main task of a worker’s occupation. The left panel shows results for the entire sample, the right panel restricts the sample to men age 35-45. In both panels, bars show the interquartile range, the black horizontal line shows the mean job level, and the vertical lines indicate the range between the 10th and 90th percentile of the observed job level by main task. Five task components are constructed and used to categorize an occupation as mainly being: non-routine analytic (A-NR), non-routine interactive (I-NR), routine cognitive (C-R), routine manual (M-R), and non-routine manual (M-NR). The main task is the task-based category with the largest task share as defined by [Dengler et al. \(2014\)](#).

their main tasks are analytical, interactive, cognitive, non-routine manual, or routine manual.¹³ Figure 3 shows that on average the jobs in those occupations that mainly execute analytical non-routine tasks are the most CAR intensive, i.e., they have the highest average job level. In panel (a), we consider the entire sample and in panel (b), we restrict the sample to mid-age men as a robustness check. Jobs in occupations with mainly manual routine tasks are the least CAR intensive. Cognitive routine, interactive non-routine, and manual non-routine are in between. However, in line with Figure 2, there is substantial heterogeneity, even conditional on the main task type. In fact, conditional on the main task, we find again that a task group typically spans three job levels. In Appendix D, we discuss that the CAR intensity of occupational tasks is the main explanatory variable of task-based occupational wage differences. These results highlight that job levels provide an additional dimension that allows one to refine the task-based approach to job descriptions and wage determination.

4.4 Job levels and job-leveling factors (CAR intensity)

Respondents in the SES receive instructions on how to assign job levels to jobs based on their CAR intensity (Section 4.2). Each element, complexity, autonomy, and responsibility, plays a role in this assignment; wages do not. However, one might be concerned about reverse causality in the form of respondents not following the survey instructions and providing job-level information based on wage information.¹⁴ To address this concern, we use additional, independent

¹³This classification is based on 3-digit occupation codes. [Spitz-Oener \(2006\)](#) uses an alternative based on BIBB/BAuA survey responses from workers. In the German 5-digit occupation coding (KldB2010) the fifth digit refers to the complexity of tasks, obviously a concept related to job levels. Appendix D.2 provides more details on the joint distribution of job levels and the fifth digit of the occupation that captures complexity. We find that these are correlated but far from identical.

¹⁴This point is important with respect to theories of tournament models of career progression, such as those pioneered in [Lazear and Rosen \(1981\)](#). There, wages and promotions are a reflection of past success and ultimately

Table 3: Wage regressions for white-collar workers (Angestellte)

controls	adj. R^2
job-leveling factors	0.441
+ occupations	0.486
+ employer characteristics + region	0.612
occupation + employer characteristics (w/o job levels)	0.517

Notes: Adjusted R^2 from different regressions of log wages on different sets of observables (see text for details). The regression sample always contains 3,027 observations for white-collar workers.

survey data containing task execution descriptions to construct job-leveling factors based on CAR intensity. This job-leveling yields a similarly high explanatory power for wages. Naturally, we perform this job-leveling without considering wage data. The BIBB/BAuA data provide additional details on task execution that are not included in the occupational data. It also provides wage data, worker demographics, industry, and employer size. From the task execution data, we select the information that is used for job leveling. We refer to this information as job leveling factors. We show that these job-leveling factors in the BIBB/BAuA data have the same explanatory power for the wage data as the coded job levels in the SES data. In other words, we show that job levels have economic content and can be constructed from task execution information. We restrict our analysis to white-collar workers, but report the results for blue-collar workers in Appendix E.2.

As job-leveling factors, we select eight survey questions that we identify to be informative about a job’s CAR intensity. We encode answers to these questions using dummy variables and refer to them as job-leveling factors. The selection of questions is based on the job-leveling scheme from the German steel- and metalworker bargaining agreement.¹⁵ We report the detailed survey questions in Appendix E.1. First, we explore the explanatory power of job-leveling factors by running a series of linear wage regressions. Table 3 reports the R^2 from these regressions. Job-leveling factors alone account for 44% of the wage variation.¹⁶ This high explanatory power aligns closely with our results for the SES data. In the SES data, job levels alone account for 47% of the overall wage variation. Adding occupation information to the job-leveling factors increases the explanatory power only slightly to 49%.¹⁷ This aligns well with our findings in Table 2. If we further add employer characteristics, we account for 61% of the wage variation. In the SES data, the corresponding regression on plant characteristics and job levels accounts for 62% of the wage variation. We conclude that also in the detailed BIBB/BAuA data there is the same strong relationship between CAR intensity constructed directly from information on task execution and wages as in the SES data, where CAR intensity is summarized by job levels.


unrelated to the tasks of the current job.

¹⁵The questions summarize the complexity of and skills required for the job, the autonomy in organizing workflow, the degree of communication, and whether the job involves supervisory duties. Importantly, none of the information is on worker characteristics such as age or highest degree of education.

¹⁶The regression involves 18 dummies for answers to the eight questions and a constant.

¹⁷A regression on occupations alone accounts for only 34% of the wage variation.

Figure 4: Wages by job-level points



Notes: *Left*: Scatter plot of a worker's implied job-level points from statistical job-leveling scheme against the worker's job-level points from union bargaining scheme (ERA scheme, white-collar workers). The statistical job-level scheme is based on the regression of wages on survey answers. The solid line shows the linear fit and the legend reports R^2 . *Right*: Distribution of wages by job level (groups of 5 points to reduce sampling noise). Job-level points have been constructed from survey questions on job characteristics (see text for details).

As a second, more constrained exercise, we assign workers to a job-level scale in the BIBB/BAuA data using the information from the job-leveling factors. We do so based on a job-leveling scheme from an existing union wage agreement.¹⁸ Whereas one can understand our first exercise as assigning point values to the job-leveling factors to maximize the explanatory power for (log) wages, we now take point values from a job-leveling scheme of an existing union wage contract. Concretely, we use the *ERA scheme* from the German steel- and metalworker bargaining agreement (*ERA-Punktebewertungsbogen zur Bewertung von Arbeitsaufgaben*), which is typically seen as the reference bargaining agreement in Germany. We describe our mapping of survey answers to the job-leveling scheme in Appendix E.1 and test this mapping using a specific collective bargaining agreement as a case study in Appendix E.3. In Figure 4(a), we first compare the job-level points assigned to each worker based on the ERA scheme to the implied job-level points from the regression of log wages on job-leveling factors. We derive implied job-level points from the job-leveling factors using predicted wages from the regression and standardizing them to have the same mean and standard deviation as the points based on the ERA scheme. We find that the two job-level point schemes align very closely. A linear regression yields a R^2 of 0.86. Given this close alignment, we also get that job-level points from the ERA scheme account for 39% of the wage variation in a regression with log wages as the dependent variable, and, hence, only slightly less than the 44% from the more flexible regression on job-leveling factors in Table 3. Figure 4(b) shows the distribution of wages by assigned job-level points. The mean wage is increasing in the job-level points and the dispersion is roughly constant around the mean over the entire point range. Although there is dispersion conditional on job-level points, the data

¹⁸ Ellguth and Kohaut (2021) report that in 2014 about half (47%) of West German private-sector employees were covered by union bargaining contracts, for East Germany they report a share of about one-third (28%). Importantly, they also report that employers who do not pay according to a union wage still align their wages to existing union wage contracts. In 2020, 40% of employers without union bargaining agreement reported such an approach. Union wage contracts are very transparent in how pay is assigned to jobs and are a prime example of a job-leveling scheme.

show a clear positive relationship between job-level points and (log) wages.¹⁹

Two points are important to emphasize regarding these results. First, the coding of job-level points involves only eight survey questions regarding complexity, autonomy, and responsibility (CAR). Second, neither information about worker characteristics nor wages has been used for assigning points to jobs.²⁰ These two points address the important question of reverse causality that job levels could be just a recoding of wages (e.g., wage quintiles).

4.5 Job levels in theory

To offer an interpretation to our empirical findings from an economic theory perspective, we formalize the concept of job levels within a task-based framework, in which the job-level structure emerges endogenously and rationalizes key empirical properties. First, higher job levels entail higher minimum skill requirements, reflecting greater complexity. Second, they encompass a broader set of tasks, which necessitates increased autonomy. Third, they generate higher marginal products for appropriately qualified workers.

4.5.1 Setup

As in [Freund \(2022\)](#), we consider the planning problem of a firm that needs to allocate its workforce to a range of tasks. For tractability, we assume that the firm (plans as if it) has a continuum of workers of measure one and maximizes output. Total output, Y , is produced according to the CES production function in tasks $\tau \in [0, 1]$:

$$Y = \left\{ \int_0^1 \left[y(\tau) \right]^{\frac{\eta-1}{\eta}} d\tau \right\}^{\frac{\eta}{\eta-1}}. \quad (1)$$

To produce output $y(\tau)$ of task τ the firm assigns its continuum of workers indexed by i that differ in their task-specific productivity $z_i(\tau)$. The output of a task is the productivity weighted number of hours of all workers spent on task τ

$$y(\tau) = \int z_i(\tau) l_i(\tau) di. \quad (2)$$

A firm cannot perfectly (without costs) observe the productivity of worker i , but the firm knows that productivity on task τ is given by $z_i(\tau) = z(\tau, \alpha_i) = 1 - \tau(1 - \alpha_i)$, where $\alpha_i \in [0, 1]$ characterizes the (unobserved) skills of worker i . A firm can design assessment tests of the worker's skills, e.g., through interviews, that partition the skill space, $\{[a_1, a_2), [a_2, a_3), \dots, [a_M, a_{M+1})\}$, with $a_1 = 0$ and $a_{M+1} = 1$. These allow the firm to measure the skill partition in which a worker falls, but we assume that a finer filtration of skills creates increasing costs of assessment

¹⁹It has to be taken into account that the wage data come from a survey so we also expect substantial measurement error on wages that accounts for some of the dispersion conditional on job-level points.

²⁰An additional point is also worthwhile to re-iterate: the assignment of job-level points is based on our reading of one specific job-leveling scheme. This makes it clear why the predictive power is already lower compared to the dummy regression. In Appendix E.1, Figure A8, we demonstrate, however, that our job leveling successfully recovers the bargained union wages, except for jobs at very low levels, where there is strong compression in union wages. The fact that job levels can be derived independently of the wage structure has also been shown in case studies ([Dohmen et al., 2004](#)).

such that the firm only distinguishes M levels of skills. In the following, we denote skill group j as the group that corresponds to the skill set $[a_j, a_{j+1})$. We assume that in the population the skills of workers a are distributed according to the distribution function $F(a)$ with density $f(a)$. By grouping workers into skill categories, firms cannot directly condition their decisions on a worker's identity i , but only on the corresponding group j . Consequently, worker i 's work schedule $l_i(\tau)$ is the schedule of her skill group $l_j(\tau)$. This allows us to define a “job”:

Definition 1 (Job). *A job (of level) j is a skill requirement for a worker, such that the worker's productivity is in the interval $\alpha \in [a_j, a_{j+1})$, and an assignment of (positive) hours to tasks under a work schedule $l_j(\tau) \geq 0$, such that $\int l_j(\tau) d\tau = 1$.*

The minimum skill requirement can be viewed as a description of the complexity of the job. The firm requires that the worker's assessed skills exceed a minimum level of productivity given the “complexity” of the tasks associated with the job. The “autonomy” of a job is, of course, related to the range of tasks in a job's work schedule. However, jobs differ not only in task scope, but also in the mass of workers working on the assigned tasks $\Delta F_j := F(a_{j+1}) - F(a_j)$. With a continuum of tasks, the firm is indifferent to who exactly performs which task within a skill group (as they are homogeneous from the firm's perspective). Thus, the exact set of tasks assigned to an individual worker remains indeterminate in the model. To resolve this indeterminacy, we assume that the firm assigns $\mathcal{A}_j := \frac{\mu(\mathcal{T}_j)}{\Delta F_j}$ “tasks per worker” within the job level j , where $\mu(\mathcal{T}_j)$ denotes the mass of tasks in $\mathcal{T}_j$. We refer to $\mathcal{A}_j$ as the autonomy of the job. Intuitively, one can think of this as minimizing “number” of tasks of a representative worker.²¹

4.5.2 Optimal work schedule

First, we discuss how firms with M job levels determine the autonomy of a job with a given complexity (minimum skill requirement), i.e., we fix $\{a_j\}_{j=1}^{M+1}$. The central element is to characterize the optimal work schedule in terms of the total number of hours spent by all workers on each task $l_j(\tau)\Delta F_j$ at job level j . This planning problem is given by

$$\max_{\{l_j(\tau)\}_{j=1}^M} \left\{ \int_0^1 \left[\sum_{j=1}^M \int_{a_j}^{a_{j+1}} [1 - \tau(1 - \alpha)] l_j(\tau) f(\alpha) d\alpha \right]^{\frac{\eta-1}{\eta}} d\tau \right\}^{\frac{\eta}{\eta-1}} \quad (3)$$

$$\text{s.t. } \int l_j(\tau) d\tau \leq 1 \wedge l_j(\tau) \geq 0 \quad (4)$$

Given the distribution of worker skills, a firm knows for each job the average skills of the representative worker $\bar{a}_j$ of group j . Since hours depend only on the job, not on the actual identity of the worker, the firm's planning problem simplifies to:

$$\max_{\{l_j(\tau)\}_{j=1}^M} \left\{ \int_0^1 \left[\sum_{j=1}^M \underbrace{[1 - \tau(1 - \bar{a}_j)]}_{\text{av. output per hour of task } \tau \text{ in job } j} \underbrace{\Delta F_j}_{\text{share of workers in job } j} \underbrace{l_j(\tau)}_{\text{number of hours in task } \tau \text{ on job } j} \right]^{\frac{\eta-1}{\eta}} d\tau \right\}^{\frac{\eta}{\eta-1}}, \quad (5)$$

²¹Strictly speaking, the firm minimizes the maximal ratio of the mass of tasks to a mass of workers in any measurable sample of workers within group j . The minimum is then trivially achieved by assigning a constant density ratio $\mathcal{A}_j$.

which has the first-order conditions:

$$l_j(\tau) : \quad \left\{ \int_0^1 \left[\sum_{m=1}^M [1 - t(1 - \bar{a}_m)] \Delta F_m l_m(t) \right]^{\frac{\eta-1}{\eta}} dt \right\}^{\frac{1}{\eta-1}} \quad (6a)$$

$$\times \left[\sum_{m=1}^M [1 - \tau(1 - \bar{a}_m)] \Delta F_m l_m(\tau) \right]^{\frac{-1}{\eta}} \quad (6b)$$

$$\times [1 - \tau(1 - \bar{a}_j)] \Delta F_j = \lambda_j + \phi_j(\tau). \quad (6c)$$

Here λ_j is the Lagrangian multiplier on the total hours constraint and $\phi_j(\tau)$ is the multiplier on the non-negativity constraint of hours worker group j spends on task τ . The first part (6a) reflects the marginal value of output, (6b) is the marginal product of task τ and (the left-hand side of) (6c) reflects the marginal productivity of a level j worker on task τ .

Lemma 1 (All tasks). *The optimal job design problem assigns positive hours to all tasks.*

Proof. This follows trivially from the Inada condition. $\square$

This implies, together with the first-order condition for hours, that jobs partition the task space.

Proposition 1 (Jobs partition the task space). *No two jobs have positive hours simultaneously on more than one task such that jobs partition the task space. This means that all $\mathcal{T}_j := \{\tau | l_j(\tau) > 0\}$ are intervals $[\tau_j, \tau_{j+1}]$ with increasing cutoffs $\tau_j < \tau_{j+1}$ and $\tau_1 = 0, \tau_{M+1} = 1$.*

Proof. Assume to the contrary that jobs j and k have strictly positive hours in two distinct tasks τ_1 and τ_2 . Then, $\phi_j(\tau_1) = \phi_j(\tau_2) = \phi_k(\tau_1) = \phi_k(\tau_2) = 0$. And from the first-order conditions

$$1 - \tau(1 - \bar{a}_j) = \frac{\lambda_j / \Delta F_j}{\lambda_k / \Delta F_k} [1 - \tau(1 - \bar{a}_k)] \quad (7)$$

needs to hold for both τ_j and τ_k . For a linear equation in τ to have more than one solution, both sides have to be identical. This would require, in particular, $\frac{\lambda_j / \Delta F_j}{\lambda_k / \Delta F_k} = 1$ and $(\bar{a}_j - 1) = \frac{\lambda_j / \Delta F_j}{\lambda_k / \Delta F_k} (\bar{a}_k - 1)$, which in turn implies $\bar{a}_j = \bar{a}_k$, contradicting the assumption of $j \neq k$.

As all tasks are produced (Lemma 1), it must be that (7) holds exactly for one τ_j for neighboring jobs j and $k = j + 1$, so it follows that $\mathcal{T}_j$ are intervals which partition $[0, 1]$. The fact that the cutoffs τ_j are increasing follows directly from more skilled workers having an absolute advantage in all tasks and a comparative advantage in more difficult tasks. $\square$

Note that this proposition means that a job can be equivalently described as a skill requirement, $[a_j, a_{j+1}]$, as a range of tasks $[\tau_j, \tau_{j+1}]$, and third as a work schedule $l_j(\tau)$. Both the skills and the tasks give rise to a hierarchical ordering of jobs as tasks of jobs do not overlap.

Building on Proposition 1, the optimal hours choice for any given partition is given by

Lemma 2 (Optimal Hours Allocation). *Assume any partitioning of the task space in intervals $\{T_j\}_{j=1 \dots M}$ with $\cup_j T_j = [0, 1]$ such that only job j supplies positive hours on $\tau \in T_j$ and zero*

hours outside T_j . Then the optimal work schedule for all $\tau \in T_j$ is given by

$$l_j(\tau) = \left[\frac{1 - \tau(1 - \bar{a}_j)}{S_j} \right]^{\eta-1} \quad \text{with} \quad S_j := \left\{ \int_{T_j} [1 - t(1 - \bar{a}_j)]^{\eta-1} dt \right\}^{\frac{1}{\eta-1}}, \quad (8)$$

where S_j is the expected productivity over tasks assigned to a worker in group j .

Proof. Given a partitioning the objective function can be rewritten as

$$Y = \left\{ \sum_{j=1}^M \int_{T_j} \left[\sum_{k=1}^M (1 - \tau(1 - \bar{a}_k)) \Delta F_k l_k(\tau) \right]^{\frac{\eta-1}{\eta}} d\tau \right\}^{\frac{\eta}{\eta-1}} \quad (9)$$

which simplifies to

$$Y = \left\{ \sum_{j=1}^M \int_{T_j} [(1 - \tau(1 - \bar{a}_j)) \Delta F_j l_j(\tau)]^{\frac{\eta-1}{\eta}} d\tau \right\}^{\frac{\eta}{\eta-1}} \quad (10)$$

because of the assumption $\forall \tau \notin T_j : l_j(\tau) = 0$. The maximizer $l_j(\tau)$ of Y subject to (4) is, therefore, characterized by the first-order condition

$$l_j(\tau) : \quad \left\{ \sum_{m=1}^M \int_{T_m} [(1 - t(1 - \bar{a}_m)) \Delta F_m l_m(t)]^{\frac{\eta-1}{\eta}} dt \right\}^{\frac{1}{\eta-1}} \quad (11a)$$

$$\times [1 - \tau(1 - \bar{a}_j)]^{\frac{\eta-1}{\eta}} (\Delta F_j)^{\frac{\eta-1}{\eta}} l_j(\tau)^{\frac{-1}{\eta}} = \lambda_j \quad (11b)$$

The right-hand side of the first-order condition is independent of a specific task. This implies that any task served by job j has to have, in terms of the final good, the same marginal product (of an hour worked taken into account worker productivity). As only one job serves one task, we can use the FOC to back out the optimal work schedule $l_j(\tau)$ of job j , expressing the first-order condition relative to the one for a reference task $\hat{\tau}_j \in T_j$ in job j :

$$l_j(\tau) = \left[\frac{1 - \tau(1 - \bar{a}_j)}{1 - \hat{\tau}_j(1 - \bar{a}_j)} \right]^{\eta-1} l_j(\hat{\tau}_j). \quad (12)$$

This yields the desired result by using the total labor supply constraint $\int_{T_j} l_j(\tau) = 1$. $\square$

Given the optimal hours choice for given τ_j , it is straightforward to see that total output can be written in a compact form as a function of task cutoffs and skill requirements

$$Y(\{\tau_1, \dots, \tau_{M+1}\}, \{a_1, \dots, a_{M+1}\}) = \left(\sum_{j=1}^M (S_j \Delta F_j)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}. \quad (13)$$

The representative worker on job j produces a bundle of tasks aggregated to S_j , which the firm then aggregates weighted by the employment shares ΔF_j .

The optimal job design problem then concludes with the firm maximizing (13) by choosing task cutoffs τ_j (with the boundary conditions $\tau_1 = 0, \tau_{M+1} = 1$). The first-order conditions

characterizing τ_j (and thus the optimal $\mathcal{T}_j := [\tau_j, \tau_{j+1}]$) are

$$0 = \frac{Y^{\frac{1}{\eta}}}{\eta - 1} \left\{ (S_{j-1} \Delta F_{j-1})^{\frac{\eta-1}{\eta}} \left[\frac{1 - \tau_j(1 - \bar{a}_{j-1})}{S_{j-1}} \right]^{\eta-1} - (S_j \Delta F_j)^{\frac{\eta-1}{\eta}} \left[\frac{1 - \tau_j(1 - \bar{a}_j)}{S_j} \right]^{\eta-1} \right\}. \quad (14)$$

4.5.3 Optimal skill requirements

Using the compact form of (13), it is also straightforward to describe the optimal skill requirement in terms of first-order conditions for a_j (with the boundary conditions $a_1 = 0, a_{M+1} = 1$),

$$0 = Y^{\frac{1}{\eta}} \left\{ (S_{j-1} \Delta F_{j-1})^{-\frac{1}{\eta}} \left[\frac{\partial S_{j-1}}{\partial a_j} \Delta F_{j-1} + S_{j-1} f(a_j) \right] + (S_j \Delta F_j)^{-\frac{1}{\eta}} \left[\frac{\partial S_j}{\partial a_j} \Delta F_j - S_j f(a_j) \right] \right\}. \quad (15)$$

The two sets of first-order conditions ((14) and (15)) can now be further simplified in four steps. First, we use the fact that the Lagrangian multipliers λ_j are proportional to $(S_j \Delta F_j)^{\frac{\eta-1}{\eta}}$. Second, for (14), we use the optimal labor input schedule $l_j(\tau)$ from Lemma 2. Third, for (15), we use that $\frac{\partial S_{j-1}}{\partial a_j} = \frac{\partial S_{j-1}}{\partial \bar{a}_{j-1}} \frac{\partial \bar{a}_{j-1}}{\partial a_j}$ and $\frac{\partial S_j}{\partial a_j} = \frac{\partial S_{j-1}}{\partial \bar{a}_j} \frac{\partial \bar{a}_j}{\partial a_j}$, together with $\frac{\partial \bar{a}_{j-1}}{\partial a_j} = (a_j - \bar{a}_{j-1}) \frac{f(a_j)}{\Delta F_{j-1}}$, and $\frac{\partial \bar{a}_j}{\partial a_j} = (\bar{a}_j - a_j) \frac{f(a_j)}{\Delta F_j}$. Finally, we define $\epsilon_j := \frac{\partial S_j}{\partial \bar{a}_j} S_j^{-1}$. This yields the more compact form

$$\lambda_{j-1} l_{j-1}(\tau_j) = \lambda_j l_j(\tau_j) \quad (16)$$

$$\lambda_{j-1} \frac{1 + \epsilon_{j-1}(a_j - \bar{a}_{j-1})}{\Delta F_{j-1}} = \lambda_j \frac{1 - \epsilon_j(\bar{a}_j - a_j)}{\Delta F_j}. \quad (17)$$


Under the optimal allocation, higher-skilled workers will work on harder tasks (Proposition 1) and will have a higher expected product:

Proposition 2. *The expected product of a job is increasing in the associated skill requirement.*

Proof. The average product of a worker in skill group/job j , w_j , can be written as $w_j = p_j S_j$ the product of the quasi-price $p_j \propto (S_j \Delta F_j)^{-\frac{1}{\eta}}$ of the j -th task groups and the expected output of this worker on her tasks S_j . Define $\mathcal{S}_{k,j} := \left\{ \int_{\mathcal{T}_k} [1 - t(1 - \bar{a}_j)]^{\eta-1} dt \right\}^{\frac{1}{\eta-1}}$ as the output in task group k that a worker of skill group j would have as her expected output if we allowed her to optimize her hours spent on that task group. In an optimal assignment of workers, we must have $\forall k, j : p_j S_j > p_k \mathcal{S}_{k,j}$. For $j > k$, we get $\mathcal{S}_{k,j} > \mathcal{S}_k$ because the worker in the higher skill group has an absolute advantage on all tasks, and hence $w_j = p_j S_j > p_k \mathcal{S}_k = w_k$ follows. $\square$

To show further the properties of the job design problem, we solve the problem for $M = 5$ job levels, a flexible class of distribution functions, $F(\alpha)$, and various degrees of substitutability, η , in the production process. We consider η in an interval from 0.05 to 10 going from very low to very high substitutability of tasks. For the skill distribution, $F(\alpha)$, we assume a Beta distribution and present two concrete alternatives. First, we consider a uniform distribution on $[0, 1]$ so that $F(\alpha) = \alpha$ and second, we consider parameters $a = b = 5$, which is relatively similar to a normal distribution in shape but bound to $[0, 1]$. We calculate the optimal skill $\{a_j\}_{j=1}^{M+1}$ and task cutoffs $\{\tau_j\}_{j=1}^{M+1}$ from the first-order conditions in (16) and (17). Job levels are ordered by their complexity from job level 1 being the least complex (lowest skill requirement).

Figure 5: Results from the Job Design Problem for two alternative Skill Distributions $F(\alpha)$



Notes: Change in productivity with skills, optimal task range $[\tau_j, \tau_{j+1}]$, and optimal group size ΔF from job design problem with different distributional assumptions for $F(\alpha)$. The first row corresponds to a Beta distribution with $\alpha = \beta = 1$ (uniform distribution), while the second row corresponds to $\alpha = \beta = 5$ (approximating a normal distribution). Columns represent different variables: the change in productivity with skills, optimal task range, and optimal group size, respectively. The vertical axis reflects the degree of substitutability of tasks in production.

Figure 5 reports three key implications of the firm's planning problem. The first finding (Figures 5(a) and 5(d)) strengthens Proposition 2 insofar as expected productivity not only increases in the average skills employed in a job level, but that the relationship is convex: the expected product increases more relative to the average skills the higher the job level, $\frac{w_j - w_{j-1}}{a_j - a_{j-1}} < \frac{w_{j+1} - w_j}{a_{j+1} - a_j}$. The second finding is (Figures 5(b) and 5(e)) that autonomy $\mathcal{A}_j$ is larger at higher job levels except for cases of low levels of substitutability. Third, we find that the size of the top job level group (Figures 5(c) and 5(f)) tends to be the smallest and its size decreases in substitutability η and vice versa, the lowest job level tends to be the largest group, increasing in η . However, when we compare the different worker skill distributions, we find that the supply of workers shapes the job design and that with fewer low-skilled workers the firm can have non-pyramid-shaped organizational structures (different to [Garicano and Rossi-Hansberg, 2004](#)).

In summary, this theoretical framework rationalizes the existence of job levels by an information friction, where assessing worker skills precisely is prohibitively expensive. The result of this friction is that firms assess worker types only within ranges and group workers. For each group of workers, they design a job that has a specific job level.²² The model then predicts that, equating wages to expected productivity, (1) all workers of a job level earn the same wage, (2) wages are convex in underlying skills, (3) job levels are characterized by a positive association of complexity (minimum skill requirements a_j) and autonomy (the range of tasks per worker $\mathcal{A}_j$).

²²This setup abstracts from occupations. An interesting extension would be to consider occupations that, at the same general difficulty level, cover a range of specific tasks, giving rise to a trade-off between general and specific skills (and their accumulation).

5 Job levels as mediator in wage dynamics

In the next step, we embed the above static assignment problem in a dynamic environment to derive implications for life-cycle wage dynamics. We confront these implications with the data and show that the data confirm key predictions of the model. Our novel job-level perspective provides a new interpretation of several widely documented between-worker wage differences. Age, gender, education, and seniority are all mediated to wages through job levels.

5.1 Job levels and career dynamics in theory

We start from the solution to the static job design problem in Section 4. In order to maintain analytical tractability, we assume that even in a dynamic setting and with a finite number of workers, the firm relies on the optimal solution for minimum skill requirements and task ranges for a continuum of workers. One interpretation of this assumption is that job design is an irreversible, one-time technology choice. Given this assumption, we show that when faced with a small-sample skill distribution, firms typically assign a worker without matching skills to a job level rather than changing the fraction of workers at a particular job level ΔF_j . However, if workers are paid according to their current job level, firms will leave a job level j vacant and continue to search if the only available worker i is underqualified ($\alpha_i < a_j$). Firms only hire an appropriately qualified or overqualified worker i to a job level j ($\alpha_i \geq a_j$).

5.1.1 Worker allocation with a fixed job design

Following our assumption, a firm makes its irreversible job design decision, expecting to draw a continuum of workers according to distribution F . In any small sample of workers, the realized fractions will no longer be deterministic. We consider the case of a realized fraction $\Delta F_j + \delta$ of workers in skill group j and $\Delta F_k - \delta$ in group k . In this case, the firm must decide whether to employ all additional workers δ on the job at level j , which matches their skills, or to shift them to the understaffed job at level k . Allowing for re-optimization of hours for “mismatched” workers, we obtain the output of a worker of type j on job k as (see the proof of Proposition 2)

$$\mathcal{S}_{k,j} = \left\{ \int_{\mathcal{T}_k} [1 - \tau(1 - \bar{a}_j)]^{\eta-1} d\tau \right\}^{\frac{1}{\eta-1}}. \quad (18)$$


The firm will accept a deviation from the planned worker allocation, if and only if

$$(\mathcal{S}_{j,j}(\Delta F_j + \delta))^{-\frac{1}{\eta}} \mathcal{S}_{j,j} \geq (\mathcal{S}_{k,k}(\Delta F_k - \delta))^{-\frac{1}{\eta}} \mathcal{S}_{k,j}, \quad (19)$$

which compares the value-weighted output of the matched job-skill assignment (j, j) to the unmatched one (k, j) . This inequality can be expressed equivalently as a comparison of population share deviations to value-weighted output ratios under the optimal allocation:

$$\frac{\frac{\Delta F_k - \delta}{\Delta F_k}}{\frac{\Delta F_j + \delta}{\Delta F_j}} \leq \left(\frac{\mathcal{S}_{k,j}}{\mathcal{S}_{j,j}} \right)^\eta \frac{\mathcal{S}_{j,j}}{\mathcal{S}_{k,k}} \frac{\Delta F_j}{\Delta F_k} = \left(\frac{\mathcal{S}_{k,j} p_k}{\mathcal{S}_{j,j} p_j} \right)^\eta =: \mathcal{C}_{k,j}, \quad (20)$$

Figure 6: Maximum acceptable deviation of workers per job level in percent of the workforce



Notes: Maximum acceptable deviation δ of workers per job level in percent of the workforce. Vertical axis is the tolerable increase in employment (in percent of total) at job level j at the expense of job level $j - 1$ before a firm reallocates a worker of skill group j to job level $j - 1$. The left panel shows results for a beta distribution with $\alpha = \beta = 1$ (uniform), while the right panel shows results for $\alpha = \beta = 5$ (approximately normal).

where $p_k := (\mathcal{S}_{kk} \Delta F_k)^{-\frac{1}{\eta}}$ is the quasi-price of task k in the optimal allocation, see the proof of Proposition 2. If the inequality holds as a strict inequality and in the other direction, the firm has an incentive to reallocate a marginal worker to the understaffed job. The same right-hand-side criterion applies to reallocating from a fully-staffed job level to an understaffed one.²³


In Figure 6, we plot the maximum acceptable deviation $\bar{\delta}_{k,j} = \Delta F_k \frac{1 - C_{k,j}}{1 + C_{k,j} \frac{\Delta F_k}{\Delta F_j}}$, assuming the firm meets too many workers with level j -skills at the expense of too few workers with level $j - 1$ -skills ($k = j - 1$). If the firm meets an underqualified workforce, the critical values are very similar. We find that for most levels of substitutability, a firm with as few as 50 workers (97% of all German firms are smaller) would not accept even a one-person deviation from the expected number of workers per job level, $\bar{\delta}_{k,j} \leq 2\%$. Therefore, we will assume that, as in the deterministic case, a firm always maintains an optimal number of workers at each job level. This behavior then creates a discrepancy between a worker's skills and her job (level).

5.1.2 Implications for worker productivity

In this situation, the marginal product of a worker ($p_k S_{k,j}$) depends on her skills and, importantly, also on her job. Figure 7 gives two examples of how a worker's marginal product depends on her skills and her job for settings where tasks are complements ($\eta = 0.5$) or substitutes ($\eta = 10.0$). At job level one, the effect of skills on the marginal product is negligible. For the case of complementarity ($\eta = 0.5$), moving skills from the lowest level to the highest level increases the marginal product by only 10%. For the case of substitutability ($\eta = 10$), this increase is only 3%. In contrast, moving a worker with skill level five across job levels from lowest to highest increases the marginal product by 46% (22%). From an individual's perspective, there would be significant benefits to moving to a job level that matches her skills. However,

²³This may be optimal if a cascade of reallocating workers is better than direct reallocation. For example, if there are too few workers in level one and too many in level five, it is usually better to move workers in level five to four, workers in level four to three, and so on, minimizing the mismatch along the way.

Figure 7: Marginal products of workers in different skill groups across jobs



Notes: The figure displays the marginal product of a worker for all job-level k and skill level j combinations $p_k S_{k,j}$. The distribution of worker skills against which the firm has optimized its job design is assumed to be uniform. In a dynamic setting, only combinations highlighted by bold edges will be observed, see Section 5.1.3 (Proposition 3) for details. *Left panel:* Task are complements, $\eta = 0.5$. *Right panel:* Task are substitutes, $\eta = 10.0$.

as we have seen, the firm trades off these gains against the losses from over- or underproducing a set of tasks.²⁴ Therefore, once we allow for a non-deterministic distribution of worker skills ($\exists : \delta > 0$), we obtain that skill requirements and actual skills of workers in a job do no longer need to be identical and a concept of *human capital utilization* emerges when workers of the same skill type work on different job levels.

5.1.3 Hiring decisions with a fixed job design

In a dynamic context, firms do not face a constant small-sample distribution of skills in their workforce. Rather, workers leave firms and firms hire new workers. To study this situation, we simplify further and consider a case in which there is exactly one vacancy at job level k . In such a situation, a firm will need to decide whether or not to hire a worker even if the assessment test indicates that a specific worker's skills do not fall within the skill band of an open position $[a_k, a_{k+1}]$. The worker can be under- or overqualified for the vacancy. If the assessment test shows that the worker is overqualified, the firm is trivially willing to hire the worker as long as it has to pay at most the marginal product (Figure 7). For example, if the firm pays a wage equal to the marginal product of an ideally matched worker, then it actually makes a profit by hiring an overqualified worker. Hence, the interesting case is when a firm faces an underqualified worker at the hiring stage. We formalize these maintained assumptions as follows:

Assumption 1. *The firm has a fixed ex-ante job design (task ranges and skill groups).*

Assumption 2. *The firm has N positions, $N - 1$ are filled. The open position is at level k . Workers on the $N - 1$ filled position will never leave the firm.*

Assumption 3. *Workers are remunerated based on their job level. Higher job levels pay better.*

²⁴ Assigning a worker to her “ideal” job level would increase the output of that job level $S_k(\Delta F_k + \delta)$ to an extent that the marginal product of that job, p_k , declines so much that this outweighs the relative worker productivity differences $\mathcal{S}_{k,j}$, see (19).

Assumption 4. Any newly hired worker leaves the firm again with probability $\nu + \chi_{k,j}$, where ν is an exogenous destruction rate and $\chi_{k,j}$ reflects the probability of a skill- j worker on job level k to find a better paying job. With probability λ the firm meets a worker with skills of level k .

Proposition 3. Under Assumptions 1 - 4, there exists an equilibrium in which firms will never fill a job with a strictly underqualified worker ($\alpha_i < a_j$) when jobs are sufficiently stable, i.e., when ν is small, and the firm's discount factor is sufficiently close to one.

Proof. First, if all other firms do not hire underqualified workers, there is no wage increase that an underqualified worker could achieve by meeting another firm because of Assumption 3. Hence, from Assumption 4 we know that $\chi_{k,j} = 0$ if $k \geq j$. Let $\mathcal{V}$ be the value of the firm with the level k position unfilled (Assumption 2). Let $\tilde{V}_{k,j}$ be the value of the firm after hiring a worker of skill group j for the vacant position at level k . Let $\underline{V}_k$ the value of the firm that maintains the level- k position unfilled until it meets a weakly overqualified worker. Let $\tilde{Y}_{k,j}$ and $\underline{Y}_k$ be the corresponding output. Finally, denote with W the wage sum over all workers when all positions are filled. We then have

$$\tilde{V}_{k,j} = \tilde{Y}_{k,j} - W + \beta \left(\nu \mathcal{V} + (1 - \nu) \tilde{V}_{k,j} \right) \quad \forall j \leq k \quad (21)$$

$$\underline{V}_k \geq \underline{Y}_k - (W - w_k) + \beta \left(\lambda \tilde{V}_{k,k} + (1 - \lambda) \underline{V}_k \right) \quad (22)$$

where (21) follows from the fact that no weakly underqualified worker will ever leave the firm for another employer ($\chi_{k,j} = 0$ if $k \geq j$). Inequality (22) reflects that looking for any overqualified worker has to be a better strategy than looking for an exactly qualified worker.

The firm will never hire an underqualified worker if for all $k > j : \underline{V}_k > \tilde{V}_{k,j}$. Taking the limit of $\beta \rightarrow 1$, this condition can be rewritten such as to obtain

$$\underbrace{\nu^{-1}}_{\text{expected match duration}} \underbrace{(\tilde{Y}_{k,k} - \tilde{Y}_{k,j})}_{\text{output loss from underskilled workforce}} > \underbrace{\lambda^{-1}}_{\text{search duration}} \underbrace{[(W - w_k) - \underline{Y}_k]}_{\text{wage bill net of output}} \quad (23)$$

as a sufficient condition. The left-hand side shows the difference in output between a firm with an ideally skilled workforce and one with an underskilled workforce, weighted by the expected duration of a match. Since $\tilde{Y}_{k,j}$ increases in j , the left-hand side is always positive. The right-hand side is the expected shortfall of output relative to the wage bill while searching for the ideal match (λ^{-1} being the expected search duration) and finite for any fixed $\lambda > 0$. Thus, there exists a ν sufficiently small such that the inequality always holds and the proposition follows. $\square$

Assumption 3 could be relaxed to allow wages to depend on worker type, as long as output losses from hiring an underqualified worker are imperfectly passed through to wages. The firm's trade-off remains the same: foregone persistent profit differences between an ideal and an underqualified workforce vs. the cost of search, losses of employing only $N - 1$ workers.

Proposition 3 implies for a dynamic setting no jobs with underqualified workers but only workers of skill type j on job levels $k \leq j$. As a result, we will not observe the declining productivity path of skill type one when moving from job level one to five in Figure 7. By contrast, we will

observe career dynamics of skill type five moving from job level one to job level five. We indicate the worker type job combinations that will be observed by bars with bold edges in Figure 7.

5.1.4 A stylized dynamic model of job and worker search

In expanding our job-level model to capture career dynamics, we maintain that firms design their job-level structure once and for all (Assumption 1) and that wages are determined by the job level (Assumption 3). Additionally, we assume constant worker skills, thus deliberately abstracting from skill accumulation and focusing on skill utilization. Given the job-level structure, contact rates and separations define career dynamics. Workers accept job offers in order to be promoted. However, if a coworker leaves, there is also career progression through internal promotion.

We further assume that a firm hires a discrete number of $N = M$ workers (rather than a continuum), one for each job level.²⁵ Ordering jobs by job levels $k = 1, \dots, M$, this means that we can compactly describe the state of the firm by a vector of $x = (x_1, \dots, x_M)$ where $x_k = j$ denotes the (tested) skill level j of the worker employed at job level k . If the job is not filled, we denote this as $x_k = 0$. We can describe a worker's state, $s = (\theta, x)$, using her job and skill level pair, $\theta = (\theta_1, \theta_2)$, and the firm she works for, x . Without any further restrictions, we would obtain a very rich state space of $(M + 3)!$ different states.

We assume that, in each period, only one worker within a firm may become “mobile” (with probability ν). However, to simplify the state space, we assume that only workers in firms without missing workers ($x_k > 0$ for all k) can become mobile. Consequently, we only observe firms with all job levels filled or firms with exactly one missing worker.


An employed mobile worker either separates exogenously into unemployment at rate δ or meets a firm with a vacancy (a firm with some $x_k = 0$) at rate λ_1 . If the vacancy is at a higher level than her current job, the worker will switch employers if she also has the required skills or more. Both types of separation turn the worker's old firm into a firm with a vacancy at level θ_1 . If a worker leaves a firm, this potentially generates the opportunity of promotions of incumbent workers if their skills at least match the now-open position. This defines a promotion mapping $\mathcal{P}(x)$. This mapping fulfills that $(\forall k : x_k > 0) \Rightarrow \mathcal{P}(x) = x$ and $\mathcal{P}(\mathcal{P}(x)) = \mathcal{P}(x)$, i.e., there are no promotions in full firms and the firm applies the promotion operator until no further promotions are possible. The promotions then change the firm description for all the remaining incumbent workers, $x' = \mathcal{P}(x)$. This firm description x' will be the relevant state when searching for a new worker next period.

Unemployed workers are always mobile and meet firms with vacancies at the same rate λ_1 . Firms with a vacancy meet a searching worker at rate λ_2 and the ratio of the two contact rates needs to equal labor market tightness. This structure restricts the state space to a more tractable dimensionality and induces a first-order Markov chain given the equilibrium distribution of (searching) firms (over states x) and searching workers (over states θ).

To bring the model to the data, we assume, first, that there are five job levels ($M = 5$). Second, we assume promotions, $\mathcal{P}(x)$, to move workers up by at most one level, which is guided by the empirical evidence in Section 5.2. Third, we allow for chains of promotions beginning

²⁵This can be viewed as an implicit assumption on $F(\alpha)$ rendering $\Delta F_j = \frac{1}{M}$ optimal.

Figure 8: Average wages over the life cycle by skill group in the model




from the top. This leaves us with 1,085 attainable states. We calibrate the model to quarterly frequency and set the separation rate into unemployment $\nu\delta$ in the model to match the empirical employment to unemployment transition rate of 0.93% equating this with the separation rate of mobile workers.²⁶ The model endogenously features a large difference in unemployment levels between skill groups, consistent with the data (see Bayer et al., 2025, for a detailed discussion of the concentration of labor market risks). Grouping workers by lifetime income reveals that those in the bottom quintile have a 24.3 times higher unemployment rate than those in the top quintile. We calibrate the probability of mobility ν to match this difference in the model. Finally, we set the contact rate of firms with vacancies $\lambda_2 = 1.0$ and set λ_1 according to the equilibrium labor market tightness. We determine the equilibrium as the induced ergodic distribution of worker-firm pairs.

In order to simulate a life cycle of workers, we model a mass zero of workers entering the labor market matched with a firm but at the lowest job level. The other workers of that firm work at the job level that fits their skills. To reduce the impact of the choice of the entry match, we drop the first 12 quarters of the simulation. Figure 8 plots the resulting average wage profile by skill group. Here, we use empirical wage differences between job levels to map job levels to wages (details of the estimation are discussed in Section 5.3). As one can see, our model, which includes job-level structure, skill requirements for jobs, and skill-job mismatch and search, generates a pronounced life-cycle profile for wages. Workers move from low-level jobs to higher-level jobs if their skills permit. This allows them to utilize their human capital more and more over time. This process of climbing the career ladder leads to life-cycle wage growth and differences in life-cycle wage growth across skill groups. The increasing human capital utilization will show up as increasing productivity over the life cycle despite fixed skill levels.

Next, we will document empirically the extent to which the life-cycle is mediated to the profile of wages through the level of jobs, and we will compare this quantitatively to our model predictions. These predictions come from the described parsimoniously parameterized model calibrated to simple flows of unemployment and differences in unemployment across skill groups.

²⁶We rely on German social security data from the old-age pension system, as in Bayer et al. (2025). The data provide monthly labor market states and allow us to observe entire work histories starting at age 14, enabling us to construct average lifetime incomes. See Appendix F for further details.

Figure 9: Promotion and demotion rates by age in data and model



Notes: Annual promotion and demotion rates by age from model and data. Data show promotion and demotion rates for men based on SOEP data, years 1984-2015. All rates are shown in percentages. The left panel shows promotion and demotion rates for men in the data, the right panel compares the shape of promotions in model and data, the dashed lines in each panel show a quadratic fit to the data. A promotion/demotion is a positive/negative, persistent, change in the the job level of an interviewed worker. Job levels in the SOEP are coded on the [Hoffmeyer-Zlotnik \(2003\)](#) definition. In the model, we calculate promotions/demotions as changes in the job level comparing every fourth quarter. Rates are calculated relative to the share of workers in “complete” firms.

5.2 Job levels and career dynamics in the data

Climbing from low job levels to higher job levels is the driver of wage growth in the model. The panel dimension of the SOEP data allows us to document how workers move up (and down) job levels during their working life. Thus, it allows us to empirically check this key prediction of our model in this dimension to the data. Promotions (demotions) in the SOEP are naturally coded as a change in the job level from the current survey date to a higher (lower) job level at the next survey date.²⁷ We apply the same measurement to model-simulated data using the fourth quarter as our observation period to match the annual frequency of the SOEP data.

Figure 9 reports estimated life-cycle profiles of annual promotion and demotion rates in the data and the corresponding model predictions. Consistent with the model calibration, we focus here on men and report results for promotions and demotions of women in Appendix A.2 and discuss differences to men in detail in Section 5.4.1. We find declining promotion rates during working life, in line with a concave wage profile. Demotion rates are strikingly constant throughout working life and substantially below promotion rates at the beginning of working life. Around the age of 45, the rates of promotion and demotion converge, suggesting that there will be no further average net career progression. Our model, calibrated only to the flows into unemployment and the unemployment differences between skill groups, matches the empirical promotion dynamics very well. Even the level of promotions and demotions is fitted well—the model produces 75% of the empirical promotion rate. The lower rate emerges as we abstract from other (non-pecuniary) motives for job mobility that induce workers to change jobs, even if this implies a demotion and subsequently opens up the possibility of promotion for another worker.

²⁷To remove measurement error, we only consider job level changes as promotions/demotions if they are neither reversals of a change in the preceding year nor are undone by a reversal in the following year. We find that 88% of all promotions are upward movements by one job level. Male promotions involve slightly more often a multiple

Table 4: Promotions and demotions by stayer and occupation status

	non-stayer (%)		stayer (%)	
	different occ.	same occ.	different occ.	same occ.
demotion	32.4	5.6	3.4	58.6
no change	6.6	5.0	1.4	87.0
promotion	20.5	6.3	5.4	67.8
population share	9.8	5.1	2.0	83.2

Notes: Shares of all promotions and demotions by type of labor market transition. Columns distinguish between non-stayers and stayers, with further breakdowns by occupational change. Each row sums to 100%.

A key implication of our model is that most promotions occur without the worker changing employer (62.3%). The SOEP data allow us to compare model and data along this dimension as well. In Table 4, we distinguish between stayers, who are employed by the same employer for two consecutive full years, and those who change employers. Employer changers are individuals who either experienced a nonemployment period between the two survey dates or were employed by their current employer for less than one year on the second survey date. We further split the groups based on whether the worker changed occupation.²⁸

Consistent with the model, we find that over 68% of promotions occur for workers who stay with their employer and remain in the same occupation. An additional 6% of promotions involve changing employer but not occupation, and conversely 5% involve an occupation change but no employer change. Only 20% of promotions involve both an employer and an occupation change. The importance of promotions at the same employer for life-cycle career dynamics in the SOEP data support the key model prediction that most career advancement occurs without labor market mobility. The findings on job level and occupation changes also reinforce the view that job levels and occupations are only weakly correlated.

An additional aspect of the data is that we find a similar distribution between stayers and changers for demotions: about 60% of demotions occur within the same employer, while 40% involve a change of employer. Furthermore, among workers without a promotion or demotion, we find that 88% stay with their employer. Our simple model does not provide a rationale why 12% of all worker-year observations without a change in the job level involve an employer change (no lateral moves), nor does it, by assumption, capture worker demotions that occur within the same firm (e.g., in lieu of a separation). We are confident that a slightly richer theory with, e.g., idiosyncratic (non-pecuniary) shocks can account for these aspects of career dynamics, too.

In a final step, we use the SOEP data to explore how career progression changes with labor market mobility. In Table 5, we report the share of promotions, demotions, and lateral moves

level change (13.8%) compared to female promotions (8.8%).

²⁸For occupation changes, we consider workers who answer the question whether “there has been a change in their job” affirmatively and report an occupation change. We condition on the information of job change to reduce measurement error in the occupation codes. It is well known that occupation codes are prone to be recorded with error so that occupational changes are too prevalent in survey data (Kambourov and Manovskii, 2013).

Table 5: Promotions and demotions for labor market transitions

	stayer	employer change	non- employment	any occ. change	stayer, occ. change	all workers
demotion	2.1	6.7	10.8	11.0	6.0	3.0
no change	93.6	84.3	77.3	75.6	77.6	92.0
promotion	4.3	9.0	11.9	13.5	16.4	5.0
net promotion	2.2	2.3	1.1	2.5	10.3	2.0

Notes: Promotions and demotions shares in % for different mobility events (see text for details). Each column shows a mobility event and the share of workers conditional on this mobility event who have a promotion or demotion. The row *net promotion* reports the difference between promotion and demotion rates for each mobility event. The first three rows (excluding net promotions) of each column sum to 100%.

conditional on employer changes, transitions through nonemployment, and occupation changes. Although most promotions occur within the same employer, workers who stay with their employer are less likely to be promoted (4.3%) or demoted (2.1%).²⁹ Workers who change their employer have both higher promotion (9.0%) and demotion (6.7%) rates but their net promotion rate (promotion minus demotion rate) is close to that of stayers with (2.3% vs. 2.2%). Zooming in on employer changes that involve a non-employment spell, we see an increase of *both* promotion (11.9%) and demotion (10.8%) rates relative to employer changes, but the lowest net promotion rate of all groups (1.1%).³⁰ Finally, we look at the group of workers who change occupations, either while staying with the same employer or while changing employers. Two important observations emerge: First, they have the highest promotion (13.5%) but also the highest demotion rates (11.0%) of all groups. In particular, the small group of workers that stay with their employer but change occupation show the highest net promotion rate. However, and this is the second observation, the bulk of all occupation changes, three out of four, come without a change in job level, just as the bulk of job level changes (Table 4) come without an occupation change. This result further corroborates our finding on the differences between job levels and occupations.

5.3 Job levels and wage growth and inequality over the life cycle

The prevalence of promotion dynamics in the data, the large differences between wages at different job levels in our descriptive analysis, and the predictions of our newly developed model directly lead to the question of the quantitative importance of job level changes for life-cycle wage dynamics in the data. To explore this question, we return to SES data as our main data source and develop a methodology to quantify the importance of job-level changes during working life for wage growth and for the increasing dispersion of wages with age as the two

²⁹If we condition both on employer and occupation stayers as in the last column of Table 5 the corresponding numbers are 2.1%, 93.9%, and 4.0% respectively.

³⁰Non-employment considers the case where we observe an intervening spell of non-employment between two employment spells.

salient life-cycle wage dynamics.³¹ The main challenge in using the SES data is to deal with its repeated cross-sectional structure, where unobserved individual heterogeneity may drive career dynamics. We present an established methodology that allows us to decompose life-cycle wage dynamics in the presence of such heterogeneity given this data structure.

5.3.1 Methodology

We start from the following empirical model of log wages w_{ipt} of worker i at plant p at time t

$$w_{ipt} = \gamma_i + \zeta_{pt} + \beta_J J_{ipt} + \beta_I I_{ipt} + \epsilon_{ipt}, \quad (24)$$

where J_{ipt} are the characteristics of the job of individual i at plant p at time t , I_{ipt} are the characteristics of the individual itself, γ_i is a worker-fixed effect, and ζ_{pt} is a plant-year effect. The *individual component*, $\beta_I I_{ipt}$, captures the wage effect of observable worker characteristics, i.e. education and experience, which are included as education dummies and gender-specific age dummies.³² It does not contain unobserved worker characteristics captured by γ_i . The *job component*, $\beta_J J_{ipt}$, captures the characteristics of a job. For the job level, we use a full set of dummies for two-digit occupations alongside five job-level dummies (without dummy interactions). Later, we will decompose this job component into its two subcomponents and demonstrate that the job level is the dominant one.

One challenge for the decomposition of life-cycle wage dynamics is that unobserved individual characteristics might jointly affect wages and the career progression of workers. If workers are also paid for their individual productivity, a simple OLS estimate of wages on job levels would then be inflated because more able workers in line with our theory are more likely to end up at higher job levels. We deal with the challenge of unobserved heterogeneity by relying on two different approaches. First, we estimate a synthetic panel specification that exploits the fact that aggregating microdata to the cohort level creates a panel structure so that we can control for unobserved heterogeneity of cohorts in the decomposition (see [Deaton \(1985\)](#) and [Verbeek \(2008\)](#) for an overview of the method).³³ Aggregating the data to the cohort level eliminates purely individual unobserved heterogeneity and thus is likely to reduce the general importance of unobserved heterogeneity. Taking advantage of the matched employer-employee nature of our data, we also control for firm-year fixed effects. Thus, the identifying variation is net of cohort and firm-year (and thus region-year) effects. In other words, we exploit how wages and (job) characteristics evolve over time within a cohort while controlling for variation that affects all cohorts in a firm and region.

An example of the type of variation that we use is the entry of a new plant into a region for which this plant has an atypical organizational structure. If this affects the job characteristics of the cohorts of workers who are young at the time of the plant's entry more than those of older cohorts, we get a variation that identifies the job effect. Such an effect should be strongest around the entry date of a plant because younger workers are more mobile, and thus more likely

³¹The high data quality, the detailed information on workers' earnings and hours, and the matched employer-employee information make the SES the preferred dataset over the SOEP for this analysis.

³²We combine three-year age groups to identify cohort effects despite the four-year interval between surveys.

³³We provide results based on pooled worker-level OLS in Appendix H.5.

to take advantage of new job opportunities. At the same time, the cohort-fixed effect controls for a potential correlation of the atypical new entrant with unobserved skills of newly entering cohorts. Any observable changes in the job composition (occupations, job levels) or worker composition (educational attainment) will be directly controlled for as well. Another example of identifying variation would be (regional) variations in industry composition with heterogeneous effects on cohorts. More generally, identification comes from changes in the structure of job opportunities within a region over time, but since this affects different age groups differently, the variation is not captured by the region-year effect. We rely on the identifying variation from the second example and estimate an instrumental variable regression as a robustness analysis using shift-shares (Bartik, 1993). We provide a further discussion of the identification challenges of our baseline synthetic panel approach and the alternative instrumental variable regressions in Appendix G. We opt for OLS estimation with cohort-fixed effects in the main part because it is easier to interpret and has favorable small-sample properties.

Specifically, in our baseline approach, we first control for plant-year effects by subtracting plant-level averages, $X_{.pt}$, for the plant p at which worker i works in year t :

$$w_{ipt} - w_{.pt} = (\gamma_i - \gamma_{.pt}) + \beta_J(J_{it} - J_{.pt}) + \beta_I(I_{it} - I_{.pt}) + (\epsilon_{it} - \epsilon_{.pt}), \quad (25)$$

as a result, the plant component, ζ_{pt} , drops out of the regression—we explain below how we re-construct the estimate of the plant component, $\tilde{\zeta}_{pt}$ later. Subsequently, we define cohorts based on workers' gender, year of birth, and regional information (north-south-east-west),³⁴ and we aggregate the variables to the cohort level to obtain

$$\hat{w}_{ct} = \hat{\gamma}_c + \beta_J \hat{J}_{ct} + \beta_I \hat{I}_{ct} + \hat{\epsilon}_{ct} + (\hat{\gamma}_{ct} - \hat{\gamma}_c), \quad (26)$$

where $\hat{X}_{ct}$ is the average difference ($X_{it} - X_{.pt}$) within cohort c . The coefficients of interest for our decomposition are β_J and β_I that we estimate from (26) rather than (24). Unobserved heterogeneity at the individual level, which could cause cohorts to differ in their unobserved skills, is absorbed by the fixed effect $\hat{\gamma}_c$. Therefore, we rely on the key idea of Deaton's (1985) synthetic panel estimator and use between-group variation in outcomes and observables to identify the coefficients of interest. For our estimation, we assume further that there is no systematic variation in the co-worker composition of a cohort over time, that is $(\hat{\gamma}_{ct} - \hat{\gamma}_c) \perp \hat{J}_{ct}, \hat{I}_{ct}$.³⁵ Under this assumption, OLS with cohort-fixed effects obtains unbiased estimates $\tilde{\beta}_J$ and $\tilde{\beta}_I$.

Using the coefficient estimates $\tilde{\beta}_J$ and $\tilde{\beta}_I$, we directly construct the worker and job component as $\tilde{\beta}_J J_{ipt}$ and $\tilde{\beta}_I I_{ipt}$. The estimated plant component, $\tilde{\zeta}_{pt}$, is constructed as the residual plant-level wage after accounting for observable worker and job characteristics. It is given by

$$\tilde{\zeta}_{pt} = \zeta_{pt} + \gamma_{.pt} + \epsilon_{.pt} = w_{.pt} - \tilde{\beta}_J J_{.pt} - \tilde{\beta}_I I_{.pt}. \quad (27)$$

³⁴Using social security data (SIAB 1975-2021), we find an annual gross migration rate of 0.49% for workers aged 25 to 55 years for the time period 1975 to 2021.

³⁵Given that cohorts are large, the typical set of co-workers of an entire cohort is a broad mixture of workers and it is unlikely that this mixture changes largely over time. In the special case, of fixed co-workers of a cohort over time, the difference $(\hat{\gamma}_{ct} - \hat{\gamma}_c)$ would be strictly zero. This holds, for example, if employers replace retiring workers by equally skilled workers from the new cohort entering the labor market.

This construction implies that the plant component corrects the average wage at a plant (w_{pt}) for differences in organizational structure and workforce composition by removing the average individual ($\tilde{\beta}_I I_{pt}$) and job components ($\tilde{\beta}_J J_{pt}$) across plants. Hence, a high-wage plant is a plant that pays on average more than other plants after accounting for worker and job observables at that plant. Unlike [Abowd et al. \(1999\)](#), we do not have individual-level panel information to identify residual plant fixed effects ζ_{pt} , so that the average plant effect is not separately identified from the average worker fixed effect γ_{pt} of workers working at plant p . Conversely, the average cohort fixed effects are not separately identified from the average co-worker fixed effects of that cohort. For our estimator $\tilde{\zeta}_{pt}$, $\mathbb{E}(\tilde{\zeta}_{pt} | \zeta_{pt} = \zeta) \gtrless \zeta$ holds only if there is positive / no / negative assortative matching.³⁶

The minimum number of observations across cohort-year cells is 265, the maximum is 8,383, the median is 3,159, and the mean is 3,285. Identifying assumptions for our regression are that all coefficients, in particular the pure experience effects on life cycles (captured by β_I), are stable across cohorts and that regressors have overlapping support across cohorts.

5.3.2 Decomposing wage growth and wage dispersion over a worker’s life

Based on our estimation results, we decompose the average wage growth over the life cycle. We decompose the wage growth of male and female workers separately because these decompositions show very distinct patterns. We begin with a discussion of men’s wage dynamics and devote the separate Section 5.4.1 to the differences between men and women.

In the empirical decomposition, the estimated worker component, $\tilde{\beta}_I I_{ipt}$, and job component, $\tilde{\beta}_J J_{ipt}$, include worker and job characteristics that can still contain cohort effects. Therefore, we remove cohort effects from the estimated components by regressing them on a full set of cohort and age dummies. We report the coefficients on the age dummies as our life-cycle profiles and always normalize the log wage components of a 25-year-old worker to zero.


Figure 10(a) reports the decomposition of mean log wages for men.³⁷ On average, wages grow by approximately 56 log points over the life cycle and the job component accounts for more than 50% of this wage growth. Moving to better-paying plants, the plant component, contributes approximately 25% to life-cycle wage growth (see also [Topel and Ward, 1992](#); [Bagger et al., 2014](#)). The remaining part, the individual component, captures a pure experience effect.

The important mediating role of changing job levels for wage growth becomes apparent when looking at the decomposition of the job component. (Figure 10(b)). We find that an increase in the average job level accounts for most of the wage growth in the job component (82%) and that movements across occupations contribute less than 20% to the wage growth in the job component once we control for job levels. Hence, the single most important component of the life-cycle wage growth is accounted for by workers taking on jobs at higher job levels (“climbing the job ladder”). Figure 10(c) compares the job-level component from panel (b) to the model prediction. We simulate the model as described in Section 5.1.4 and set the level of wages at

³⁶The estimate by [Card et al. \(2013\)](#) for Germany is based on the [Abowd et al. \(1999\)](#) approach, which is not directly comparable to our results, as it does not control for the organizational structure at the firm. They find a modest positive contribution to cross-sectional wage inequality from assortative matching.

³⁷The estimated life-cycle profiles are estimated from both cross-sectional and within-cohort variation. Out of the 30-year life cycle, we typically observe a cohort over 12 years.

Figure 10: Wage and job component decomposition for men



Notes: Left panel: Decomposition of log wage differences by age relative to age 25 for male workers. The dashed line corresponds to the individual, the dotted line to the plant, and the dash-dotted line to the job component; the solid line (total) equals the sum of the three components. The horizontal axis shows age, and the vertical axis shows the log wage difference. Middle panel: Decomposition of the job component (solid line) into the contribution of occupations (dotted) and job levels (dashed). Right panel: Comparison of model prediction for job-level component (solid line) to profile from the empirical decomposition (dashed).


each job level to the estimated job-level wage from the SES data. The fit of the model and the data is very close in terms of shape and scope, which reiterates the fact that the model aligns closely with the empirical promotion patterns (Section 5.1.4 (b)).

The close fit to average growth raises the question of the extent to which changing job levels in general, and our model in particular, can explain the life-cycle pattern of wage inequality. To this end, we decompose in a second step rising wage inequality over the life cycle. The high degree of statistical determination of 81% in our data (see Table 2) implies that this decomposition into observables is comprehensive. By contrast, existing microdata based on cross-sectional regressions typically explain about 30% of wage inequality by observables, leaving most of the wage inequality unexplained and unrelated to observable job, firm, or worker characteristics.

Figure 11(a) shows that, in the SES data, observables account for virtually the entire increase in wage inequality over the life cycle. In particular, job characteristics captured in the job component are the prime explanatory variables. Overall, the variance of log wages increases from about 12 log points to 24 log points. The variance of the plant component contributes about 6 log points to the level of wage dispersion, but is virtually flat over the life cycle. The job component, on the other hand, shows an increase in its variance of 11 log points, from 6 to 17 log points. Thus, almost all of the increase in the variance of wages is due to workers becoming increasingly different in the types of jobs they hold. As for average wages, the job level is the main driving variable (not shown). The individual component contributes virtually nothing to wage dispersion. Since the distribution of individual characteristics by age largely reflects the distribution of education, this implies that education has a negligible direct effect on wage differences across workers once job level is controlled for (see Section 5.4.2 for details).

Figure 11(b) complements these results by adding the covariances of the job, individual, and plant components by age. We find that the covariance terms are on average close to zero, and the two covariance profiles including the individual component are also flat over the life cycle. The plant-job component shows a systematic life-cycle pattern. This increasing correlation implies that young workers are mostly in high-level jobs in plants that do not pay well on average, and

Figure 11: Variance-covariance decomposition for men



Notes: Left panel: Decomposition of the variance of log wages by age for male workers. Variances of all components are calculated by age-cohort cell. The solid line is the variance of total wage, the dashed line is the individual, the dotted line is the plant, and the dash-dotted line is the job component. Right panel: Covariance components for variance decomposition calculated analogously to the left panel; the solid line refers to the covariance of the individual and job component, the dashed line to the covariance of the individual and plant component, and the dotted line to the covariance of the plant and job component; all covariances are within the age-cohort cell.

as workers age, high-level jobs in well-paying plants become more common. In other words, only when young is there a trade-off between plant type and job level; when old, workers in well-paying plants also face organizational structures that favor more CAR-intensive jobs. The plant component in isolation does not show such systematic variation over the life cycle. The additional covariance term between the plant and job components increases from slightly less than -0.5 log points to slightly more than 0.5 log points over the life cycle. Thus, the covariance terms contribute another 2 log points to the increase in variance over the life cycle (twice the increase in the covariance term over a worker's life), so that the job component and its covariance terms account for virtually all of the increase in wage dispersion over the life.


Figure 11(c) compares the predictions from the model simulation and the data. By construction, the job level explains all wage differences in the model, so we compare only the life-cycle dynamics of the job-level component. Like the fit of the average wage, the fit of the life-cycle dynamics of inequality is very tight. That is, the model predicts not only the average promotion and demotion dynamics and their wage consequences from the data, but also how career dynamics translate into more dispersion over workers' lives.

In summary, these findings demonstrate that, in the data, most of the life-cycle wage growth results from changes in the CAR-intensity of a worker's job; workers on average progress to more complex jobs that require jobholders to execute more autonomy and take on more responsibility. The decomposition thus supports the view embodied in our theory of an important role of job levels in accounting for wage differences. We next use this framework to examine how job levels also have a key mediating role in other salient wage facts.

5.4 The mediating role of job levels in between-group wage differences

The gender wage gap, returns to education, and returns to seniority have been widely studied as wage differences between groups. For each of these, we provide a new interpretation as arising from differences in the job level composition of these groups and show that these differences arise naturally in our structural model.

Figure 12: Average wages, their decomposition and wage inequality over a female worker's life



Notes: Top-row: log wage levels by age relative to age 25. The horizontal axis shows age, and the vertical axis shows the log wage difference. Left panel: Decomposition for women into components—individual (dashed line), plant (dotted line), and job (dash-dotted); the solid line (total) equals the sum of the three components. Middle panel: Gender wage gap in the job component from model and data. Solid line shows job component for men. Dashed line shows job component for women. Dotted lines show the model implies job components. Right panel: Cumulative net promotions from SOEP data for men and women. Bottom row: log wage dispersion by age. The horizontal axis shows age. Left panel: Decomposition of log-wage variance for women into the variance of components—individual (dashed line), plant (dotted line), and job (dash-dotted); the solid line (total) equals the variance of wages (including covariances). Middle panel: Variance of the job component in model and data. Solid line refers the variance of the job component of men in the data. Dashed line refers to women in the data. Model implied job component variances are dotted. Right panel: Covariances of individual, job, and plant component for women in the data, see Figure 11.

5.4.1 Gender wage gap

At age 25, women in our sample earn about 7% less than men. By age 50, the gap has widened and women earn 30% less than men. To understand this widening of the average wage gap, Figure 12(a) decomposes the life-cycle wage growth for women according to the decomposition for men in Figure 10. As we did for men, we find that the job component accounts for most of the life-cycle wage growth. However, comparing the job component with that of men—Figure 12(b)—we find that the job component (mainly driven by the job level, not shown) flattens out for women after age 30, exactly the age at which the wage gap starts to widen.³⁸ This widening of the job level gap finds its counterpart in a net promotion gap in the SOEP data, see Figure 12(c). In terms of variances and covariances, see Figure 12(d) and (f), we find the same qualitative picture for women and men, namely that the increasing variance of the job component is the dominant driver of the increase in wage variance over the life cycle. Quantitatively, however, the overall increase in wage inequality is one third smaller for women

³⁸One might be concerned that female careers have changed across cohorts in recent decades. The SOEP data on cumulative net promotions (Figure 13(c)), covering more than 30 years for an individual, are congruent with the SES life-cycle patterns that load more heavily on cross-sectional information.

than for men (8 log points vs. 12 log points), and it levels off earlier. Including the establishment-job covariance, the job component again accounts for virtually all of the increase in life-cycle wage inequality. Taken together, this shows that job levels appear to be central to the life-cycle wage gap between men and women and their differences in wage inequality.

The differences in the growth of the job component and the net promotion pattern also arise naturally in our model when calibrated to the different labor market dynamics of women (instead of calibrating to men). In the German labor market, we see that in and after childbearing age, the differences between men and women in terms of part-time and marginal employment increase dramatically. By the age of 40, 60% of women are in part-time jobs, compared to only 14% of men.³⁹ Similarly, we observe a sharply increasing gap between men and women in marginal employment (“minijobs” of less than 10 hours per week), reaching 16% at age 40 (24.2% vs. 8.7%). The model captures this increasing gap as the risk of an “immobility shock” affecting only women. “Immobile” women neither receive outside offers nor are promoted until they recover from the immobility shock. We calibrate a probability of recovering from immobility that corresponds to an expected time in the immobile state of 3 years. We set the quarterly probability of entering the immobile state at 7% to match the part-time gap at age 40. When immobile, women face a small probability of a career interruption shock ι , which we interpret as a career restart. We calibrate the employment gap in the lowest job level to the 16% gap in marginal jobs at age 40, yielding $\iota = 0.015$. Figure 12 (b) and (e) compare the implied wage dynamics over the life cycle in terms of means and variances for model and data. The model results for men are as discussed in the previous subsection. The close fit of the immobility-shock augmented model to the data for women leads us to interpret the rising gender wage gap as a gender promotion gap due to women’s career interruptions. Importantly, the lower wages for women in the model do not result from lower skills (human capital), but from an underutilization of the skills of high-skilled women who work in low-level jobs.

5.4.2 Returns to education

The returns to education, i.e. wage differentials between workers with different education, are another prominent between-group wage differential (e.g. [Goldin and Katz, 2009](#); [Katz and Murphy, 1992](#); [Juhn et al., 1993](#)). In our model, skill differences affect wages only when mediated through (faster) progression to higher-level jobs.⁴⁰ This mediating role aligns with the empirical finding that the estimated returns to education from (26) are close to zero, when we control for job levels as in our baseline. First column of Table 6 shows this for college education.

The columns of Table 6 decompose the mediating role of job levels. Column 1 shows that college-educated workers have only 10 log points higher wages than workers with vocational training, once we control for all other observable job and firm characteristics. Leaving out only occupational information, but maintaining job-level information (column two, *w/o occupations*) leaves the return to education unchanged. Leaving out job-level information (column three, *w/o job levels*), but maintaining occupations that correlate with job levels (Section 4.3) increases the

³⁹These results are based on social security data (SIAB 1975-2021) for the period 2006 to 2017, which is consistent with the period covered by the SES data.

⁴⁰Table A1 provides descriptive evidence of career progression by education group over the life cycle.

Table 6: Transmission of returns to education through jobs

	data				model
	baseline	w/o occupations	w/o job levels	w/o all job info	
College	0.10**	0.09***	0.27***	0.49***	0.41

Notes: The table displays the coefficients of dummies for college education in a regression of log wages on worker and job characteristics using cohort fixed effects across four different specifications: first our baseline, second a specification that leaves out occupation but keeps job-level information, third a specification that leaves out job-level but keeps occupation information, and fourth a specification that leaves out all job information (levels and occupations). The baseline education category is vocational training. *, **, *** indicate significance at the 10%, 5%, and 1% levels, respectively.

estimated return to education to 27 log points. Without any job information, but maintaining all other controls (column four, *w/o all job info*), increases the estimated return to 49 log points. In the model, the return to education, when controlling for job-level differences, is zero by construction. For the model simulation, we assume that firms screen for $N = 5$ worker skill groups. To construct a model equivalent to the return to education, we compare the simulated wages of workers with skill level five to those of skill levels two to four. This grouping matches the population shares of college educated workers of roughly one-fifth and the share of workers with vocational training that is roughly three-fifth. Table 6 reports this model-based estimate for the return to experience. College-educated workers in the model (skill level five) have an average premium of 41 log points. This skill premium moderated through job levels is well in line with the estimated education effect we document in the data that is moderated through job levels (39 log points, the difference between *w/o all job info* and *baseline*).

5.4.3 Returns to seniority

Returns to seniority are the third wage differentials between groups of workers, which we show to be mediated through job levels, too. Returns to seniority refer to an *additional* effect of having more seniority than one’s peers, over and above the effect of tenure and experience itself. In fact, [Buhai et al. \(2014\)](#) provide an identification strategy and empirically establish that not only a worker’s own tenure, but also the relative ranking among coworkers determines workers’ wages. Similarly, and consistent with the internal promotion in our model, we know from [Jäger and Heining \(2022\)](#) that workers’ wages and the probability of moving within a plant to better-paying jobs increase when coworkers leave the plant (in their case, because of death). In the data, we estimate the effect of seniority ranking among peers within a plant who may be competing for career advancement. We consider two measures of wages. The first is the log wage reported in the data. The second measure is the job-level wage, which is the wage predicted by a worker’s current job level using the coefficient estimates from the baseline regression (26). We also consider two measures of seniority ranking. In the first case, we include a dummy variable for only the most experienced worker in each peer group (based on tenure with the firm). The estimated coefficient quantifies the “silverback effect”—the effect of being the most experienced member of the peer group on job-level wages. In the second case, we use what we refer to as the seniority rank. For the seniority rank, we follow [Buhai et al. \(2014\)](#) and calculate the distance

Table 7: Being the silverback: the effect of experience ranking on job-level wages

Wage measure	Relative experience concept					
	Silverback effect			Seniority rank		
	Log	Job level	Model	Log	Job level	Model
More experienced	4.4***	2.9***	2.0	1.6***	1.8***	1.6
adj. R^2	0.70	0.50		0.70	0.50	
N	288,881	288,881		288,881	288,881	

Notes: The table displays the coefficients of an OLS regression of a log wage measure of a worker (multiplied by 100) on two sets of controls for experience ranking within peer groups of workers. We use two different wage concepts: first, the *raw* log hourly wage of the worker; second, the wage predicted by the worker's *Job level*. A worker's peer group is composed of all workers at the same plant who are at least as old as, and up to five years older than, the worker and have the same educational attainment. Experience ranking controls are described in the text. The regression sample includes all male workers ages 45 to 50. All regressions include a constant, education dummies (coefficients not reported), tenure and tenure squared, and plant-fixed effects. *, **, *** indicate significance at the 10%, 5%, and 1% levels, respectively. The column *model* repeats the exercise on simulated data from our model, see main text for details.

between ranks as $\log(N_i + 1 - r_i) - \log(N_i)$, where r_i is the seniority rank of worker i within their peer group and N_i is the number of members in worker i 's peer group. For example, the most experienced worker in a peer group has a seniority rank of $r_i = 1$, and the least experienced worker has a seniority rank of N_i . Within each peer group, the distance between seniority ranks varies between $[-\log(N_i), 0]$. We restrict the sample to male workers with vocational training or a college education, excluding women due to their different career dynamics after age 30. We define a worker's competitive peer group at a plant as a group of workers who are no more than five years older than the worker in question and who have the same level of education. Within each age-education cell of the plant, we construct the silverback dummy and the distance of seniority ranks. We regress log (job-level) wages on a set of education fixed effects, a quadratic tenure polynomial, firm fixed effects, and our two alternative measures of seniority. Table 7 reports the estimated seniority coefficients in columns *Log* and *Job level*.

On average, we find the *silverback* and *seniority rank* effect to be statistically significant. The higher the tenure of the worker relative to his peers, the higher is his wage. In the first case, the effect of being the most experienced worker, we find that these workers obtain a statistically highly significant 4.4% wage premium for seniority in log wages. Their job-level wages are also 2.9% higher, and consequently, there is only a small residual seniority premium of 1.5% left once we control for job levels. For the second case, using the distance between seniority ranks, we also get a highly significant coefficient of 1.6% for log wages (close to [Buhai et al.](#)'s estimates for Denmark and Portugal) and 1.8% for job-level wages, implying no residual effect. In other words, we find that seniority affects wages primarily by giving senior coworkers an edge over their peers in being assigned to higher job levels—assigned to more CAR-intensive tasks.

We repeat the estimation in our model by first drawing from our simulation 40,000 workers between ages 45 and 50. From these samples, we first construct five-worker production units, selecting workers who have the same coworker description x . We randomly combine five of these

production units to form a “firm” and define a worker’s peer group within the firm. We construct 400,000 of these firms. Using the same method as in the previous subsection, we map skills 2-4 to vocational training and skill 5 to college. This yields an average peer group size of ten, close to the empirical average of eleven peers. The columns *Model* in Table 7 report the results of the estimation in the simulated data. We get, consistent with the data, an important mediating role of job levels with 2.8% for the *silverback effect* and 1.2% for the *seniority rank*.

5.5 Sensitivity and extensions

We provide an extensive sensitivity analysis to our analysis in Appendix H. In particular, we explore several extensions to our baseline specification from the baseline regression (26).

First (Appendix H.1), we check for heterogeneity in the job component of wages across worker groups. Concretely, we look at differences for workers covered by collective bargaining, workers working full-time, and workers working in large plants. In summary, we find that the importance of the job component in accounting for wage dynamics increases for workers not covered by collective bargaining and decreases in large plants. Qualitatively, the results are the same. Results for wage growth are also quantitatively very similar for full-time male workers, while the effect becomes slightly lower for female workers. For the increase in wage dispersion, we find again that the job component becomes more important for workers not covered by collective bargaining and less important in large plants, without a qualitative change in results. The contribution to increasing wage dispersion for full-time workers is slightly lower than in the baseline for both male and female workers. We also explore the sensitivity of our results when we include public employers and publicly controlled firms. When including public employers, we find a 30% larger job component for female wage growth over the life cycle. This finding suggests that public employers offer more opportunities for female career dynamics, in line with over 60% of employees being female at these employers. Overall, we find that our results on the importance of the job component are robust across specifications and sample selections.

Second, we explore more flexible specifications to (26) where we allow the returns to experience to be education-specific (Appendix H.2) and occupation-specific (Appendix H.3). Again, we find that the key finding regarding the importance of job levels for life-cycle wage dynamics remains robust. In the decomposition, we attribute the flexible experience profiles to the individual components and find that more flexible experience profiles hardly affect the decomposition results. These more flexible specifications do not provide any indication that job components are systematically inflated in our more restricted setup.

Third, we explore the effect of not including job-level information for the life-cycle results (Appendix H.4). Without job-level information, the other components absorb some of the wage variation. Consequently, firm-level differences, the individual component, and residual wage dispersion become more important drivers of observable wage differences.

Finally, we estimate in Appendix H.5 the regression in (26) by pooled OLS using cohort-fixed effects only, but we do not control for individual fixed effects. We find that the result of the job component being the key driver of wage dynamics also holds under this specification, but results also suggest that there is a substantial omitted variable bias if we do not control for individual fixed effects. In that sense, the results support our approach based on a synthetic panel.

6 Conclusions

Job levels describe differences in jobs based on the complexity, autonomy, and responsibility (CAR intensity) involved in executing tasks within and across occupations. At least since the 1950s, they have been used in labor market reports by statistical agencies and form the basis for union bargaining agreements and compensation schemes provided by consulting firms for private businesses. Our study is the first to comprehensively investigate their role in accounting for wage differences at a macroeconomic level.

Our paper makes two significant contributions: one to economic theory and the other to empirical analysis. First, we develop a new theory of the employer job-design problem. The solution to this problem gives rise to jobs at different levels and guides our empirical analysis. Regarding the empirical analysis, we use high-quality matched employer-employee microdata from Germany to demonstrate the empirical relevance of job levels, showing that they differ from occupations and education. Using survey data, we construct job-leveling factors based on task execution in order to determine the economic content of the statistical concept of job levels, as well as their explanatory power with respect to wages. While our empirical analysis focuses on Germany, we demonstrate that job levels play a comparable role in explaining wage differences in U.S. data. Combining empirical analysis and theory, we document that observed worker differences are mediated by job levels to observed between-worker wage differences. We find that changes in job level are the main driver of life-cycle wage growth and virtually the only driver of increasing wage dispersion over the life cycle. Additionally, we demonstrate that, although career progression is associated with labor market mobility, most changes of the job level occur within the same employer.

Together, our theory and our empirical findings provide a new, unified interpretation of important differences in wages between workers. Gender, education, and seniority wage differences largely result from differences in progression over job levels. In particular, the gender wage gap emerges as a gender promotion gap, reflecting slower advancement for women at key stages of working life. Importantly, we find that wage differences can result from inherent skill differences, as in the case of the education wage premium, but also from differential utilization of human capital across job levels of equally skilled workers, as in the case of the gender wage gap.

These results support a view of wage determination in which workers are paid based on the content and execution of their jobs rather than solely on their skills. This job-level perspective is based on a key idea of the task-based approach, refining it by incorporating how tasks are organized and executed. This refinement substantially improves the approach's explanatory power in accounting for wage differences, even within occupations. It also suggests a macroeconomic importance of job design and organizational structures as these determine the income risks workers face. Future research may build on this new approach to study the relationship between job design, occupational choice, and specific versus general skill accumulation. Another interesting direction may be to study the impact of labor market policies and technological advancements on job design, employment, and wage structures.

References

- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. Econometrica, 67(2):251–333.
- Acemoglu, D. (2002). Technical change, inequality, and the labor market. Journal of Economic Literature, 40(1):7–72.
- Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Handbook of Labor Economics, volume 4, pages 1043–1171. Elsevier.
- Acemoglu, D. and Autor, D. (2012). What Does Human Capital Do? A Review of Goldin and Katz’s “The Race between Education and Technology”. Journal of Economic Literature, 50(2):426–463.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. The Quarterly Journal of Economics, 118(4):1279–1333.
- Autor, David, H., Katz, L. F., and Kearney, M. S. (2006). The polarization of the US labor market. American Economic Review, 96(2):189–194.
- Bagger, J., Fontaine, F., Postel-Vinay, F., and Robin, J.-M. (2014). Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics. American Economic Review, 104(6):1551–96.
- Baker, G., Gibbs, M., and Holmstrom, B. (1994). The internal economics of the firm: evidence from personnel data. The Quarterly Journal of Economics, 109(4):881–919.
- Bartik, T. J. (1993). Who benefits from local job growth: migrants or the original residents? Regional studies, 27(4):297–311.
- Bayer, C., Hildebrand, S., Hintermaier, T., Kuhn, M., and Ploj, G. (2025). Unequal lives: Heterogeneity in income risk and household portfolios. University of Bonn Working Paper.
- Bloom, N., Sadun, R., and Van Reenen, J. (2016). Management as a technology? Technical report, National Bureau of Economic Research.
- Bronson, M. A. and Thoursie, P. S. (2018). The Lifecycle Wage Growth of Men and Women: Explaining Gender Differences in Wage Trajectories. 2018 Meeting Papers 923, Society for Economic Dynamics.
- Buhai, I. S., Portela, M. A., Teulings, C. N., and van Vuuren, A. (2014). Returns to tenure or seniority? Econometrica, 82(2):705–730.
- Caicedo, S., Lucas, Robert E., J., and Rossi-Hansberg, E. (2019). Learning, career paths, and the distribution of wages. American Economic Journal: Macroeconomics, 11(1):49–88.
- Caliendo, L., Monte, F., and Rossi-Hansberg, E. (2015). The Anatomy of French Production Hierarchies. Journal of Political Economy, 123(4):809–852.
- Card, D., Heining, J., and Kline, P. (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. Quarterly Journal of Economics, 128(3):967.
- Caroli, E. and Van Reenen, J. (2001). Skill-biased organizational change? evidence from a panel of british and french establishments. The Quarterly Journal of Economics, 116(4):1449–1492.
- Carroll, C. D. and Samwick, A. A. (1997). The nature of precautionary wealth. Journal of Monetary Economics, 40(1):41–71.
- Cohen, L. H., Gurun, U. G., and Ozel, N. B. (2026). Too many managers: The strategic use of titles to avoid overtime payments. The Review of Financial Studies, page hhag016. Advance online publication.
- Deaton, A. (1985). Panel data from time series of cross-sections. Journal of Econometrics, 30(1-2):109–126.
- Deaton, A. and Paxson, C. (1994). Intertemporal choice and inequality. Journal of Political Economy, 102(3):437–467.
- Dengler, K., Matthes, B., and Paulus, W. (2014). Occupational tasks in the German labour market. FDZ Methodenreport, 12.
- der Rentenversicherung (FDZ-RV), F. (2020). Versichertenkontenstichprobe 2020 (version 1.0). [Dataset].
- Doeringer, P. B. and Piore, M. J. (1985). Internal labor markets and manpower analysis. ME Sharpe.
- Dohmen, T. J., Kriechel, B., and Pfann, G. A. (2004). Monkey bars and ladders: The importance of lateral and vertical job mobility in internal labor market careers. Journal of Population Economics, 17(2):193–228.
- Ellguth, P. and Kohaut, S. (2021). Tarifbindung und betriebliche Interessenvertretung: Ergebnisse aus dem iab-betriebspanel 2020. WSI-Mitteilungen, 74(4):306–314.
- Fox, J. T. (2009). Firm-size wage gaps, job responsibility, and hierarchical matching. Journal of Labor Economics, 27(1):83–126.

- Freund, L. (2022). Superstar teams: The micro origins and macro implications of coworker complementarities. Available at SSRN 4312245.
- Garicano, L. and Rossi-Hansberg, E. (2004). Inequality and the organization of knowledge. American Economic Review, 94(2):197–202.
- Garicano, L. and Rossi-Hansberg, E. (2006). Organization and inequality in a knowledge economy. The Quarterly Journal of Economics, 121(4):1383–1435.
- Gibbons, R., Katz, L., Lemieux, T., and Parent, D. (2006). Comparative advantage, learning, and sectoral wage determination. Journal of Labor Economics, 23(4):681.
- Gibbons, R. and Waldman, M. (2006). Enriching a theory of wage and promotion dynamics inside firms. Journal of Labor Economics, 24(1):59–107.
- Goebel, J., Grabka, M. M., Liebig, S., Kroh, M., Richter, D., Schröder, C., and Schupp, J. (2019). The German Socio-Economic Panel (SOEP). Jahrbücher für Nationalökonomie und Statistik, 239(2):345–360.
- Goldin, C. and Katz, L. F. (2009). The race between education and technology. harvard university press.
- Guvenen, F. (2009). An empirical investigation of labor income processes. Review of Economic Dynamics, 12(1):58–79.
- Guvenen, F., Kaplan, G., and Song, J. (2021). The glass ceiling and the paper floor: Changing gender composition of top earners since the 1980s. NBER Macroeconomics Annual, 35(1):309–373.
- Guvenen, F. and Smith, A. A. (2014). Inferring labor income risk and partial insurance from economic choices. Econometrica, 82(6):2085–2129.
- Hall, A., Siefer, A., and Tiemann, M. (2018). BIBB/BAuA Employment Survey of the Working Population on Qualification and Working Conditions in Germany 2012. Technical report, BIBB: Federal Institute for Vocational Education and Training.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2005). Two views of inequality over the life cycle. Journal of the European Economic Association, 3(2-3):765–775.
- Hoffmeyer-Zlotnik, J. H. P. (2003). “stellung im beruf” als ersatz für eine berufsklassifikation zur ermittlung von sozialem prestige. ZUMA Nachrichten, 27(53):114–127.
- Hornstein, A., Krusell, P., and Violante, G. L. (2011). Frictional wage dispersion in search models: A quantitative assessment. American Economic Review, 101(7):2873–98.
- Huggett, M., Ventura, G., and Yaron, A. (2006). Human capital and earnings distribution dynamics. Journal of Monetary Economics, 53(2):265–290.
- Huggett, M., Ventura, G., and Yaron, A. (2011). Sources of lifetime inequality. American Economic Review, 101(7):2923–54.
- Jäger, S. and Heining, J. (2022). How substitutable are workers? evidence from worker deaths. Technical report, National Bureau of Economic Research.
- Juhn, C., Murphy, K. M., and Pierce, B. (1993). Wage inequality and the rise in returns to skill. Journal of political Economy, 101(3):410–442.
- Jung, P. and Kuhn, M. (2016). Earnings losses and labor mobility over the life cycle. Journal of the European Economic Association.
- Kambourov, G. and Manovskii, I. (2008). Rising occupational and industry mobility in the United States: 1968–97. International Economic Review, 49(1):41–79.
- Kambourov, G. and Manovskii, I. (2009a). Occupational mobility and wage inequality. The Review of Economic Studies, 76(2):731–759.
- Kambourov, G. and Manovskii, I. (2009b). Occupational specificity of human capital. International Economic Review, 50(1):63–115.
- Kambourov, G. and Manovskii, I. (2013). A cautionary note on using (March) Current Population Survey and Panel Study of Income Dynamics data to study worker mobility. Macroeconomic Dynamics, 17(1):172–194.
- Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. The Quarterly Journal of Economics, 107(1):35–78.
- Krueger, D., Perri, F., Pistaferri, L., and Violante, G. L. (2010). Cross Sectional Facts for Macroeconomists. Review of Economic Dynamics, 13(1):1–14.
- Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. Econometrica, 68(5):1029–1053.

- Kuhn, M., Luo, J., Manovskii, I., and Qiu, X. (2023). Coordinated firm-level work processes and macroeconomic resilience. Journal of Monetary Economics, 137:107–127.
- Lazear, E. P. and Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. Journal of Political Economy, 89(5):841–864.
- Lillard, L. A. and Willis, R. J. (1978). Dynamic aspects of earning mobility. Econometrica, pages 985–1012.
- Low, H., Meghir, C., and Pistaferri, L. (2010). Wage risk and employment risk over the life cycle. American Economic Review, 100(4):1432–67.
- MaCurdy, T. E. (1982). The use of time series processes to model the error structure of earnings in a longitudinal data analysis. Journal of Econometrics, 18(1):83–114.
- Meghir, C. and Pistaferri, L. (2004). Income variance dynamics and heterogeneity. Econometrica, 72(1):1–32.
- Mincer, J. (1974). Schooling, Experience, and Earnings. Columbia University Press, New York.
- Pastorino, E. (2024). Careers in firms: The role of learning about ability and human capital acquisition. Journal of Political Economy, 132(6):1994–2073.
- Pham-Dao, L. (2019). Public insurance and wealth inequality—a euro area analysis. Technical report, University of Bonn.
- Pierce, B. (1999). Using the National Compensation Survey to predict wage rates. Compensation and Working Conditions, 4(4):8–16.
- Rosen, S. (1983). Specialization and human capital. Journal of Labor Economics, 1(1):43–49.
- Rubinstein, Y. and Weiss, Y. (2006). Post schooling wage growth: Investment, search and learning. Handbook of the Economics of Education, 1:1–67.
- Song, J., Price, D. J., Guvenen, F., Bloom, N., and von Wachter, T. (2019). Firming up inequality. The Quarterly Journal of Economics, 134(1):1–50.
- Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. Journal of Labor Economics, 24(2):235–270. Publisher: The University of Chicago Press.
- Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Consumption and risk sharing over the life cycle. Journal of Monetary Economics, 51(3):609–633.
- Strub, S., Gerfin, M., and Buetikofer, A. (2008). Vergleichende Analyse der Löhne von Frauen und Männern anhand der Lohnstrukturserhebungen 1998 bis 2006. Technical report, Schweizer Bundesamt für Statistik.
- Tåg, J., Åstebro, T., and Thompson, P. (2016). Hierarchies and entrepreneurship. European Economic Review, 89:129–147.
- Topel, R. H. and Ward, M. P. (1992). Job mobility and the careers of young men. The Quarterly Journal of Economics, 107(2):439–479.
- U.S. Bureau of Labor Statistics (2013). National Compensation Survey: Guide for Evaluating Your Firm’s Jobs and Pay. US Bureau of Labor Statistics.
- Verbeek, M. (2008). Pseudo-panels and repeated cross-sections. In The Econometrics of Panel Data, pages 369–383. Springer.
- Waldman, M. (1984). Job assignments, signalling, and efficiency. The RAND Journal of Economics, 15(2):255–267.
- Waldman, M. (2012). Theory and evidence in internal labor markets. In Robert, G. and John, R., editors, The Handbook of Organizational Economics, pages 520–572. Princeton University Press.
- Winter, E. (2004). Incentives and discrimination. American Economic Review, 94(3):764–773.
- Winter, E. (2006). Optimal incentives for sequential production processes. The RAND Journal of Economics, 37(2):376–390.
- Yamaguchi, S. (2012). Tasks and heterogeneous human capital. Journal of Labor Economics, 30(1):1–53.


Online Appendix: *Job levels and Wages*

A Additional results from SOEP data

This appendix provides additional information on wages, job levels, and career progression of men and women from the SOEP data. The data cover the period 1984 to 2015.

A.1 Comparison of SOEP and SES data

Figure A1: Wage by age and job level



Notes: The left panel shows mean (log) real wage by age and job level. The right panel shows the mean (log) real wage by age and job levels from SOEP data (1990-2015). Year fixed effects have been removed in both panels. The job-level information is not directly comparable to the SES job levels. See text for details.


First, we consider wage differences by job level over the life cycle. Figure A1 compares (log) wages by age and job level from the SES and SOEP data. The data span different time periods so that average wages differ in levels and the specific coded job levels are not directly comparable because of different coding approaches (see Section 5.2). Still, wage differences over the life cycle show strikingly similar patterns in the SOEP and SES data, in particular, for the four lower job levels.⁴¹ There is roughly an 80 log point difference between average wages on job level 1 and job level 4 and a 40 log point difference between job level 1 and job level 3. The strong increase in wages on job level 5 in the first part of the working life in the SOEP data is a key difference and is related to the different coding approaches. Compared to the SES data, the SOEP job-level data have a smaller top group.

A.2 Promotion/Demotion dynamics of women

Following the discussion of the career ladder dynamics of men in Section 5.2, we report in Figure A2 the corresponding promotion and demotion rates of women. We find that compared to men, women have lower promotion rates in the first part of the life cycle. At age 55, the levels of promotion rates for women and men have converged. Demotion rates for both genders are

⁴¹Because of missing hours information in the SOEP data, we construct wage data only from 1990 onward.

Figure A2: Promotion and demotion rates by age for women



Notes: Annual promotion and demotion rates by age for women based on SOEP data, years 1984-2015. All rates are shown in percentages. A promotion/demotion is a positive/negative, persistent, change in the the job level of an interviewed worker. Job levels in the SOEP are coded on the [Hoffmeyer-Zlotnik \(2003\)](#) definition.

strikingly constant over the entire working life, levels are very close, and demotion rates are substantially below promotion rates at the beginning of working life. As for men, we find that the levels of promotion and demotion rates roughly converge around 45 years of age, implying no further average net career progression thereafter. For women, the overall lower level of promotion rates leads already to few net promotions (promotion minus demotions) after age 40. We provide a model-based discussion and comparison in Section 5.4.1 in the main part of the paper.

B Job levels and education

This appendix documents that the relationship between job levels and education is similar to the relationship between job levels and occupation. Table A1 shows how workers with different levels of education are distributed across job levels. In the table, we separate younger workers (ages 25 to 35) and prime-age workers (35 to 45) and men and women.

First, we find for all age-education groups that they have significant shares of workers ($> 10\%$) on at least three job levels. Hence, we find, as we did for occupations, that education levels span various job levels. In this sense, education is not a sufficient condition to be on a certain job level a finding that reinforces the view of job levels as human capital utilization (Section 5.1.2). Second, we find that education is positively correlated with job levels. Consistent with the prediction from our new theory (Section 5.4.2), we find that workers with more education are on higher job levels in line with the higher complexity of these jobs. Typically, more than 60% of workers with only secondary education are at the two lowest job levels (levels 1 and 2). For workers with a college education, we find that typically over 60% are at the two highest job levels (levels 4 and 5). Third, the distribution across job levels shifts towards higher job levels as workers age. As they age, workers from all education groups move to higher job levels, but the chance of being promoted to the highest job level is if we consider the relative increase in the share, the highest for college-educated men. However, all education groups typically span

Table A1: Share of job levels within formal education and age groups

Education	at ages 25-35 (in %)					at ages 35-45 (in %)				
	1	2	3	4	5	1	2	3	4	5
Men										
Secondary	24.5	37.3	27.9	8.2	2.1	17.6	39.3	30.3	9.4	3.4
Vocational	5.0	14.9	61.9	15.6	2.6	3.5	12.5	53.3	24.0	6.7
College	1.4	2.8	28.0	48.0	19.8	0.4	1.2	14.1	45.1	39.3
Other	18.9	28.8	37.7	11.8	2.9	13.5	28.1	35.9	15.5	7.1
Women										
Secondary	27.7	32.4	28.8	9.4	1.7	32.8	36.2	22.3	6.6	2.0
Vocational	5.0	12.3	66.4	14.5	1.9	5.8	13.3	58.9	19.0	3.0
College	1.9	4.3	35.3	40.9	17.6	0.9	2.5	25.5	44.4	26.8
Other	19.8	24.2	42.9	11.0	2.2	26.6	25.5	34.4	10.5	3.0

Notes: Relative frequencies across job levels in percentage points for different age groups. The top part of the table shows men, the bottom part women. Shares sum within age groups to 100. “Secondary” refers to workers with secondary education but no vocational training. “Vocational” refers to workers with secondary education as well as a vocational degree. “College” refers to all workers with a university or technical college degree. Workers without reported education are in the “Other” group.

three to four job levels if we use a ten percent population share as a cutoff.

C Evidence from the United States

In this appendix, we first discuss additional evidence based on the National Compensation Survey (NCS) for the United States. These results corroborate our conclusions from the German SES data about the importance of job levels in accounting for wage dispersion. Second, we discuss the relationship of job levels and occupations based on the U.S. data corresponding to our analysis for Germany in Section 4.3. Third, we look at wages of assemblers and fabricators in the U.S. and Germany as a case study.

C.1 Job levels in the National Compensation Survey

The NCS is a nationally representative employer survey conducted by the Bureau of Labor Statistics (BLS) that collects information from private industry and state and local government establishments. The survey collects detailed job characteristics that are encoded as job levels using the BLS job-leveling system.⁴² For the job leveling, the BLS interviewers evaluate the duties and responsibilities according to their required knowledge, job controls and complexity, contacts (nature and purpose), and physical environment.⁴³ The BLS job-leveling system relies on *point factor leveling* that assigns points to particular aspects of duties and responsibilities of the job and the required skills, education, and training to execute the job tasks. The job level is the sum of level points from all (four) individual factors. Importantly, job leveling is based on duties and responsibilities and not on assigned job titles in establishments. The distinction to job titles is important as Cohen et al. (2026) highlight the change of job titles by employers in response to labor market regulation without changing the tasks and duties of jobs. The BLS groups jobs in up to 15 job levels. Occupations are coded using the Standard Occupational Classification (SOC) System. The NCS data do not contain worker-level information, but only information about employers and jobs.

Pierce (1999) provides a detailed study of the NCS microdata. He studies the explanatory power of different job-leveling factors for wages and our analysis of BIBB/BAuA data in Section 4.4 is inspired by his original work. He runs cross-sectional wage regressions on different combinations of job and establishment attributes and job-leveling factors. Because the data are collected at the employer-job level, reported wages do not include individual components from overtime pay, bonuses, or other sources so within-job-level variation is absent at the establishment level. This likely explains the even higher explanatory power of observables for cross-sectional wage dispersion compared to the SES data. When all employer and job information is included, observables account for 85% of cross-sectional wage dispersion ($R^2 = 0.847$, Pierce (1999), Table 4), and job-leveling factors alone account for 75% of wage variation. These results corroborate

⁴²See Bureau of Labor Statistics, National Compensation Survey, <https://www.bls.gov/ncs/home.htm>, for a detailed discussion of the NCS data and the job-leveling scheme. The BLS job-leveling scheme is distinct from its occupational coding, although some of the information used for the occupational coding and job leveling overlaps. Occupational classification schemes such as the Standard Occupational Classification (SOC) System differentiate jobs horizontally according to the executed tasks but not vertically according to the CAR intensity of task execution. We provide corresponding evidence based on the German occupational coding (KldB) discussed in Appendix D.2.

⁴³We provide a case study for assemblers and fabricators in Appendix C.3 below to demonstrate that the BLS job levels summarize job differences that are similar to the job levels in the German data.

Table A2: Mean wages in 2015 by job level and occupational group

Level	Occupational groups (SOC)					All
	11-29	31-39	41-43	45-49	51-53	
All	38.22	12.58	17.34	23.09	17.87	23.25
1		8.55	9.63		10.01	9.25
2		9.63	10.53	14.26	12.09	10.48
3	13.01	11.15	12.83	14.78	15.62	12.89
4	15.42	13.67	16.32	18.23	19.67	16.39
5	18.80	18.84	20.14	21.11	20.95	20.13
6	20.96	21.83	24.42	27.47	24.92	23.77
7	24.63	28.03	30.56	30.67	31.27	27.17
8	32.11	33.14	38.82	34.12		32.92
9	37.50		62.13			38.32
10	42.68					44.55
11	50.65					53.26
12	69.37					73.13

Notes: Mean wages by job level and occupational groups from the 2015 National Compensation Survey. Occupational groups follow the 2010 SOC codes. The different occupational groups correspond roughly to Management, Business and Finance, IT and Engineering, Education, Legal, Healthcare (11-29), Service (31-39), Sales and Administration (41-43), Farming, Construction, Maintenance (45-49), and Production and Transportation (51-53). See SOC classification for further details. Missing fields indicate the case of too few observations for a combination of job level and occupational group to be reported by the BLS. These estimates are currently not published by the BLS and have been provided by the BLS upon request.

key findings from our analysis of SES data. First, employer surveys with information on the CAR intensity of jobs deliver high explanatory power on wage dispersion, and second, the job levels are a key contributor to the high explanatory power of wage dispersion in these data. The results from the U.S. data corroborate our findings from Section 4 and thus show that this finding is not a particularity of the German labor market and its institutions.


C.2 Job levels and occupations in the United States

In the second step, we explore, corresponding to our analysis in Section 4.3, the relationship between occupational wage differences and job-level wage differences in the NCS data. The BLS provides information on average wages by job level both across and within occupations. Table A2 shows mean wages by job level and occupational group from the 2015 NCS.⁴⁴ We see that there is a wide variation in wages across job levels within occupational groups. For example, looking at all jobs, we see that going from job level 3 (paying on average \$13) to job level 8 means a wage increase of \$20 per hour. Climbing further to job levels 10, 11, and 12 will lead to stellar wage increases of \$30, \$40, or \$60 per hour. If anything, these data suggest that climbing the career ladder to higher job levels is more important in the United States than in Germany. We also note that when looking across occupation groups, the first occupation group (11-29),

⁴⁴These estimates are currently not published by the BLS and have been provided by the BLS upon an individual data request.

which includes management occupations, has on average much higher wages (\$38.22) than the average over all groups (\$23.25). Strikingly, once we condition on the job level, the “high-wage” occupation group (11-29) tends to have below-average wages. For example, at job level 7 management occupations pay \$24.63, which is less than the average overall occupations at level 7; the latter being \$27.17. Generally, we find that relative wage differences across occupation groups are small and (with one exception) less than 20% once we condition on job levels.

Figure A3: Occupation wage premia and within-occupation wage dispersion




Notes: Left panel: estimated occupation wage premia after controlling for employer and job characteristics with and without job-leveling factors in the National Compensation Survey (NCS). See text for details. Right panel: residual within-occupation wage variance after controlling for employer and job characteristics with and without job-leveling factors in the NCS. All estimation results are taken from Table 7 in [Pierce \(1999\)](#).

The fact that raw differences in occupational wages are largely driven by differences in the average job level of an occupation is also shown in [Pierce \(1999\)](#). [Pierce](#) explores occupational wage premia and within-occupation wage differences with and without controlling for job-level factors. The results are striking. He finds that most occupational wage differences disappear once job-leveling factors have been taken into account and that even within-occupation groups, on average 50% of the wage dispersion is accounted for by job-leveling factors. These findings align closely with our findings from Section 3. Figure A3 visualizes results from Table 7 in [Pierce \(1999\)](#). Figure A3(a) shows occupational wage premia that are estimated as wage differences to an average occupation in a (log) wage regression that includes and excludes job-leveling factors. Figure A3(a) sorts occupations by their estimated occupation-wage premium for the specification without job-leveling factors. We find large occupational wage premia relative to the average wage ranging from almost -50 to +100 log points (dark bars). After including the job-leveling factors, the wage premia decline substantially (light bars). This suggests that a large part of occupational wage differences comes from different distributions across job levels within each occupation and that the job levels themselves account for a large share of wage dispersion (Table A2). Closely related to that, [Pierce \(1999\)](#) finds that if he compares within-occupation wage dispersion without accounting for job-level factors to a specification including job-level factors, then within-occupation wage dispersion in the latter case is largely reduced. Figure A3(b) shows within-occupation wage dispersion for the two specifications. On average, the results show that including job-leveling factors reduces within-occupation wage dispersion by 50%. These results corroborate and strengthen our findings from Section 3 on the distinction

between job levels and occupations.

Figure A4: U.S. wage density across occupations by job level



Notes: Density estimates for residual wages by occupation and job level from U.S. NCS data. *Within all* shows residual wage density after removing the average wage, *within job levels* removes average job level wages, and *within occupations* removes average wages by occupation. Wage observations are for occupation-job-level cells. See text for further details. We observe 269 occupations and 15 job levels.

In Section 4.3, we report results from SES data comparing the explanatory power of job levels, ISCO occupation codes, and finer five-digit KldB occupation codes for wage differences. We find that five job levels account for as much of the wage differences across occupation-job-level cells as 1,077 occupation dummies. Here, we use equivalent occupation-job-level cell data from the NCS to conduct the same analysis on U.S. data. Figure A4 shows the decomposition results for the 2010 NCS data where we observe 269 occupations and 15 job levels. Hence, we have 18 times as many occupations as we have job levels in the decomposition. Figure A4 shows density estimates for residual (log) wages for the equivalent three decomposition cases from Figure 1. In the first case, we remove average wages; that is, we show the variance of (log) wages. This is shown as the case *within all*. Second, we remove average wages by job level. This is shown as the case *within job levels*. Finally, we remove average wages by occupation. This is shown as the case *within occupations*. The legend also reports the estimated standard deviation for each case.⁴⁵ The 15 job levels account for roughly two-thirds of the cross-sectional standard deviation, whereas 269 occupations account for only about a third of the cross-sectional wage variation. Hence, we find an even more striking difference in the explanatory power of job levels in the United States. In general, we corroborate the relationship between job levels and occupations that we document in Section 4.3 based on German data.


C.3 Case study of within-occupation job-level differences across countries

To further substantiate the differences between occupations and job levels, as well as highlight that these differences apply beyond the German case, we present a case study of a narrowly defined occupational group: *Assemblers and Fabricators in Production*. For this study, we begin with the German collective bargaining agreement for metal and steel workers in North Rhine-Westphalia that also forms the basis of our analysis in Section 4.4. At its core, this union

⁴⁵We use unweighted estimates across cells because the BLS does not release cell sizes for these data.

bargaining agreement uses an analytic job-leveling scheme to assign workers to wage scales, which is closely comparable to the BLS job-leveling scheme. Along with the job level, we observe the bargained wage for each level.⁴⁶ For assemblers (*Montierer*) and fabricators (*Maschinen- und Anlagenbauer*), we have job-leveling information that distinguishes these occupations at six different job levels: four for assemblers and two for fabricators.⁴⁷ Starting from the German job-leveling information, i.e., the specific job descriptions of the tasks and duties of the jobholder, we assign job levels based on the BLS job-leveling guide. Using the resulting U.S. job levels, we also assign wages for full-time workers from the 2010 National Compensation Survey (NCS) tabulations for production occupations to the German workers.⁴⁸ Within the NCS data, we remain within a single occupation group according to the classification of the 2000 SOC system and only consider wages at different job levels. After applying the BLS procedure to level the German jobs, we eliminate the mean wage differences between Germany and the United States, ensuring that the average assigned wage is the same in both countries. Thus, we classify German workers as if they worked in the U.S. labor market, comparing their relative pay differences to that of their U.S. counterparts in identical occupations and on the same job level.

Figure A5: Leveling wage structures for assemblers and fabricators in production



Notes: Standardized wages for assemblers and fabricators in production for the United States and Germany. German wages are taken from the union bargaining agreement for metal- and steelworkers in North-Rhine Westphalia. Wages for the United States are derived using the BLS job-leveling approach and NCS wage information by occupation and job level. The job levels are taken from the metal- and steelworkers' bargaining agreement. See text for details.

Figure A5 shows the standardized wage differences across job levels for Germany and the United States. We find that wage structures show a very similar shape across countries, with the key difference being that the German wage structure shows more wage compression in the lower part, which is typically associated with union wage bargaining. Overall, we find wages to be roughly flat across the first four groups in Germany and the first three groups in the United States, and find a positive gradient across the upper three groups. Hence, qualitatively, the

⁴⁶These bargained wages are lower bounds and are usually supplemented by performance components that vary by worker and firm.

⁴⁷One occupation has no directly assigned occupation title but comes from the same task section (*Aufgabenfamilie*).

⁴⁸United States Department of Labor, Bureau of Labor Statistics, National Compensation Survey — Wages, Table 8: Civilian workers: Mean hourly earnings for full-time and part-time workers by work levels, <https://www.bls.gov/ncs/ocs/sp/nctb1482.txt>.

estimates for the corresponding U.S. jobs show a very similar pattern but show more wage dispersion overall. Put differently, differences in how tasks are executed within the organization structure of U.S. firms result in very similar pay differences to the German labor market, a finding that is consistent with the idea that organization-technological differences have the same wage effects across countries. Part of the remaining differences might be because job-level wages in the German collective bargaining agreement only include base pay, whereas they also include incentive and performance pay in the data for the United States. In addition, the wages for Germany are only wages under the specific union bargaining agreement in one state that likely features wage compression. Despite these caveats, we take this case study of a narrowly defined occupational group as further evidence for the importance of job levels for determining wages and wage differences in Germany and the United States.

D Job levels and task-based classification of jobs

In Section 4, we have documented that job levels capture the CAR intensity of a job’s task execution and that they have strong explanatory power for wages. In this appendix, we complement the analysis of Section 4.3 by providing a detailed analysis on the relationship between job levels and their occupational task content.

D.1 Further details on task based occupation grouping versus job levels

The high explanatory power of job levels for wages complements the idea of the task-based approach by [Autor et al. \(2003\)](#) that task execution determines a jobholder’s pay. In contrast to job levels, the task-based approach typically aggregates task information from occupations and classifies jobs depending on the executed tasks along the dimensions of cognitive versus manual tasks and routine versus non-routine tasks. The task-based approach formalizes the idea that some tasks can be executed by computers because task execution follows a fixed set of routines (routine tasks) while others are not amenable to being put into a computer program (non-routine tasks). In fact, categorizing jobs in terms of complexity, autonomy, and responsibility (CAR) has the flavor of ranking jobs along their cognitive-non-routine intensity dimension. In their description of jobs, [Autor et al. \(2003\)](#) focus on the amenability of tasks to be automated using computer software but in addition to routine and non-routine tasks, they distinguish manual and cognitive, analytic and interactive tasks. In total, their task-based classification of occupations consists of five groups: non-routine analytic, non-routine interactive, routine cognitive, routine manual, and non-routine manual.

There are two key differences in the standard task-based approach of the literature to job leveling. First, the standard task-based approach is derived from occupation-level information (not from the Cartesian product of occupations and job levels) so that it does not differentiate within occupations, while job levels provide within-occupation differentiation (Section 4.3 and Appendix C). Second, one way to interpret the task-based approach is that it projects occupational tasks on their amenability to being executed by a computer.⁴⁹ This projection aligns

⁴⁹Examples of tasks from Appendix Table 1 in [Autor et al. \(2003\)](#) are “computes discount, interest, profit, and

most closely with autonomy that enters the CAR intensity measurement of the job level, but it does not relate directly to responsibility and complexity.

D.1.1 Correlation between job levels and occupational tasks

In Section 4.3, we have documented that most occupations span many job levels, but not all occupations are alike in terms of their average job level. Workers in some occupations have, on average, higher job levels than workers in other occupations.⁵⁰ We have also seen that *on average* workers in more analytical non-routine occupations have higher job levels, workers in manual-routine occupations have lower job levels (Figure 3). As explained briefly in the main text, we rely on previous work that has implemented the task-based approach (Spitz-Oener, 2006; Dengler et al., 2014).

Here we provide more details. When bringing together job levels and a task-based classification of occupations, we follow Dengler et al. (2014), who themselves closely follow the original approach by Autor et al. (2003), who, in turn, rely on expert assessments of job task contents.⁵¹ We use the classification by Dengler et al. (2014) based on 2013 occupational tasks to impute the main task content to the 2018 SES data (using the same sample selection as before). We then aggregate to the three-digit occupation level. Our final occupation sample has information on 140 occupations (3-digit KldB2010), their mean log wages and mean job levels from the SES data and the task contents for non-routine analytic (A-NR), non-routine interactive (I-NR), routine cognitive (C-R), routine manual (M-R), and non-routine manual (M-NR) tasks and the main task category from Dengler et al. (2014). Task contents are measured as task shares summing to 100% within each occupation.

Table A3: Task components and average job levels

	A-NR	I-NR	C-R	M-R	M-NR
job level	0.71	0.20	0.14	-0.46	-0.52

Notes: Correlation coefficients between average job level and occupation task shares for non-routine analytic (A-NR), non-routine interactive (I-NR), routine cognitive (C-R), routine manual (M-R), and non-routine manual (M-NR). Data for 140 occupations (3-digit KldB2010) from 2018 SES and Dengler et al. (2014).

In Table A3, we look at correlations between the average job level of an occupation and task shares, complementing Figure 3 which shows the distribution of job levels across workers by main task of the worker’s occupation. The key conclusion of the task-based approach is that routine tasks can be replaced by computers and that non-routine tasks are relative complements to computer capital. In line with the fact that autonomy is one of the key components of job levels and at the same time captures how much workers have to follow a fixed set of rules and cannot make individual decisions on the workflow, we find that the non-routine analytic (A-NR) and non-routine interactive (I-NR) correlate the most positively with the average job level. Manual routine (M-R) correlates the most negatively with the job level but also for non-routine


loss,” “mixes and bakes ingredients according to recipes.”

⁵⁰See Table A2 for this fact based on U.S. data.

⁵¹Spitz-Oener (2006) instead classifies occupations based on BIBB/BAuA survey data on workers to assign tasks to occupations.

manual (M-NR), we find a negative correlation. This latter negative correlation aligns with the fact that there are multiple dimensions entering into job leveling as job levels also capture the complexity and skill requirements of a job, and these are typically low for manual jobs.

Figure A6: Tasks and job levels



Notes: Panels (a) to (e) show average occupation job levels against the five components constructed by the task-based approach: non-routine analytic (A-NR), non-routine interactive (I-NR), routine cognitive (C-R), routine manual (M-R), and non-routine manual (M-NR). Each dot represents one 3-digit KldB2010 occupation. Data are aggregated for 140 occupations from 2018 SES data and data provided by [Dengler et al. \(2014\)](#). The task shares are defined based on [Dengler et al. \(2014\)](#).

To explore these correlations at the occupation level in more detail, Figure A6 shows scatter plots of the average job level and the shares of the different task components across the 140 occupations. Looking at Figure A6(a), we find a clearly upward-sloping relationship between job levels and the share of analytic non-routine tasks. Yet, there is also substantial dispersion. For the interactive non-routine component (I-NR) in Figure A6(b), the data are much more dispersed and a positive relationship is less striking. The cognitive routine tasks (C-R) in Figure A6(c) show a positive relationship, yet again, there is also substantial dispersion. For the manual routine tasks (M-R) in Figure A6(d), we observe that occupations with average job levels of 3 and higher hardly comprise any manual routine tasks. There is a strong decline in the share of jobs with average job levels between 2 and 3. The pattern for the manual non-routine tasks (M-NR) in Figure A6(e) largely resembles the pattern for the manual routine tasks. This similarity likely highlights that, within manual routine occupations, there are also foremen and group leaders who have to act autonomously in the production process and have responsibility for the work of their group members. As the task-based classification is coded from descriptions of occupations and their typical tasks, by construction, it does not allow for within-occupation differences in task content. For example, an architect who “plans and designs private residences, office buildings, factories, and other structures” is carrying-out non-routine interactive tasks as

can be seen in Appendix 1 table of Autor et al. (2003). Yet, there are likely differences in job levels across architects. While the architect at job level 5 decides how the building is going to look, an architect at job level 4 has to work out the planning details according to the plan of the architect at level 5. Job levels capture this additional distinction of CAR intensity within occupational task execution.

D.1.2 Wages, job levels, and occupational tasks

Given the observed correlation between occupational task contents and job levels from Figure A6, we next ask how much each component contributes to occupational wage differences. In particular, we are interested in understanding whether it is rather the average job level or the task shares that determine observed wage premia *across occupations* in a statistical sense. This complements our previous evidence that *across workers*, job levels rather than occupations determine wage differences and that *across occupation-job-level cells* it is the job levels that are important (Figure 10 and Figure 1).

We run a simple linear regression at the occupation level of log wages on average job levels and task contents of occupations

$$\tilde{w}_i = \alpha + \beta x_i + \sum_c \gamma_c z_{i,c} + \varepsilon_i$$

where $\tilde{w}_i$ is the average log wage in occupation i , x_i is the average job level of occupation i and $z_{i,c}$ are the task shares of occupation i . As task shares sum to 1, i.e. $\sum_{c=1}^5 z_{i,c} = 1$, we drop the cognitive routine share if necessary to avoid collinearity. Table A4 shows the regression results for different specifications of the regression above.

The first striking observation is the high explanatory power of the average job level for inter-occupational wage differences in the first specification (column (1) *only JL*) where we only regress on the average job level of an occupation. The next striking observation is that adding information from the task-based approach (column (2)) adds little to the explanatory power of the regression. If we only consider the task-based approach in column (3), the explanatory power is less than half that of the job levels alone. In terms of coefficients, most notably, manual non-routine (M-NR) tasks have a large negative effect on wages that is highly statistically significant. When we run the different task components in isolation, we find that analytic non-routine and manual non-routine have the highest explanatory power for inter-occupational wage differences. Finally, we note that the point estimate for the average job level remains largely unaffected when we include the information from the task-based approach (columns (1) and (2)). These results corroborate our findings from Section 4.3 and Appendix C of the large explanatory power of job levels on between-occupation wage differences.

D.2 Fifth occupation digit and job levels

The latest revisions of five-digit occupation codes have started to also measure and encode job complexity (Helper/Trained/Specialist/Expert) and whether some management and supervisory duties are associated with the job (ISCO-08 or KldB-2010 for Germany). We observe the latest revision of these occupation codes in the 2018 SES data and compare them against the

Table A4: Wages, tasks, and job levels

	(1) only JL	(2) JL + TBA	(3) only TBA	(4) A-NR	(5) I-NR	(6) M-R	(7) M-NR	(8) C-R
job level	0.47*** (0.00)	0.54*** (0.00)						
A-NR		-0.23** (0.01)	0.25 (0.06)	0.63*** (0.00)				
I-NR		-0.23* (0.02)	-0.35* (0.03)		0.11 (0.50)			
M-R		0.08 (0.35)	-0.27 (0.05)			-0.32** (0.00)		
M-NR		-0.16* (0.04)	-0.51*** (0.00)				-0.58*** (0.00)	
C-R								0.34* (0.01)
<i>N</i>	140	140	140	140	140	140	140	140
adj. R^2	0.72	0.76	0.34	0.26	0.00	0.05	0.23	0.04

Notes: Regression coefficients from regressing mean occupation log wages on average job levels and task-based components. Wage and job level data are aggregated for 140 occupations from 2018 SES data and task-based components are taken from [Dengler et al. \(2014\)](#). For each specification, the number of observations and adjusted R^2 are shown at the bottom of the table, p -values in parentheses, and *, **, *** indicate the significance of coefficients at the 5%, 1%, and 0.1% levels, respectively. See text for further details.

job-level information in these data. Table A5 shows the cross-tabulation of the last digits of the occupational classification system (KldB 2010) of the German employment agency against job-level information in the 2018 SES data. We find a clear positive correlation between the information from the occupation code and the job level, but we also see that there is substantial mass off-diagonal. Although there is a correlation of job levels with the very detailed occupation classification, the correlation is weak. Hence, job levels contain additional information even beyond what the very fine-grained occupational codes contain.

Table A5: Cross-tabulation of job levels measured directly and job levels inferred from occupation codes

Complexity measured by occupation	Fraction of occupation (in %)	Fraction of job level ...within occupation (in %)				
		1	2	3	4	5
All	100	6.5	15.0	51.4	18.4	8.7
from last digit (KldB 2010)						
Helper	14.9	27.7	41.9	27.4	2.1	0.9
Trained	57.2	3.8	14.1	66.5	12.5	3.0
Specialist	14.2	1.0	3.3	44.6	37.9	13.2
Expert	13.7	0.5	1.5	21.3	40.6	36.1
using management occupations (KldB 2010)						
Supervisors	2.5	0.8	3.1	30.5	42.8	22.9
Managers	3.0	0.4	1.9	17.4	34.2	46.1

Notes: Cross-tabulation of job levels and occupation information from the 2018 Structure of Earnings Survey. Occupational information is extracted from five-digit occupational code (KldB 2010). The first part of the table (*last digit*) shows the distribution of workers by occupational complexity across job-level groups. Shares sum to 100 within each row. The first column (*total*) shows the population share of the occupation group. The second part of the table (*management occupations*) shows the distribution of occupations coded as supervisors or managers across job-level groups. Shares sum to 100 within each row. The numbers in the columns refer to the share of workers coded as supervisors or managers in the total population.

E Additional details on job leveling for Germany

In this section, we provide additional details for the analysis of CAR intensity and job-leveling factors in Section 4.3. First, we explain the details of the implementation of the job-leveling scheme that we apply to the BIBB/BAuA data. Second, we provide additional results for blue-collar workers. The analysis in Section 4.3 focuses on white-collar workers. Finally, we compare the wages by job level constructed from the survey data with the actual bargained wages by job level.

E.1 Mapping of job-leveling scheme to survey questions

We use eight questions from the 2012 BIBB/BAuA employment survey to construct job-leveling factors and to implement the ERA job-leveling scheme (Hall et al., 2018). Point values for the ERA job-leveling scheme are taken from the leveling scheme of the bargaining agreement for the steel and metal industry (Germany’s largest industry) in North-Rhine-Westphalia (Germany’s largest state). The collective bargaining agreement is the largest single one in the private sector in terms of workers covered ($\approx 700,000$). The point system can be downloaded in English.⁵² The job-leveling system has four components: required skills and knowledge, autonomy, cooperation

⁵²See METALL NRW: Verband der Metall- und Elektro-Industrie Nordrhein-Westfalen e.V., “Salary Schedule 2010/2012 (ERA),” page 6, “Point System for Evaluating Job Functions” https://metall.nrw/fileadmin/_migrated/content_uploads/Tarifikarte_ERA_2010-2012_englisch_01.pdf (accessed May 22, 2019).

and communication, and supervision. We identify the questions from the BIBB/BAuA survey that we consider to most closely correspond to the different components of the job-leveling system. We use the following eight specific questions for our job-leveling approach:


1. What kind of training is usually required for performing your occupational activity? (four answers)
2. Is a quick briefing sufficient to perform your occupational activity, or is a longer working-in period required? (two answers)
3. How often does it happen in your occupational activity that one and the same work cycle/process is repeated in the minutest details? (four answers)
4. How often does it happen in your occupational activity that you improve existing procedures or try out something new? (four answers)
5. Question on type of task performed (simple, qualified, highly qualified)
6. How often does it happen in your occupational activity that you have to communicate with other people in your occupational activity? (three answers)
7. Do you have colleagues to whom you are the immediate supervisor?
8. And how many are they?

To apply the job-leveling scheme, we have to assign job-level points to answers from the BIBB/BAuA survey. The point range of the job-leveling system is from 10 to 170 points and we apply the following assignment of points. For the skills part, we assign 10 points if a quick briefing is sufficient and no vocational training is necessary to execute the tasks and duties of the worker's current job. We assign 30 points if a longer working-in period is required but still no vocational training, 50 points if the job requires apprenticeship training, 80 points if the job needs a master craftsperson or technician certificate, and 100 points if the job requires a university or technical college degree. Note that the requirements are typical minimum requirements and do not imply that only workers with such skill level work in jobs with these job levels. We further assign 6 points if the job involves complex/qualified tasks and 12 points if it involves highly complex/qualified tasks. For autonomy, we assign 2 points if the same work cycle is repeated in detail often, 10 points if this is sometimes the case, and 18 points if this is rarely the case. For jobs where the same activity is never repeated, we assign 30 points if it is a complex/qualified job and 40 points if it is a highly complex/qualified job. For communication and cooperation, we assign 2 points if the job requires no communication with other people, 4 points if this is sometimes the case, and 10 points if this is often the case but the job rarely or never requires improving on existing procedures or trying something new. We assign 15 points if the job requires communicating often and sometimes requires improving on existing procedures, and we assign 20 points if it is often the case that the job requires improving on existing procedures or trying something new. Finally, for responsibility, we assign 10 points if the job includes supervisory duties and 10 additional points if the job involves supervising more than 20 other workers. We sum these job-level points to the total job-level points for each observation in the data. We refer to the sum of points to *job-level points from the ERA scheme*.

E.2 Results for blue-collar workers

In Section 4.4, we restricted the sample to white-collar workers, Figure A7(a) reports corresponding results for blue-collar workers. We report separate results for white- and blue-collar workers because of different job complexity variables. After implementing the job-leveling scheme for blue-collar workers, we also find a close alignment between the statistical scheme (Section 4.4) and the assigned job-level points from the ERA scheme. We also see there are fewer blue-collar workers in the data, so estimates are less precise.

Figure A7: job-level points and average wages by job-level points (blue-collar workers)




Notes: Left: Scatter plot of a worker's implied job-level points from statistical job-leveling scheme against the worker's job-level points from union bargaining scheme (ERA scheme). The statistical job-level scheme is based on the regression of wages on survey answers. The solid line shows the linear fit and the legend reports R^2 . Right: Distribution of wages by job level (groups of 5 points to reduce sampling noise). Job-level points have been constructed from survey questions on job characteristics (see text for details).

Figure A7(b) visualizes the distribution of wages for each job-level point (in groups of 5 points each). We find variation in wages at each point level, but the variation across job levels dominates the variation within job levels. For blue-collar workers, the variation across job-level points is somewhat smaller, but there is still a clearly positive relation between wages and job-level points.

E.3 Job leveling and bargained wages

Finally, we explore how well our implementation of the point-leveling scheme aligns with reported wages from the union bargaining contract. For this, we focus on workers from North-Rhine-Westphalia in the BIBB/BAuA data and compare their average wages by point level to the reported wages by job level from the union bargaining agreement for steel- and metalworkers. Figure A8 shows wages from the BIBB/BAuA data by point level together with wages taken from the union bargaining agreement. Overall, we find a good fit between wages by job levels from the microdata in comparison to the wages from the union bargaining agreement. The BIBB/BAuA data are for 2012 and also include workers not covered by a union bargaining contract and not working in the steel and metal industry. The data for wages from the union bargaining contract are for 2018 and have been adjusted for inflation and average real wage growth. The close fit suggests that our implementation based on the selected survey questions provides a close approximation to how base wages of workers are set in practice.

Figure A8: Average and bargained wages by job-level points for North-Rhine-Westphalia (blue-collar workers)



Notes: Average (log) wages by job-level points and bargained wages for steel- and metalworkers. Workers in BIBB/BAuA data from North-Rhine-Westphalia. Bargained wages for steel- and metalworkers for North-Rhine-Westphalia for 2018 have been adjusted to 2012 euros for CPI and average real wage growth. Job-level points have been constructed from survey questions on job characteristics. The lines represent the average log wage for the job-level points (in groups of 5 points).

F Social Security Data for model calibration

For the model calibration, we rely on administrative data from German social security (pension) records, the *Versichertenkontenstichprobe* (VSKT 2020) [der Rentenversicherung](#) (FDZ-RV). [Bayer et al. \(2025\)](#) provide a detailed description of the data. In the data, we observe entire employment history of workers starting at 14 years of age with monthly earnings and labor market status. The VSKT data come as 2% annual samples of employment histories for workers with a social security record and from birth cohorts that are at the time of sampling between 15 and 70 years old. The scientific use file contains data on 187,113 individuals and weights to preserve representativeness of the data. Earnings in the data are reported as pension points. These pension points express monthly earnings relative to macroeconomic annual average earnings, an index updated since 1957 based on gross earnings data from the national accounts.

We restrict the data to male workers whose labor market histories between 25 and 55 years of age are covered by the data. We exclude workers with employment periods in East Germany. We require that workers have labor market information in 240 out of the 360 labor market history months (excluding workers who became public servants or self-employed). We construct permanent income as total accumulated pension points from employment divided by the number of employment and unemployment months, and group workers into quintiles based on their permanent income. Using the panel dimension of the data, we compute the average separation rate of workers in the sample and for each permanent income group, we compute the average unemployment rate. For the separation rate, we require that the worker spends most of the quarter including the last employment spell in unemployment. We derive the quarterly separation rate of 0.93% by multiplying the monthly separation rate by a factor of 3. Table A6 reports the unemployment rate by permanent income quintile. We use the ratio of the bottom quintile to top quintile unemployment rate in our calibration ($16.40\%/0.68\% = 24.3$).

Table A6: Unemployment Rates by Permanent Income Quintile

Permanent Income Quintile	Unemployment Rate
1 (Bottom)	16.40%
2	3.40%
3	1.66%
4	1.43%
5 (Top)	0.68%

G Identification and instrumental variable regression


G.1 Identification challenges from other theories of career progression

Our empirical analysis follows from the newly develop model in the paper. It also addresses two identification challenges that are motivated by other existing models of career progression. The first challenge results from the seminal work by [Waldman \(1984\)](#) and refined by [Gibbons and Waldman \(2006\)](#). In this model framework, employers learn about workers’ abilities and promote good (highly productive) workers to jobs with potentially higher skill requirements and higher skill complementarity. High wages are the means by which employers prevent other employers from poaching their highly productive workers. A worker’s productivity is the key determinant of wages and high-paying jobs are only a signal that the jobholder is a highly productive worker. According to this view, each job is set up based on the skills of the individual worker. We address the arising challenge that unobserved individual heterogeneity is accounting for the wage differences across job levels in three ways. First, we aggregate the data to the cohort level so that we exploit only the differential distribution of cohorts across job levels for identification. Second, these cohorts might still differ in their (average) individual fixed effects and career progression. Controlling for fixed effects in our panel regression removes this challenge for identification. Third, we apply an instrumental variable approach relying on a Bartik-style instrument ([Bartik, 1993](#)) based on shifts in industry composition over time. In the next section, we provide details on how we construct the instrument and discuss estimation results. In Appendix H.5, we report results when not including fixed effects to control for individual heterogeneity. These results are consistent with an omitted variable bias as described. The second challenge for identification arises from the mechanism highlighted in the seminal paper by [Lazear and Rosen \(1981\)](#). [Lazear and Rosen](#) provide an alternative view on career progression that interprets promotions as the outcome of a tournament. Considering jobs and the associated wages as prizes implies that wages only represent a prize for previous performance, but not remuneration for task execution on the current job. If wages are prizes, differences in a job’s tasks will not be systematically related to wages. In Section 4.4, we provide evidence based on the BIBB/BAuA data that differences in task execution are systematically related to wages and that job levels can be constructed from the CAR intensity in task execution. This finding supports the task-based approach that postulates that the executed tasks determine a worker’s wage. Importantly, this evidence does not rule out that residual wage differences result from performance-related pay.

G.2 Details on instrumental variable regression

The instrumental variable approach addresses the concern that differences in the organizational structure and job composition across cohorts that we use for identification in our baseline approach could be endogenous to the composition of workers in these cohorts. To address this potential endogeneity problem, we instrument job levels using a Bartik-type instrument (Bartik, 1993). To construct our instrument for the job-level component, we only exploit changes in the industry composition over time. Based on the average job composition of an industry over the entire sample period, we construct the predicted occupation and job-level composition for each cohort at each moment in time. We then estimate the synthetic cohort approach by applying these instruments. We proceed with the decomposition of wage growth and wage dispersion over the life cycle as in the baseline case. Figure A9 shows the resulting decomposition results for wage growth and wage inequality for men (Figures A9(a) and A9(c)) and for women (Figures A9(b) and A9(d)).

Figure A9: Decomposition of wage growth and wage dispersion over the life cycle using IV approach



Notes: Contribution of the job component to wage growth (top row) and wage dispersion (bottom row) for men (left panels) and women (right panels). The solid line shows the job component for the baseline from the main part of the paper; the short dashed line shows the case with no collective bargaining interaction; the dotted line shows the case with full-time interaction; and the dash-dotted line shows the case with large firm interaction. Job components have been constructed by setting all dummy variables in the interaction terms to one. As in the main text, all graphs show the coefficients of age dummies of a regression of the components on a full set of age and cohort dummies (ages defined as three-year groups).

In the decomposition of wage growth, we find that for both men (Figure A9(a)) and women

(Figure A9(b)), the relative importance of the job component remains unchanged, while the individual component decreases and the plant component increases in its relative importance. In the decomposition of the increase in wage inequality, the results become even more striking than in our baseline approach. We find that for both men (Figure A9(c)) and women (Figure A9(d)), the relative importance of the job component increases. For men, the job component even exceeds the total increase of the variance when we do not account for the covariances. In the case of women, the contribution of the job component tracks the overall increase almost one-for-one. These results demonstrate that the results of our baseline approach for the job component are robust to the potential endogeneity problem for the organizational structure and job composition of plants.

H Sensitivity analysis, extensions, and further results

In this section, we provide several sensitivity checks to our baseline analysis from the main part of the paper. In the sensitivity checks, we first show results for the wage decomposition without job component (Section H.4). We then explore the effects of not being covered by a collective bargaining agreement, considering only full-time work, and focusing on large establishments. We also show the results if we do not drop public employers from the sample or do not control for individual fixed effects using the synthetic panel regression. We discuss these results in Section H.1. As extensions to our baseline results, Sections H.2 and H.3 explore more flexible specifications for the wage equation. Section H.5 reports results if instead of a synthetic cohort panel approach, we rely on a pooled OLS regression when decomposing wages.

H.1 Heterogeneous returns to job and individual characteristics

For the first set of sensitivity checks, we interact variables from the baseline regression (26) with dummy variables for not being covered by a collective bargaining agreement, for working full-time, and for working in a large establishment. We also report results for a sensitivity analysis in which we do not drop observations from public employers and publicly controlled firms. In columns 1 to 4 of Table A7, we compare the baseline sample to the part of the sample that gets a positive dummy in the sensitivity analysis. Overall, there are differences in the job-level composition in the alternative groups compared to the baseline sample, but they are not striking. The last column of Table A7 shows characteristics of workers and jobs at public employers that we drop for the baseline analysis. Two observations are noteworthy for this sample of public employers. First, the share of women is large: 60% of employees at public employers are women. Second, the job composition at public employers has fewer jobs at job levels 1 to 3 but more jobs at the two top job levels.

In the first step, we consider the sensitivity analysis with respect to collective bargaining agreements, full-time workers, and large establishments and test whether the estimated coefficients on the additional interaction terms are statistically significant. Table A8 shows test statistics for three tests for the three different interaction specifications. The first row jointly tests all interaction coefficients. We find that insignificance can always be strongly rejected.

Table A7: Summary Statistics

	baseline	no collective bargaining	only full-time	large plants	public employers
wage	19.3	18.0	20.3	22.4	20.0
age	41.1	40.6	40.8	41.3	41.9
female	39.0	37.9	27.1	37.6	60.2
1	8.1	7.0	6.2	7.6	4.9
2	15.9	18.2	14.7	13.8	7.3
3	45.9	50.0	45.5	41.2	39.1
4	21.4	17.6	23.5	25.5	27.5
5	8.7	7.2	10.2	11.8	21.3
N (million)	2.7	1.5	2.1	1.0	0.6

Notes: Descriptive statistics of sample composition for baseline sample and subsamples considered in sensitivity analysis. The rows *wage* and *age* refer to the sample averages. The row *female* refers to the share of females in the sample; Rows labeled 1 to 5 show the shares for workers at the different job levels in the samples; and *N* is the number of observations in millions of the different samples.

This finding means that potentially there is a layer of heterogeneity that is deeper than what our baseline treatment explores. Yet, the test results in Table A8 only talk about statistical, not economic, significance. The same careers (e.g., across job levels and occupations) can potentially mean something different when the coefficients (i.e., the returns to occupation and job level) are much different for full-time workers or workers not covered by collective bargaining.

Given the importance of the job component, we focus here on the changes in the job component when discussing the economic significance and sensitivity of our results. Figures A10(a) and A10(b) show the job component from the baseline specification together with the specifications from the different sensitivity specifications (no collective bargaining, full-time, large plants). We show the case in which we keep the evolution of the characteristics of jobs over the workers' life cycle as in the baseline sample but treat them with the wage schedule for the subgroup for which we estimated the interaction terms. That is, we ask, what would the wage profile of workers look like if all workers got non-collectively bargained wages? Of course, this assumes that neither the career paths nor the wage schedule of non-collectively bargained wages would change when there is no collective bargaining. This has to be taken into account when comparing the different job components.⁵³ Similarly, Figures A10(c) and A10(d) show the contribution of the job component to the increase in the variance of log wages over the life cycle for the baseline and the different sensitivity specifications using the same technique. In contrast to the presentation in the main part of the paper, we removed level differences at age 25 for easier comparison.

Looking first at the case of no collective bargaining, we find the age-wage profile (for the job

⁵³This assumes that there are no equilibrium effects on the organizational structure if there are, for example, only plants without collective bargaining agreements in the market.

Table A8: Test statistics for coefficient tests

	no collective bargaining		only full-time		large plants	
	<i>p-value</i>	<i>F-stat</i>	<i>p-value</i>	<i>F-stat</i>	<i>p-value</i>	<i>F-stat</i>
all	0.00	2.4	0.00	3.2	0.00	1.6
individual	0.00	3.0	0.00	2.2	0.00	2.3
job	0.00	2.5	0.00	3.0	0.02	1.5
job level	0.00	8.6	0.00	4.3	0.01	3.4


Notes: Test statistics for joint significance of interaction coefficients with wage component coefficients. Row *all* shows test results for joint significance of all interaction terms, row *individual* shows test statistics for coefficients of individual component, row *job* shows test statistics for coefficients of job component, and row *job level* shows test statistics for the joint significance of the job-level interaction dummies. See text for further details.

component) would look steeper if no worker had collectively bargained wages. When looking at variances, we also find that job-level returns in wages are more diverse when the worker is not covered by a collective bargaining agreement, so without collectively bargained wages, wage dispersion would increase much more over the life cycle. This reflects the fact that there is wage compression in collectively bargained wages (Appendix C.3). When looking at large plants, we find results that are opposite to *no collective bargaining*. Wage growth profiles are less steep, and wage dispersion increases less. The likelihood is that these plants have a larger fraction of workers with collectively bargained wages.

The effect of working full-time is negligible for wage growth and for the increase of the variance, we get a slightly stronger increase for women and the same increase for men. Here, it is important to note that we keep the distribution of workers across job levels unchanged and only change the estimated job-level wage. Importantly, this result is consistent with our model-based analysis of the gender wage gap as it is *qualitatively* different. Here, we change job-level wages, but keep the distribution over job levels the same.


Figure A11 shows the effects of including public employers in the baseline sample. We perform the same decomposition for the larger sample that includes workers at public employers as in the baseline analysis and compare the results for the job component to the baseline sample. Effects for men are negligible. The more notable effect is for women. Including public employers adds slightly less than a third to the job component for female wage growth. This finding suggests that public employers are an important contributor to female career progression after age 35 and that women seem to select public-employer careers. The results including public employers further suggest that there is substantially less dispersion in career progression at public employers. The increase from the job component for women is substantially smaller once we include public employers in our sample.

Figure A10: Contribution of job component to wage growth and wage dispersion over the life cycle



Notes: Contribution of the job component to wage growth (top row) and wage dispersion (bottom row) for men (left panels) and women (right panels). The solid line shows the job component for the baseline from the main part of the paper; the short dashed line shows the case with no collective bargaining interaction; the dotted line shows the case with full-time interaction; and the dash-dotted line shows the case with large firm interaction. Job components have been constructed by setting all dummy variables in the interaction terms to one. As in the main text, all graphs show the coefficients of age dummies of a regression of the components on a full set of age and cohort dummies (ages defined as three-year groups).

Figure A11: Contribution of job component to wage growth and wage dispersion at public employers




Notes: Contribution of the job component to wage growth (top row) and wage dispersion (bottom row) for men (left panels) and women (right panels). The solid line shows the job component for the baseline from the main part of the paper; the dashed line shows results for a sample including public employers and publicly controlled firms. As in the main text, all graphs show the coefficients of age dummies of a regression of the components on a full set of age and cohort dummies (ages defined as three-year groups).

H.2 Education-specific returns to experience

Heterogeneity in returns to experience has been proposed as an explanation for the higher wage growth of better-educated workers (Gibbons et al., 2006). In our baseline regression, we allow for differences in experience only between men and women but not across education groups, so that it could be the case that heterogeneity in returns to experience across education groups gets absorbed by the job component as better-educated workers are also more often found further up on the career ladder (Section 5.4.2). To explore this possibility, we augment our baseline regression by adding linear education-specific experience profiles. In the decomposition, we attribute these education-specific experience components to the individual component. We decompose life-cycle wage growth and the increase in the variance as in the baseline case. Figure A12 shows the decomposition of life-cycle wage dynamics for men and women for this extended wage regression.

Figure A12: Life-cycle wage dynamics with education-specific slopes




Notes: Top left panel: Decomposition of log wage differences by age relative to age 25 for male workers. The dashed line corresponds to the individual, the dotted line to the plant, and the dash-dotted line to the job component; the solid line (total) equals the sum of the three components. The horizontal axis shows age, and the vertical axis shows the log wage difference. Bottom left panel shows the same decomposition for female workers. Top right panel: Decomposition of the variance of log wages by age for male workers. Variances of all components are calculated by age-cohort cell. The solid line is the variance of total wage, the dashed line is the individual, the dotted line is the plant, and the dash-dotted line is the job component. Bottom right panel shows the same decomposition for female workers.

We find our decomposition results to be very similar under this extended wage specification. For wage growth in Figures 13(a) and 13(c), the job component declines slightly for men and

women but remains for men by far the most important driver of wage growth. For women, now, all three components account for a third of wage growth at the end of working life. For men and women, the plant component gains slightly in importance. For men, it becomes more important than the individual component. For women, we observe a convergence of the three components. For the increase in wage inequality in Figures 13(b) and 13(d), we find, if anything, that the contribution of the job component increases. Ignoring covariance terms, the variance of the job component alone accounts virtually for the entire increase in wage inequality over the life cycle for both men and women. Hence, we do not find evidence that the job component is inflated by picking up an education-specific skill accumulation effect.

H.3 Occupation-specific returns to experience

Figure A13: Decomposition of life-cycle wage dynamics with occupation-specific experience



Notes: Top left panel: Decomposition of log wage differences by age relative to age 25 for male workers. The dashed line corresponds to the individual, the dotted line to the plant, and the dash-dotted line to the job component; the solid line (total) equals the sum over the three components. The horizontal axis shows age, and the vertical axis shows the log wage difference. Bottom left panel shows the same decomposition for female workers. Top right panel: Decomposition of the variance of log wages by age for male workers. Variances of all components are calculated by age-cohort cell. The solid line is variance of total wage, dashed line the individual, dotted line the plant, and dash-dotted line the job component. Bottom right panel shows the same decomposition for female workers.

The individual component in the baseline specification only includes general experience, but it could be that returns to experience differ across occupations and might in our baseline specification be absorbed by the occupation dummies that go into the job component. To explore

this possibility, we augment the baseline specification by occupation-specific experience profiles that we specify as occupation-specific linear experience profiles. These occupation-experience interaction terms cannot be unambiguously assigned to one of the three components as they interact with variables from the individual and job component. To be conservative for the job component, we include the interaction terms in the individual component for the decomposition. We proceed otherwise as in the baseline decomposition. Figure A13 shows the decomposition of life-cycle wage dynamics for men and women for this specification.

Looking at wage growth in Figures 14(a) and 14(c), we find that, as in the baseline regression, the contribution of the job component accounts for more than 50% of wage growth for men and a third for women. The contribution of the individual and plant components remains largely unaffected for men. For women, we find an increase of the individual component. For the increase in wage inequality in Figures 14(b) and 14(d), we find no notable effect for the contribution of the job component as the key driver of rising wage inequality. The individual components for men and women become U-shaped over the life cycle with a countervailing effect (not shown) from the individual-job covariance term.


H.4 Wage decomposition without job component

To explore the importance of the job component for the wage decomposition, we repeat the wage decomposition from (26), but drop the job component $\beta_J J_{ipt}$. We then compare the resulting plant and individual components to those from our baseline decomposition.

Figure A14 shows the decomposition of life-cycle wage dynamics for men. Comparing the decomposition of wage growth in panel (a) to our baseline in Figure 10, we draw qualitatively very different conclusions about the sources of life-cycle wage growth. We find that the experience effect of the individual component accounts now for a substantially larger part of wage growth. Experience acts as the residual of the wage growth decomposition and its rising importance implies that a larger part of wage growth remains unexplained when we do not include the job component in the decomposition. Figure A14(b) contrasts the contribution to wage growth at age 55 from the individual and plant components in the baseline decomposition to the decomposition without the job component. The comparison highlights the striking increase of the individual component becoming more than twice as large if differences in job levels are not accounted for. Ignoring the job component also implies that we abstract from differences in the organizational structure across plants what increases the importance of the plant component in the decomposition.

Figure A14(c) shows the decomposition of rising wage inequality over the life cycle when we drop the job component from the decomposition. Qualitatively, we get similar results to the wage growth decomposition. We find that after dropping the job component a large fraction of wage differences are no longer accounted for and that this fraction grows over the life cycle. The implied rising dispersion of the residual wage component has been traditionally interpreted as persistent labor market risk. Our new finding that most of this dispersion is accounted for by career ladder dynamics does not invalidate the interpretation as risk—indeed, our model relies on this interpretation—but opens up new opportunities to understand the sources and drivers of

Figure A14: Decomposition of wage growth and variance of wages by age (men), ignoring job controls




Notes: Top panels show the decomposition of male wage growth in the individual and plant components. The bottom panels show the corresponding decomposition of wage variances for men. The left panels show the life-cycle profiles when estimating the components without job controls. The right panels compare the components at age 55 to the baseline decomposition that includes job controls (job components not shown here).

labor market risk. Finally, Figure A14(d) shows the contribution to the level of wage inequality at age 55 accounted for by the individual and plant components. We find that individual and plant components both become more important for the level of wage inequality but that only the individual component accounts for a sizable increase of wage inequality over the life cycle.

Figure A15 reports the results for the life-cycle wage dynamics for women and the changes in the individual and plant component compared to the baseline decomposition. Looking at the decomposition results for wage growth in Figures A15(a) and A15(b), we draw generally the same conclusions as from the corresponding decomposition for men. The individual component for wage growth picks up almost all wage growth. For the variance in Figures A15(c) and A15(d), we also get that as for men the individual component increases and accounts now for a sizable fraction of the life-cycle increase in wage inequality.

These results provide important new insights for why employers are important for life-cycle wage dynamics as they point to an important mediating role of job levels also for between-employer wage differences. Recall that the plant component in our baseline decomposition captures whether plants pay better *at all* job levels; that is, the plant component in our baseline decomposition is not driven by having a larger share of top-level jobs or high-wage occupations at

Figure A15: Decomposition of wage growth and variance of wages by age (women), ignoring job controls




Notes: The top panels show the decomposition of female wage growth in individual and plant components. The bottom panels show the corresponding decomposition of wage variances for women. The left panels show the life-cycle profiles when estimating the components without job controls. The right panels compare the components at age 55 to the baseline decomposition that includes job controls (job components not shown here).

the plant. In this baseline decomposition, we find that the organizational structure of plants, i.e. the distribution of CAR intensity of jobs within the plant, and the associated career dynamics account for half of the wage growth and virtually all of the increase in wage inequality so that observed between-employer wage differences are primarily a result of differences in organizational structure and associated differences in career opportunities. That some employers pay everyone better or worse irrespective of the tasks and their execution accounts only for the much smaller part of between-employer wage differences in the baseline decomposition. Thus the result of an increased importance of the plant component when dropping the job component from the decomposition point to an important mediating role of job levels for between-employer wage differences.

This mediating role also provides an explanation to reconcile our results with the recent evidence for Germany and the United States that finds employer-wage differences to be a key driver of increasing wage inequality over time (Card et al., 2013; Song et al., 2019). If in the cross section plants differ in their organizational structure, but if this organizational structure remains unobserved then plants with many high-level jobs will appear to be “high-wage plants”. Hence, a correlation between organizational structure and the estimated plant component in our baseline

decomposition will lead to a more important role of between-employer wage differences when the job component is left out.

Figure A16: Shares of employees by job level and plant component



Notes: The figure shows the share of workers by job levels in plants with below- or above-median estimated plant component $\tilde{\zeta}_p$. The median is defined on a worker basis. 66% of all plants have a below-median plant component.

Indeed, the SES data provide evidence for such an important mediating role of job levels for understanding between-employer wage differences. Figure A16 shows the distribution of workers across job levels for plants sorted by the estimated plant component $\tilde{\zeta}_p$ from the baseline decomposition. High-wage plants with an above-median plant component offer, on average, more jobs at higher job levels (levels 4 and 5). More than one in three workers is in the top two job levels, while in the bottom half of plants, only one in four jobs has a CAR intensity that places it in the top two job levels. Conversely, the organization of the production process provides a much larger share of jobs with low CAR intensity in low-wage plants. More than one job in four is in the lowest two job levels. This result is consistent with the findings of [Tåg et al. \(2016\)](#) for Sweden.

H.5 Pooled regression without individual fixed effects


The main part of the paper uses synthetic cohorts to control for individual fixed effects that are arguably correlated with education, career progression, and potentially employer types. In this section, we run as an alternative specification a pooled OLS regression controlling for cohort effects but not controlling for individual fixed effects. Specifically, we set $\hat{\gamma}_i = \gamma_c$ in (25) and instead run the following regression on the pooled data:

$$\hat{w}_{it} = \gamma_c + \beta_J \hat{J}_{it} + \beta_I \hat{I}_{it} + \hat{\epsilon}_{it}. \quad (28)$$

We proceed otherwise as described in the main part of the paper and use the same control variables for the job component J_{it} and individual component I_{it} . We again demean at the plant level to construct $\hat{J}_{it}$ and $\hat{I}_{it}$. Figure A17 shows the decomposition of wage growth in the

individual, plant, and job components if we do not control for individual fixed effects.

Figure A17: Wage decomposition for men and women without controlling for individual fixed effects



Notes: Decomposition of log wage differences by age relative to age 25 for male (left panel) and female (right panel) workers. Decomposition based on regression without controls for individual fixed effects. The dashed line corresponds to the individual, the dotted line to the plant, and the dashed-dotted line to the job component; the solid line (total) equals the sum over the three components. The horizontal axis shows age, and the vertical axis shows the log wage difference. As in the main text, all graphs show the coefficients of age dummies of a regression of the components on a full set of age and cohort dummies (ages defined as three-year groups).

Comparing the decomposition results for wage growth to the baseline results in Figure 10 shows that the finding of a key role of the job component for wage growth over the life cycle is robust. We find that for both men and women, now all three components contribute roughly equally to wage growth. If individual fixed effects are important for labor market outcomes, we should expect that estimated coefficients change from omitting this control variable from the regression. We interpret the sizable effects on the wage components as an omitted variable bias from the individual fixed effect. The result that the job component is the driver of the increase in wage dispersion is also robust to omitting controls for individual fixed effects. We find that in the decomposition of the increase in wage dispersion, the contribution of covariance between the plant and job components becomes more important. We attribute these differences to the omitted individual fixed effect and do not report the results here. These results are available from the authors upon request.