Job levels and Wages*

Christian Bayer† Moritz Kuhn‡

March 20, 2025

Abstract

Job levels summarize the complexity, autonomy, and responsibility of task execution. Conceptually, job levels are related to the organization of production, are distinct from occupations, and can be constructed from data on task execution. We highlight their empirical role in matched employer-employee data for life-cycle wage dynamics, refine a task-based view of wage determination, and demonstrate that differences in job levels account for most of the observed wage differences. We also show, within a structural framework, that a job-level perspective provides a novel and fruitful interpretation of widely studied phenomena such as the gender wage gap and the returns to education and seniority.

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, Iourii 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 (Kuhn) as well as under CRC-TR-224 EPoS (Bayer and Kuhn). 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, CEPR, and IZA, Adenauerallee 24-42, 53113 Bonn, Germany, *email:* christian.bayer@uni-bonn.de

‡Universität Bonn, ECONtribute, CEPR, and IZA, Adenauerallee 24-42, 53113 Bonn, Germany, *email:* mokuhn@uni-bonn.de

1 Introduction

The question of what determines a worker's wage is a long-standing question of economic research. In this paper, we provide a new perspective on this topic using recent and largely unexplored data. We build on the task-based view of wage determination (Autor et al., 2003), which posits that the tasks that job holders perform determine their wages. However, we take this concept further by also considering how the tasks are executed, referred to as the *job level*. The job level encompasses the complexity of the task, the autonomy in performing it, and the responsibilities associated with it, and allows for differentiation within an occupational group, as well as comparability across occupational groups with different tasks.

The concept of job levels has a long history in labor market statistics, dating back to the 1950s, and often forms the basis of union collective bargaining agreements and corporate job-based compensation schemes.[1] The pervasive use of job levels as an organizing concept in official labor statistics, their key role in union collective bargaining, and an existing job-leveling industry provide strongly suggestive evidence of their conceptual importance. However, data limitations regarding the availability of job-level information in most datasets have prevented a systematic analysis of the relationship between job levels and wages. The goal of this paper is to use high-quality, administrative, matched employer-employee data to provide evidence of the importance of job levels in determining wage differences and to argue that understanding the economic content of job levels is crucial for a comprehensive understanding of the labor market.

What is the *job level* of a job? Occupations describe what tasks a worker executes in her job and have been widely studied in economic research (e.g. Kambourov and Manovskii, 2009a; Autor et al., 2003). Job levels provide an additional distinction of task execution within and across occupations in terms of complexity, autonomy, and responsibility (CAR). For the simplest example of differences in job levels, consider two bakers: One baker follows recipes and rules for mixing and baking dough, and the other baker also mixes ingredients and bakes dough, but develops new recipes. Both perform the occupational tasks of bakers, but their autonomy and responsibility in executing the tasks differ, and thus their job levels will differ. Importantly, job levels are designed to be independent of specific tasks so that they provide comparability across occupations. The complexity dimension of job levels defines minimum skill requirements for the job holder, in line with the idea of the task-based approach that employers need to match workers' skills to jobs and associated tasks (Acemoglu and Autor, 2011). In contrast, autonomy and responsibility are closely related to the organization of the production process as they describe the organization of work. Thus, our findings confirm a conjecture of Acemoglu and Autor (2011) that task execution is ultimately related to *"(...) the allocation of authority within the organization (...) and the nature of the responsibility system"* (p.84).

[1] 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, and the German Statistical Office has reported wages by job level since at least 1957. In addition, many unions and firms use job leveling schemes. An example of a union wage contract is the agreement of the steel and metal workers in North Rhine-Westphalia, which we will study below, and the Korn Ferry Hay point system, which is widely used in job-based compensation schemes.

Our study of job levels and wages consists of three steps. First, we explain the concept of job levels in detail, demonstrate their economic content using additional survey data, and explain how they differ from occupations and traditional measures of job tasks. Second, we decompose life-cycle wage growth and inequality and isolate the role of changing job levels in accounting for life-cycle wage dynamics. Third, we provide a search model of career dynamics and use it to provide a new perspective on the gender wage gap, returns to education, and returns to seniority as a result of differences in career progression.

Our analysis is based on four waves of the German *Structure of Earnings Survey* (SES) covering the period from 2006 to 2018. Each wave provides worker-level information on job levels and wages. In the data, job levels alone account for 47% and all observable characteristics for over 80% of the variation in wages. To establish that job levels have economic content, we complement the SES data with data on job requirements from the BIBB/BAuA Employment Survey. In these data, we construct job-leveling factors from workers' reported job requirements using a publicly available job-leveling scheme from a large collective bargaining agreement. These job-leveling factors are the building blocks for constructing job levels and we find these factors account for 44% of the wage dispersion in the BIBB/BAuA data. Our analysis also shows that job levels are different from occupational classifications and that job levels account for a substantial portion of the wage dispersion within and across occupations. Within an occupation, we typically find a significant share of workers in three (out of five possible) different job levels. Furthermore, we show that task-based wage differences (Autor et al., 2003) are largely absorbed by average job-level differences, so that task-based (occupational) wage components alone account for little of the wage dispersion when job levels are controlled for. Thus, while job levels confirm the general idea of the task-based approach, they also provide an important refinement. We generalize our findings beyond the German case using data from the National Compensation Survey (NCS) conducted by the Bureau of Labor Statistics (BLS) in the United States.[2] We find very similar results for German and U.S. data and conclude that the importance of job levels for macroeconomic wage differences is not a peculiarity of the German labor market.

In a second step, we decompose wage dynamics over the life cycle by applying synthetic panel methods (Deaton, 1985; Verbeek, 2008), a standard tool from the macroeconomic toolkit, to the repeated cross-sections of the SES data. We construct cohort-level panel data and estimate the coefficients of interest based on a cohort-level wage regression. We use the estimated coefficients to construct the worker-level wage components arising from observable *individual* and *job* characteristics plus an *employer* component. Using this decomposition, we document how much each component contributes to wage growth and increasing wage dispersion over the life cycle, with separate decompositions for men and women. We further decompose the contribution of job characteristics into a job-level component and an occupational component and find that the former accounts for the vast majority of the job component. We find that for life-cycle wage dynamics, career progression, i.e., transitions across job levels as workers age, accounts for 50%

[2]The 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).

of wage growth and almost all of the increase in wage dispersion over the life cycle. While 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.

Motivated by our findings on the importance of job levels in determining wages, we develop a labor market search model of career dynamics. The model assumes that firms have jobs at different job levels and that workers search for employment opportunities in a frictional labor market. Career opportunities arise when workers change employers or when workers leave their current employer. We calibrate the model to our new empirical evidence on job-level wages and average worker mobility between employment and unemployment. We find that the model's predictions of life-cycle wage dynamics closely match the observations from our empirical analysis. We use the model to provide new structural interpretations of widely studied wage phenomena such as the gender wage gap, the returns to education, and the returns to seniority. We trace the estimated gender wage gap to differences in life-cycle career progression so that we interpret the gender wage gap as a gender promotion gap after transitory periods during which female workers reduce their labor market mobility. The result is that female workers fall behind their male counterparts in terms of career progression. Similarly, returns to education are linked to better-educated workers being promoted more quickly to higher job levels. In the model, as in the data, such dynamic career differences account for average wage differences between college and non-college workers. Finally, we show again in the model and in the data that returns to seniority (beyond age and experience) are largely mediated by advantages in promotions to higher job levels.

To put our results in perspective, we note that coded job levels include more than our simple baker example above suggests. Job levels combine several aspects of the execution of tasks (see the BLS Job Leveling Guide US Bureau of Labor Statistics, 2013, for the US NCS data). One of these aspects is that some jobs have a (particularly) complex set of tasks, which is reflected in some minimum skill requirements. However, a minimum skill requirement still 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. This fact relates job levels to human capital because it provides a notion of human capital utilization on the job, e.g., a taxi driver with a college education would not be using all of her human capital (see Rosen, 1983, for consequences of potential underutilization of human capital for educational choices). Indeed, we find in our model that college and non-college workers hold jobs across the job level spectrum, which is consistent with the empirical evidence. It also highlights the underlying idea of wage determination, namely that workers are paid for the tasks they perform rather than for their stock of human capital. This idea is not new, nor original to the concept of job levels, but, as Acemoglu and Autor (2012) note, a key innovation of the task-based approach. They 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 and corroborate this view. 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 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 with 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. (2023).

Finally, it is important to emphasize that we do not answer the question of why workers end up in the jobs they have and why some climb the career ladder while others do not. In this sense, we study the consequences rather than the causes of career progression. We still provide descriptive evidence on career dynamics over the life cycle by complementing our cross-sectional analysis with panel evidence from the German Socio-Economic Panel (SOEP), where we observe a proxy for job levels. We document life-cycle profiles of career ladder promotion and demotion rates and explore how labor market mobility across employers and through non-employment is associated with steps up and down the career ladder. We find that employer mobility is associated with career progression, but that most career ladder moves occur while staying with the same employer.

The remainder of the paper is organized as follows: Section 2 relates our work to the existing literature. Section 4 introduces the data and provides a detailed discussion of job levels, their economic content, their relationship to occupations, and to the task-based approach. Section 5 introduces the decomposition approach used to study life-cycle wage dynamics, reports the decomposition results, and presents the results on life-cycle promotion and demotion dynamics. Section 6 introduces the career progression model and offers a new perspective on wage dynamics. We conclude with Section 7. An appendix follows.

2 Related literature

In this paper, we identify job levels as the most important determinant of observed wage differences across workers and over the life cycle. 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 fit 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 to rationalize 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 connection with 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. (2018) 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 (2019) 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. (2022) document a relationship in targeted survey data between the coordination in the production process and average worker pay.

By exploring the sources of life-cycle wage growth and inequality, 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.[3] We add 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 an important role of occupations as a determinant of wage differences in the cross-section and over time. Our results complement this work adding the importance of job-level differences to occupational or task differences. Employer differences as a source of wage differences feature prominently in the strand of 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 rising between-employer wage differentials are an important contributor to rising wage inequality in Germany. Song et al. (2015) confirm this finding in US Social Security data. Song et al. (2015) and Card et al. (2013) both argue that changes in the organizational structure of firms are the likely driver of rising between-firm wage differentials. Low et al. (2010), Hornstein et al. (2011), or Jung and Kuhn (2016) are examples that explore employer differences as a source of earnings inequality in search models.

Our findings also connect to the personnel economics literature studying internal labor markets

[3]Examples include Deaton and Paxson (1994); Storesletten et al. (2004); Heathcote et al. (2005); Huggett et al. (2006). A common practice today 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), Guvenen and Smith (2014) and Bagger et al. (2014) have taken more structural approaches to exploring the drivers of life-cycle inequality.

and career dynamics following the seminal work of Doeringer and Piore (1985). The existing research in this strand of the literature relies 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. Gibbs et al. (2003) and 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.[4] 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."

On the theoretical side, we depart from modeling the underlying frictions to career progression and focus on the life-cycle implications of career progression for wage dynamics. 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.[5] 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. We discuss in detail the relationship of theories of career progression to our novel empirical results.

3 A theory of job levels

To fix ideas for our empirical analysis of job levels, we sketch a theoretical framework where job levels exist that combines a task approach with a worker assignment problem. For tractability, we assume that the firm plans as if it faces a continuum of workers. The firm then seeks to maximize output with its labor force of measure one. Output is produced according to the CES production function in tasks $\tau \in [0,1]$. Total output Y of that firm is given by the aggregation of output by tasks:

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

To produce output 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

[4] For the United States, Guvenen et al. (2014) document persistent gender earnings gaps at the top.

[5] 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.

of hours of all workers spent on task τ

$$y(\tau) = \int z_i(\tau) l_i(\tau) di. \tag{2}$$

The key assumption that leads to the emergence of job levels it that the firm cannot perfectly (without costs) observe the productivity of worker i. It 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. We assume that firm's can design assessment tests of the worker's skills, e.g. through interviews or assessment centers, but can do so only imperfectly. They can create tests that partition the skill space,

$$\{[a_1, a_2), [a_2, a_3), \ldots, [a_M, a_{M+1}]\}$$

with $a_1 = 0$ and $a_{M+1} = 1$ and allow to figure out the skill partition a worker falls into. We use as index for the skill group, the index of the lower bound of the skill interval, i.e., skill group j corresponds to skill set $[a_j, a_{j+1})$. The higher the number M of subsets, the higher the cost $c(M)$ of the tests, with $c(1) = 0$ and $\lim_{M \to \infty} c(M) = \infty$.[6] Since the gains from specialization are bounded by output in the frictionless allocation, it is clear that firms will choose a finite number of skill groups $j = 1, \ldots, M$. A grouping of workers into skill groups further implies that firms can condition anything worker relate not on her identity i, but only at the group level j. As the firm knows the underlying distribution of worker skills, it only knows the average skill of group j's representative worker $\bar{\alpha}_j$. Worker i's work schedule $l_i(\tau)$ is therefore determined by her skill group j. The schedule of a worker 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 workers productivity is in the interval $\alpha \in [a_j, a_{j+1})$, and an assignment of hours to tasks, a work schedule $l_j(\tau) \geq 0$, such that $\int l_i(\tau) d\tau = 1$.*

3.1 Optimal hours choice

Before we turn to the decision on the organization of the firm, i.e., how to decide about the task range of jobs and the minimum skill requirements of jobs, we discuss the more narrow question of the optimal hours choice in terms of the number of hours spent on each task. We assume that in the population worker ability a is distributed according to distribution function $F(a)$ with density $f(a)$. We first consider the assignment of hours to tasks for given number of jobs

[6] This implicitly assumes that the output of a single worker is unobservable so that firms can also not identify the skills of a worker indirectly, e.g., through experimentation. What is more, the assumption of costs diverging to infinity is not a necessary one. It would suffice that costs converge to more than the output in frictionless allocation of workers to tasks.

M and skill requirements a_j. This planning problem is given by

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

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

As hours must be assigned at the job, we can factor out hours and using in addition the assumption of the uniform distribution, we can rewrite the maximization problem as

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

This leads to the first-order conditions

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

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

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

where 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 j spends on task τ and ΔF_m denotes the group mass of workers assigned to work in job m, $F(a_{m+1}) - F(a_m)$. 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 will be produced)**.** *There is no task that the optimal job design problem assigns zero output to.*

Proof. As the production function fulfills the Inada condition that marginal products of tasks go to infinity when a task is not produced, the non production of a task cannot be optimal. □

This implies together with the first order condition for hours our

Proposition 1 (Jobs partition the task space)**.** *Jobs partition the number of hours in a way that no two jobs have positive hours simultaneously in more than one task. Moreover, the optimal task sets $T_j^* := \{\tau | l_j(\tau) > 0\}$ are an increasing partition of the task space.*

Proof. Assume the contrary, that is 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

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

needs to hold for both τ_j and τ_k. A linear equation in τ two have more than one solution, requires that the linear functions on both sides are identical, i.e., both level and slope of left and right hand side are the same. This would require $\frac{\lambda_j/\Delta F_j}{\lambda_k/\Delta F_k} = 1$ and $(\bar{a}_j - 1) = \frac{\lambda_j/\Delta F_j}{\lambda_k/\Delta F_k}(\bar{a}_k - 1)$. Which would imply $\bar{a}_j = \bar{a}_k$, contradicting the assumption of $j \neq k$.

As all tasks are produced, it must be that (7) holds exactly for one $\bar{\tau}_j$ for neighboring jobs j and $k = j + 1$, such that the second part of the proposition follows. The fact that the task sets are increasing in expected productivity follows directly from more skilled workers having an absolute advantage in all tasks and a comparative advantage in more difficult tasks. □

Importantly for our empirical analysis, this means that a job can be equivalently described as a skill requirement, $[a_j, a_{j+1}]$, as a range of tasks $[\bar{\tau}_j, \bar{\tau}_{j+1}]$, 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.

Proposition 1 shows that the optimal hours schedules partition the task space. In a first step, we now consider the optimal hours choice given for any partitioning. The optimal hours choice is given by

Lemma 2 (Optimal Hours Allocation)**.** *Assume an exogenously given, arbitrary partitioning of the task space in intervals $\{T_j\}_{j=1...M}$ with $\cup_j T_j = [0, 1]$ such that only job j can supply 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*

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

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

$$\begin{aligned} Y = & \quad \left\{\int_0^1 \left[\sum_{k=1}^{M} \left(1 - \tau(1 - \bar{a}_j)\right) \Delta F_j l_j(\tau)\right]^{\frac{\eta-1}{\eta}} d\tau\right\}^{\frac{\eta}{\eta-1}} \\ = & \quad \left\{\sum_{j=1}^{M} \int_{T_j} \left[\sum_{k=1}^{M} \left(1 - \tau(1 - \bar{a}_k)\right) \Delta F_k l_k(\tau)\right]^{\frac{\eta-1}{\eta}} d\tau\right\}^{\frac{\eta}{\eta-1}} \end{aligned} \tag{9}$$

which, because of $l_j(\tau) = 0$ if $\tau \notin T_j$ simplifies to

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

Maximizing Y subject to (4) by choosing $l_j(\tau)$ is characterized by the first-order condition

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

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

The right-hand side of the first-order condition is independent of a specific task what implies that any task served by job j needs to produce the same amount in terms of final good taking into account the average productivity of a worker in job j. As only one job serves one task, we can also use it 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:

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

and thus with the total labor supply constraint $\int_{T_j} l_j(\tau) = 1$ yields the desired result. □

3.2 Job design problem

It is straightforward to see that under the optimal hours choice total output can be written in a compact form as

$$Y(\{\bar{\tau}_1, \ldots, \bar{\tau}_{M+1}\}, \{a_1, \ldots, a_{M+1}\}) = \left(\sum_{j=1}^{M} (S_j \Delta F_j)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}. \tag{13}$$

The representative worker on job j produces a bundle of tasks that can be aggregated to S_j and then the firm aggregates over these bundles weighted by the employment shares ΔF_j of workers in each job j.

By choosing skill requirements and task ranges, the firm then maximizes (13) with the boundary conditions that $\tau_1 = a_1 = 0$ and $\tau_{M+1} = a_{M+1} = 1$. The first-order conditions for this maximization problem are, for $\bar{\tau}_j$

$$0 \quad = \frac{1}{\eta - 1} \left[\sum_{m=1}^{M} (S_m \Delta F_m)^{\frac{\eta-1}{\eta}} \right]^{\frac{1}{\eta-1}} \times \tag{14}$$

$$\left\{ (S_{j-1} \Delta F_{j-1})^{\frac{\eta-1}{\eta}} \left[\frac{1 - \bar{\tau}_j(1 - \bar{a}_{j-1})}{S_{j-1}} \right]^{\eta-1} - (S_j \Delta F_j)^{\frac{\eta-1}{\eta}} \left[\frac{1 - \bar{\tau}_j(1 - \bar{a}_j)}{S_j} \right]^{\eta-1} \right\}$$

and, for a_j,

$$0 = \left[\sum_{m=1}^{M}(S_m\Delta F_m)^{\frac{\eta-1}{\eta}}\right]^{\frac{1}{\eta-1}} \times \left\{(S_{j-1}\Delta F_{j-1})^{-\frac{1}{\eta}}\left[\frac{\partial S_{j-1}}{\partial a_j}\Delta F_{j-1} + S_{j-1}f(a_j)\right] + (S_j\Delta F_j)^{-\frac{1}{\eta}}\left[\frac{\partial S_j}{\partial a_j}\Delta F_j - S_j f(a_j)\right]\right\} \quad (15)$$

These first-order conditions can be 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 in (14), we use the optimal labor input schedule $l_j(\tau)$ from Lemma 2. Third in (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

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

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

The job design problem results in an allocation of worker groups to task ranges. If the firm decides to have M jobs, then the optimal job design is a combination of $M+1$ task thresholds $\{\bar{\tau}_1, \ldots, \bar{\tau}_{M+1}\}$ and $M+1$ skill thresholds $\{a_1, \ldots, a_{M+1}\}$. The allocation will be such that worker who are identified to have higher skills will work on harder tasks (Proposition 1). To think about the empirical implications of our model, we still need to assume a rule how tasks are allocated to workers within a job level. Note that assigning the ΔF_j mass of workers in job j to task $\tau \in \mathcal{T}_j$ with $l_j(\tau)$ hours produces the same output as randomly assigning each worker a only a fraction $\mathcal{A}_j = \frac{\tau_{j+1}-\tau_j}{\Delta F_j}$ of all tasks in $\mathcal{T}_j$ but then lets each worker work $\Delta F_j l_j(\tau)$ hours on the assigned tasks. To solve this indeterminacy, we assume that the firm minimizing the number of tasks per worker within the job level and thus assigns $\mathcal{A}_j$ tasks per worker of job level j.

Moreover, we have the following

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 $S_{k,j} := \left\{\int_{T_k}[1-t(1-\bar{a}_j)]^{\eta-1}\,dt\right\}^{\frac{1}{\eta-1}}$ the output in task group k that a worker of skill group j would be expected to produce if we allowed her to optimize her hours 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 $j > k$, since the worker in the higher skill group has an absolute advantage on all tasks, $\mathcal{S}_{k,j} > S_k$, so $w_j = p_j S_j > p_k S_k = w_k$ follows. □

To demonstrate the properties of the job design problem, we use a numerical example where we vary the parameter η governing the substitutability of tasks in the production function. We

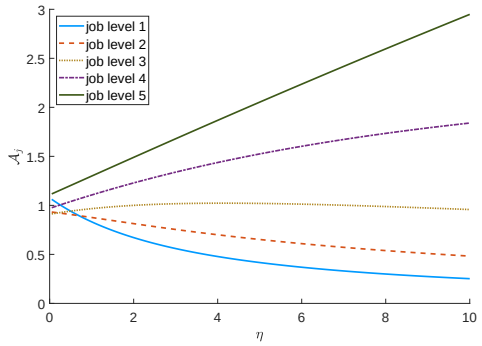
consider η in an interval from 0.05 to 10 going from very low to very high substitutability of tasks. For skills, we assume two alternatives. First, we consider a uniform distribution on $[0, 1]$ so that $F(\alpha) = \alpha$ and second we consider a Beta distribution with parameters $a, b = 5$, which is relatively similar to a normal distribution in shape but bound to $[0, 1]$.

Using the first-order conditions in (16) and (17), we solve for the optimal skill cutoffs $\{a_j\}_{j=1}^{M+1}$ and the optimal task cutoffs $\{\tau_j\}_{j=1}^{M+1}$. Figure 1 reports the results in terms of the optimal group sizes ΔF_j, the tasks per worker $\mathcal{A}_j$, and the increase in productivity relative to skills across job levels.

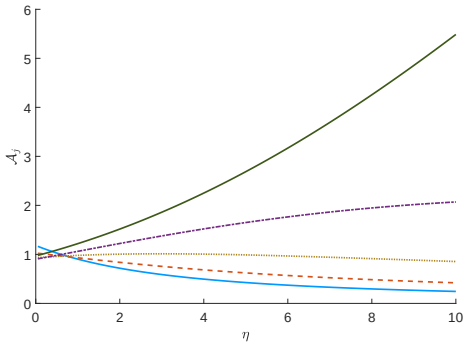
We find, first, that the tasks per worker increase in job levels as long as tasks are substitutes, $\eta > 1$. In other words, the most skilled workers are assigned not only the most difficult tasks, but have jobs with more disperse tasks. Second, the expected product of a worker is not only increasing in skill group but convex in expected skills: the expected product increases more quickly in skills the higher the job level. Exploiting the flexibility of the beta distribution, we find these first two findings to be robust to changes in the shape of the skill distribution. Third, we find that the size of the top job level group tends to be the smallest and its size is decreasing 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 **?**).

In summary, we provide a theoretical framework that 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. They design jobs with different job levels. 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 $\mathcal{A}_j$).

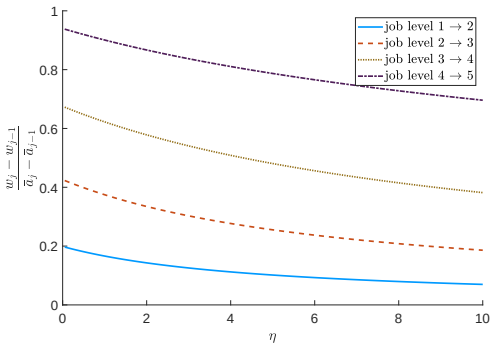
Figure 1: Results from the Job Design Problem for two alternative Skill Distributions $F(\alpha)$



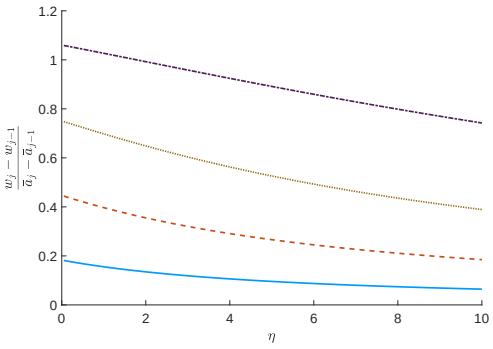
(a) Tasks per worker, $\alpha = \beta = 1$



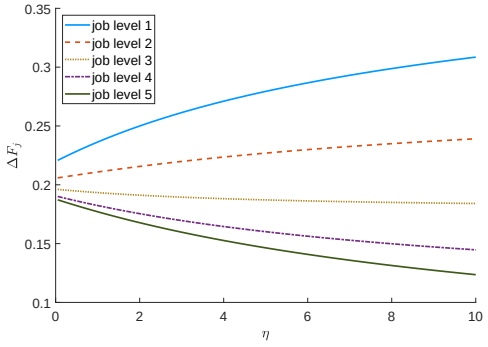
(b) Tasks per worker, $\alpha = \beta = 5$



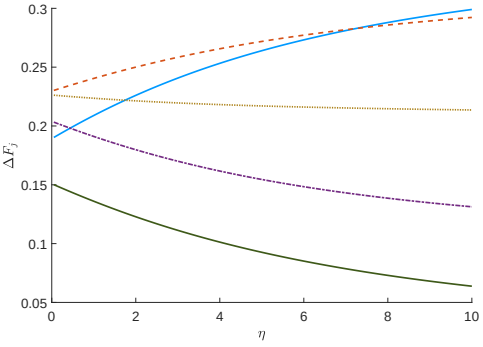
(c) Productivity increase, $\alpha = \beta = 1$



(d) Productivity increase, $\alpha = \beta = 5$



(e) Group size, $\alpha = \beta = 1$



(f) Group size, $\alpha = \beta = 5$

Notes: Optimal optimal task range $[\bar{\tau}_j, \bar{\tau}_{j+1}]$, change in productivity with skills, and optimal group size ΔF from job design problem with different distributional assumptions for $F(\alpha)$. The left column represents results for a beta distribution with $\alpha = \beta = 1$ (uniform distribution), while the right column uses $\alpha = \beta = 5$ (approximating a normal distribution). Each row represents different variables: optimal group size, optimal task range, and the change in productivity with skills, respectively. The vertical axis reflects the degree of substitutability of tasks in production.

3.3 Job level and career dynamics

For analytic tractability, we assume that firms use the solution to the job design problem from Section 3 also in a dynamic setting and when faced with a finite number of workers. One interpretation of this assumption is that job design is an irreversible one-time technology choice, and we use the optimal choice from the deterministic setting for the number of job levels, the filtration of the skill space, and the description of job levels in terms of task ranges as an approximation.

Given this assumption, we show that when faced with a small sample skill distribution that is different than expected, firms will typically accept assigning a worker with no matching skills to a job level rather than changing the fraction of workers at a particular job level ΔF_j.

We then extend this insight to a dynamic setting and argue that 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 ($\alpha_i < a_j$). However, they will accept to assign an overqualified worker i to a job level j ($\alpha_i > a_{j+1}$).

We combine these insights to build a model of job level and career dynamics that allows us to derive predictions for wage growth and inequality over the life cycle.

3.3.1 Worker allocation with a fixed job design

We return to the model developed in Section 3. We assume that firms have made an irreversible job design decision the filtration of the skill space $\{a_j\}_{j=1}^{M+1}$, and the description of job levels in terms of task ranges $\{\mathcal{T}_j\}_{j=1}^{M+1}$, and thus have fixed also the number of job levels M. Next, we consider such a firm with a realized fraction of workers different to the one expected, for example because it draws a finite sample of workers from F. We consider the case with a realized fraction $\Delta F_j + \delta$ in skill group j and $\Delta F_k - \delta$ in group k.

In this case, the firm needs to decide whether it employs all the additional workers δ in their skill-matched job at level j or shift them to the understaffed job level k. Allowing to reoptimze hours for "mismatched" workers, we obtain for the output of worker of type j in job k (see the proof of Proposition 2)

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

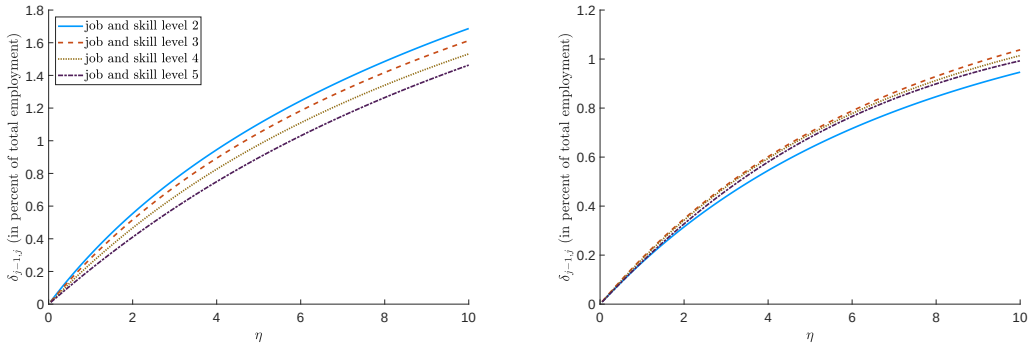
and that the firm will accept a deviation from the planned worker allocation, if and only if

$$(S_{j,j}(\Delta F_j + \delta))^{-\frac{1}{\eta}} S_{j,j} \geq (S_{k,k}(\Delta F_k - \delta))^{-\frac{1}{\eta}} S_{k,j} \tag{19}$$

which we can express equivalently as

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

Figure 2: Maximum acceptable deviation of workers per job level in percent of the workforce



(a) $\alpha = \beta = 1$ (b) $\alpha = \beta = 5$

Notes: Maximum acceptable deviation δ of workers per job level in percent of the workforce. Vertical axis 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 distribution), while the right panel shows results for $\alpha = \beta = 5$ (approximating a normal distribution).

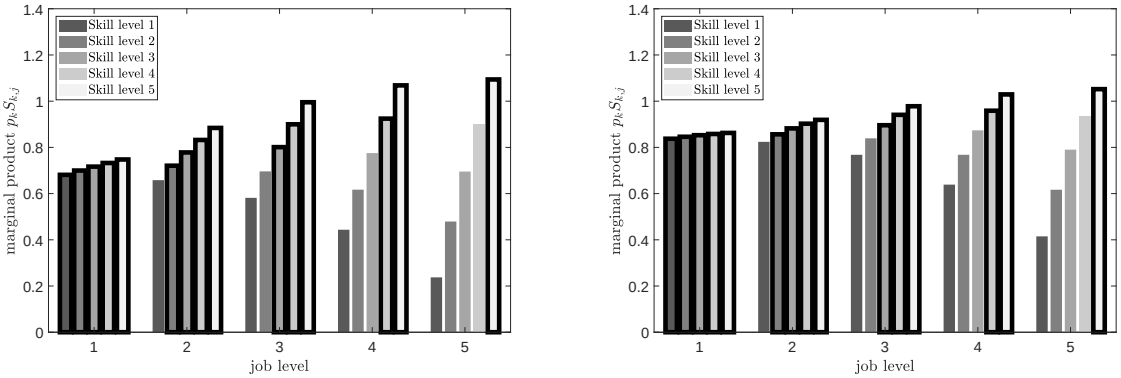
where p_k^* is the quasi-price in the optimal allocation from Section 3. If the inequality holds in the other direction, the firm has an incentive to reallocate a marginal worker to the understaffed job. The same right-hand criterion applies to reallocating a category that is staffed as expected relative to an understaffed job level.[7]

In Figure 2, we plot the maximum acceptable deviation $\bar{\delta}_{k,j} = \Delta F_k \frac{1-\mathcal{C}_{k,j}}{1+\mathcal{C}_{k,j}\frac{\Delta F_k}{\Delta F_j}}$, 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. Therefore, we assume in the following that a firm always keeps the number of workers per job level at the optimal level from the deterministic model. Hence, we get here the key distinction between a worker and her job.

In this situation, the marginal product of a worker $p_k S_{k,j}$ depends not only on the worker's skills but also, and importantly, on the job. In this sense, the model features the concept of *human capital utilization*. Figure 3 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 ($\eta = 10.0$). For the case $\eta = 0.5$ ($\eta = 10$) we get that at a low job level the effect of skills on marginal products is very small, moving up skill levels from the lowest to the highest group fixes employment at the lowest job level and increases the marginal product by only 10% (3%). In contrast, moving a worker with a skill level of five from the lowest job level to the highest job level is associated with a 46% (22%) increase in her marginal product. Despite this fact, we

[7] This 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 3: Marginal products of workers in different skill groups across jobs



(a) $\eta = 0.5$ (b) $\eta = 10.0$

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 3.3.2 for details. *Left panel:* Task are complements, $\eta = 0.5$. *Right panel:* Task are substitutes, $\eta = 10.0$.

have seen that the firm would rather underutilize a worker's skills than assign her to the "ideal" job level, because doing so would increase the output of a 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 $S_{k,j}$, see (19). 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 are no longer identical. In other words, there can be varying human capital utilization of workers of the same skill type.

3.3.2 Hiring decisions with a fixed job design

The preceding subsection considered unexpected deviations from the distribution of worker skills. A relevant and typical case is when the firm draws a finite sample of workers. We explain why firms, despite variations in the skill supply, do not deviate from their planned organizational structure: The share of workers on job level j will be approximately constant at the expected ΔF_j irrespective of the actual distribution of skills.

In a dynamic context, workers leave firms and firms hire new workers. To study such a situation, we will deviate from the continuous worker case ΔF and consider a case of discrete workers and positions at firms, where one worker can be in exactly one position. 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 $[a_j, a_{j+1}]$. The worker can be under- or overqualified for the vacancy. If the assessment test shows that the worker is overqualified, the firms is trivially willing to hire the worker as long as it has to pay at most the marginal product (Figure 3). 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 the case when a firm faces an underqualified worker at the hiring stage. We will consider this case in a stylized dynamic setting next.

To preserve tractability and focus the economic trade-off of the problem, we make some simplifying assumptions.

Assumption 1. *The firm has chosen the job design (task ranges and skill groups) permanently ex ante.*

Assumption 2. *We consider a firm that has N positions, out of which $N-1$ are filled. The workers on these $N-1$ positions will permanently remain with the firm. Given the skill set of the existing workforce, we assume that the open job is at level k.*

Assumption 3. *Workers on a job of level j are remunerated based on their job level at wage rate w_j.*

Assumption 4. *The firm meets a worker with matching skills to the job opening with probability λ, the N-th worker leaves the firm with an exogenous probability of ν plus an endogenous probability $\chi_{k,j}$ that depends on the job k and skills j and is given by the probability that this worker finds a better paying job elsewhere.*

Proposition 3. *Under these assumptions, 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 N-th position unfilled. 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 N-th 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

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

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

where the last inequality reflects the fact that looking for any overqualified worker has to be a better strategy than looking for an exactly qualified worker. The first equality follows from the fact that no weakly underqualified worker will ever leave the firm endogenously.

The firm will never hire an underqualified worker if for all $k > j$

$$\underline{V}_k > \tilde{V}_{k,j}. \tag{23}$$

Taking the limit of $\beta \to 1$, the condition can be rewritten such as to obtain

$$\underbrace{\nu^{-1}}_{\substack{\text{expected match} \\ \text{duration}}} \quad \underbrace{(\tilde{Y}_{k,k} - \tilde{Y}_{k,j})}_{\substack{\text{output loss from} \\ \text{underskilled workforce}}} \quad > \quad \underbrace{\lambda^{-1}}_{\text{search duration}} \quad \underbrace{[(W - w_i) - \underline{Y}_k]}_{\text{wage bill net of output}} \tag{24}$$

as a sufficient condition. The left hand side is the output difference between a firm with an ideally skilled workforce and an underskilled workforce weighted by the expected duration of a match. Since $\tilde{Y}_{k,j}$ is increasing 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. □

Note that Assumption 3 could be weakened in that it is sufficient that the pass through of output losses from hiring an underqualified worker to wages is imperfect. Then the firm faces qualitatively the same trade-off: foregone persistent profit differences between an ideal and an underqualified workforce and the short-run losses of employing only $N - 1$ workers.

Proposition 3 implies that in a dynamic setting we will not observe underqualified workers on any job 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 1 when moving from job level 1 to 5 in Figure 3. By contrast, we will observe the career dynamics of skill type 5 moving from job level 1 to job level 5. We indicate the worker type job combinations that will be observed by bars with bold black edges in Figure 3.

4 The Structure of Earnings Survey data

Our main data sources are the 2006, 2010, 2014, and 2018 waves of the Structure of Earnings Survey (*Verdienststrukturerhebung*, SES). The data include over seven million employee observations from over 100,000 establishments with at least 10 employees across all survey years. The survey is conducted by the German Statistical Office and establishments are legally obliged to participate. Establishments with 10-49 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 hours paid, both regular and overtime, are extracted from payroll accounting and personnel master data and transmitted via a software interface to the statistical office. Unlike German social security data, the SES reports the actual (virtually uncensored) pay and hours worked of employees. The survey also provides detailed information on workers' education, occupation, age, tenure, and job levels. Self-employed workers are not covered. The survey has information on about 3.2 million employees in 2006, 1.9 million employees in 2010, and 0.9 million employees in 2014 and 2018. The number of sampled employees decreased over time because the sampling probability of plants became smaller to reduce average bureaucratic costs. In our analysis, we equalize observation weights across surveys so that

all surveys receive equal weight.

For our baseline analysis, we restrict the data to workers ages 25 to 55. We drop very few observations where earnings are censored[8] and all observations for which the state has a major influence on the plant.[9] We drop observations from the public administration and 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.

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

	Wages (in 2010 €)					Pop. Share of Job Level (in %)					
Males	Av.	Gini	p10	p50	p90	1	2	3	4	5	N. Obs
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.

As a 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. Importantly, this group includes workers who have not completed a secondary education.[10] For occupation coding, we use two-digit 2008 ISCO codes. We rely on a

[8]The censoring limit is €1,000,000 in 2006 and €750,000 since 2010 in annual gross earnings. We impose the latter throughout.

[9]We run a robustness check in which we include publicly owned/dominated plants, too; see Appendix G. For a large set of observations, the 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, the state owning $\geq 50\%$ share, or influence arising from other regulations.

[10]Additional information in the 2014 SES data allows us to infer that the typical case in the "other" group are

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.[11] Table 1 reports descriptive statistics for males and females 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 level 1 to 5, from 1 being the lowest job level to 5 being the highest job level.

Table 2: Importance of characteristics in explaining hourly wages

	Plants	Job levels	Job levels and plants	Job levels, plants, occupations, education, experience, and sex	Job levels, plant size, region, and industry
(adj.) R^2	0.583	0.471	0.782	0.813	0.618

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

The SES data are particularly well suited to decompose wage differences across workers because they offer a very high explanatory power of observable characteristics for wages. Taken together, all of the information on workers, employers, and jobs accounts for over 81% of the observed cross-sectional variation in wages (Table 2). The high quality of the data is key for delivering this very high degree of statistical determination. Besides data quality, the other and economically more important reason for the high explanatory power is that we observe job levels. Dummies for five job levels alone account for more than 47% of cross-sectional wage variation; adding plant dummies observables accounts for 78%; and combining job levels with plant characteristics accounts for 62% of wage variation.[12] We corroborate our findings on the high explanatory power of job levels for wages in US NCS data in Appendix A.

4.1 Job levels

A key distinction between the SES data and most other data sources is that they provide information on workers' job levels. Data from other countries that include job levels and wage data are discussed, for example, in Strub et al. (2008) and Pierce (1999). 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. Similarly long reporting of wages by job levels exists in the reporting of the BLS for the United States.

workers without a completed secondary education.

[11] Crosswalk retrieved from International Labour Organization, ISCO—International Classification of Occupations "ISCO-08 Structure, Index Correspondence with ISCO-88," http://www.ilo.org/public/english/bureau/stat/isco/isco08/index.htm.

[12] Appendix Figure A15 shows the large wage differences by job level over the entire life cycle in the SES data.

The assignment of job levels differs in detail across job-leveling schemes but can be summarized as encoding the complexity, autonomy, and responsibility (CAR) in the execution of the job's tasks, given a specific task assigned to a job. The complexity of task execution relates to the minimum skill requirement that a worker will need to execute the job's tasks. As a minimum requirement, it does not rule out that higher-skilled workers do a job with lower skill requirements. Acemoglu and Autor (2011) discuss employers' allocation problem of workers to tasks of different complexity. Autonomy captures how closely a jobholder has to follow a predefined workflow and how much decision-making power is granted in the execution of tasks. Responsibility refers to the scale of operations affected by the jobholder's decisions, i.e., if task execution only affects one's own work or also the work of others. Conceptually, jobs are then described by two dimensions: on the one hand, by the occupation as describing *which* tasks are done, on the other hand, by the job level describing *how* tasks are done.

The coding instructions for the different job levels in the SES data can be summarized as follows. At the lowest job level, the 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 for some discretion in the workflow. Junior clerks or salespeople would be typical examples. Yet, the task execution in these jobs does not include responsibility for the work of others or decisions that affect the work of others like 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. For example, all jobs in research will be classified into the two highest job levels because of their complexity, autonomy, and responsibility. The 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. (2023).

Importantly, job levels in the SES or comparable data are constructed such that they offer a consistent distinction of how tasks are executed within and across occupations.[13] Focusing on

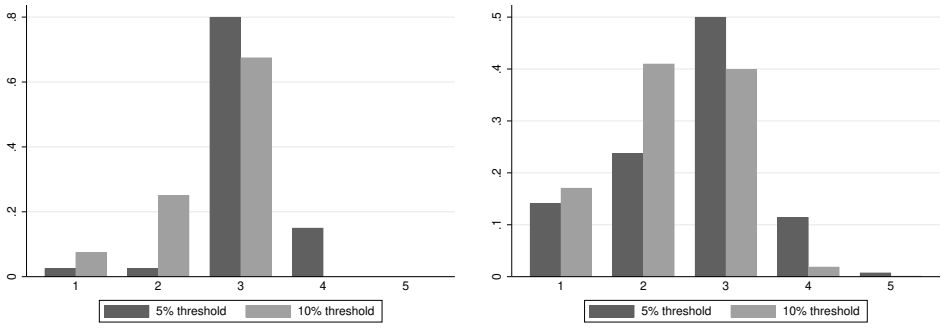
[13]The BLS job-leveling guide describes in detail the job-leveling approach for the U.S. NCS data

task execution, we follow the key idea of the task-based approach that wages are determined by executed tasks rather than by the stock of human capital of the jobholder (Autor et al., 2003; Acemoglu and Autor, 2012). We extend the task-based perspective by the refined distinction on *how* tasks are executed and find this refined perspective to be of primary importance in accounting for wage differences in the data.

4.2 Job levels and occupations

Which task is executed by a worker and how this task is executed, is 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., 5-digit occupation codes. Next, we, therefore, provide a detailed discussion on how occupations and job levels relate.

Figure 4: Share of occupations with different job-level span



(a) two-digit ISCO codes

(b) five-digit KldB codes

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.

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 4 for (a) two-digit and (b) five-digit occupation codes.[14] 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. Thus, there is a clear distinction between job levels and occupations.

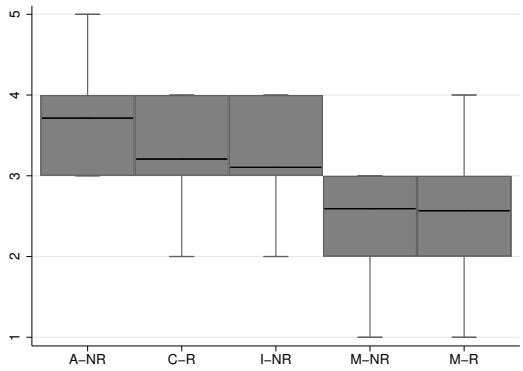
At the same time, these findings imply that an occupation does not span all job levels (equally)

(US Bureau of Labor Statistics, 2013).

[14] Results for 5-digit codes are based on 2018 SES data alone as 5-digit codes are not included in the scientific use file of the SES data.

because not all tasks can be done 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) and classify occupations into five task groups based on whether their main tasks are analytical, interactive, cognitive, non-routine manual, or routine manual.[15] Figure 5 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. 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 4, 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 spans typically three job levels. In Appendix B, we discuss that the CAR intensity of occupational tasks is the main explanatory variable of task-based occupational wage differences. Our results highlight that job levels provide a novel and important approach to refining the key idea of the task-based approach.

Figure 5: 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. Bars show the interquartile range, the black horizontal line shows the mean job level, and the vertical lines indicate the range between the 10^{th} and 90^{th} 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).

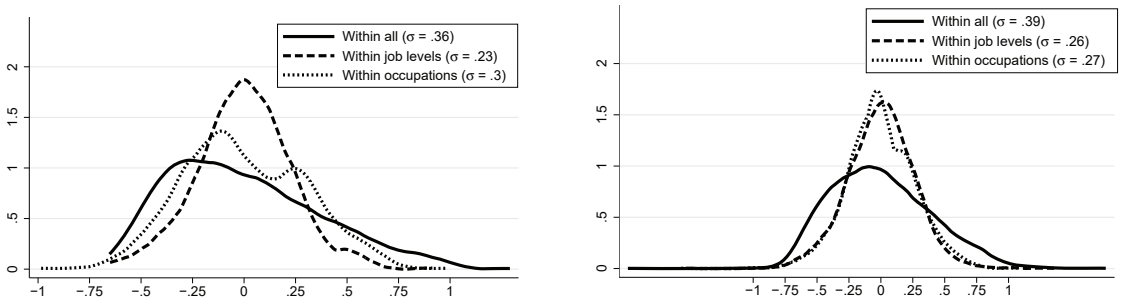
In a nutshell, this can be seen in Figure 6. The figure compares wage differences across and within occupations. For this purpose, we aggregate wage data by job-level-occupation cells. Then we either regress these aggregated data on the five job-level dummies or, the much finer-grained, occupation dummies. Figure 6 shows the distribution of the regression residuals for (a) three-digit ISCO codes (120 categories) and (b) five-digit KldB 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

[15]This classification is based on 3-digit occupation codes. 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 C 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.

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.

In Appendix A, we compare wage densities and standard deviations across occupation-job-level cells from U.S. NCS data to the corresponding German 2018 SES data. We find for the U.S. data as for the German data that job levels have higher explanatory power than occupations. We furthermore report that in the U.S. data job levels account for half of within occupation wage variation.

Figure 6: Wage density across occupations by job level for different occupation codes



(a) three-digit ISCO codes

(b) five-digit KldB 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.

4.3 Job contents and job-leveling factors

Survey respondents in the SES are given instructions on how to assign job levels to jobs based on their CAR intensity. Each element, complexity, autonomy, and responsibility, factors into this assignment; wages do not. However, one might be concerned about reverse causality in the form of survey respondents not following the survey instructions and providing job-level information based on wage information. Using additional and independent survey data, we demonstrate that job leveling with the same explanatory power for wages can be done in another dataset that has this information based on descriptions of how tasks are executed. Of course, we do this without considering wage data. The BIBB/BAuA data provide us with the necessary details on task execution beyond the task content of occupational data. It also provides wage data, worker demographics, industry, and employer size. From the task execution data, we select the information used for job leveling. We refer to this information as job-leveling factors. We demonstrate that these job-leveling factors in the BIBB/BAuA data have the same explanatory

power for the wage data as coded job levels in the SES data. In other words, we show that job levels have economic content and can be constructed from information on task execution.

We restrict the BIBB/BAuA sample to align with our SES analysis. We keep workers ages 25 to 55 who work at employers with at least 10 employees and drop workers in public service. We drop self-employed workers, freelance workers, independent contractors, and family workers. We further restrict the sample to workers who do not report second jobs and report regular working time between 35 and 45 hours per week to reduce measurement error in hours.[16] Some of the wage information in the survey has been imputed, and we drop all observations with imputed wage information. We first restrict the analysis to white-collar workers and report results for blue-collar workers in Appendix D.2.[17] The final sample has 3,027 observations with complete information for the analysis.

The survey collects data from workers on their monthly earnings and typical hours worked. We use these data to construct wages. Constructed wages in the BIBB/BAuA data likely contain substantially more measurement error than wages from the SES data, which are based on employer-reported earnings and hours. This reduces explanatory power in the regression analysis below. As job-leveling factors, we select eight survey questions that we identify to be informative about a job's CAR intensity. This selection is based on the job-leveling scheme from the German steel- and metalworker bargaining agreement.[18] We report the detailed survey questions in Appendix D.1. We encode answers to these questions using dummy variables and refer to them as job-leveling factors. First, we explore the explanatory power of the job-leveling factors by running a series of linear wage regressions. Table 3 reports the R^2 from these regressions. Second, we apply the job-leveling scheme from the bargaining agreement directly to assign job-level points that weigh the different job-leveling factors to aggregate job-leveling factors into job levels. We demonstrate that the points obtained from our first statistical analysis and the points of the bargaining agreement practically coincide.

When we run a regression on the job-leveling factors only, we account for 44% of wage variation.[19] 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%.[20] This aligns well with our findings in the previous subsection. 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

[16] Appendix G shows that our previous results based on SES data are very similar if we consider full-time workers only.

[17] Information on task complexity is coded separately for blue-collar and white-collar workers, which makes the data too intricate to aggregate and compare.

[18] The 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.

[19] The regression involves 18 dummies for answers to the eight questions and a constant.

[20] A regression on occupations alone accounts for only 34% of the wage variation.

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.

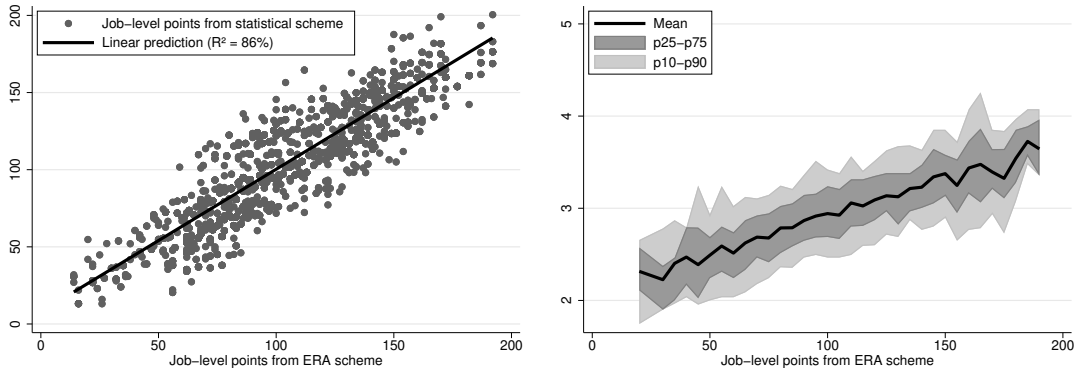
intensity constructed directly from information on task execution and wages as in the SES data, where CAR intensity is summarized by job levels.

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.[21] 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 D.1. In Figure 7(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 7(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 show a clear positive relationship between job-level points and (log) wages.[22]

[21] 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.

[22] It 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.

Figure 7: Wages by job-level points



(a) Statistical scheme

(b) Distribution (by groups of 5 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). 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).

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.[23] These two points address the important question of reverse causality that job levels could be just a recoding of wages (e.g., wage quintiles). This point is important with respect to theories of tournament models of career progression as pioneered in Lazear and Rosen (1981). With job levels differing in task execution, promotions across job levels will involve a change in the CAR intensity of jobs for promoted workers together with the wage change from the promotion. Although our results point to an important role of changes in task execution related to promotions there is obviously still room for incentive-related promotion wage dynamics.

In summary, we conclude that data on *how* tasks are executed provides important independent information over *which* tasks are executed (occupations) for determining wages. How jobs are executed is captured by what human resources and statistical offices call job levels.

[23] An 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 clear, why there is a loss in predictive power already compared to the dummy regression. In Appendix D.1 Figure A6, 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 independent of the wage structure has also been shown in case studies (Dohmen et al., 2004).

5 Job levels over the life cycle: wage growth and wage inequality

In the next step, we turn from the cross-section to the life cycle to examine the extent to which changing job levels, in addition to changing employer and worker characteristics, account for life-cycle wage growth and the increase in wage inequality. We first discuss our methodology for decomposing life-cycle wage dynamics before discussing the results.

5.1 Methodology

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

$$w_{ipt} = \gamma_i + \zeta_{pt} + \beta_J J_{ipt} + \beta_I I_{ipt} + \epsilon_{ipt}, \tag{25}$$

where J_{ipt} is the characteristics of the job of individual i at plant p at time t, I_{ipt} is 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 worker characteristics comprising education and experience that we include as education and gender-specific age dummies.[24] The *job component*, $\beta_J J_{ipt}$, captures the characteristics of a job. We use dummies for two-digit occupations and five job-level dummies.

One challenge for the decomposition of life-cycle wage dynamics is that unobserved individual characteristics could jointly affect wages and the career progression of workers. A simple OLS estimate of wages on job levels would then be inflated because more able workers obtain higher wages at any job and are also more likely to end up at higher job levels. Such unobserved worker heterogeneity as the driver of career dynamics is the focus of the seminal work by Waldman (1984) and Gibbons and Waldman (2006). 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 in the decomposition (see Deaton (1985) and Verbeek (2008) for an overview of the method).[25] The aggregation of the data to the cohort level has the further advantage that it mitigates the concern of biased estimates as the identification stems only from the variation in the job composition across cohorts rather than at the individual worker level. As a second approach, we estimate the effects of job levels on wages using a shift-share instrument (Bartik, 1993). We discuss the synthetic panel approach as our baseline approach and relegate the discussion of the results of the instrumental variable regression to Appendix E. We opt for OLS estimation with cohort fixed effects as our baseline approach because it is easier to interpret and has favorable small sample properties. We also provide additional discussion in Appendix E on potential identification challenges arising from Waldman (1984) and Gibbons and Waldman (2006) and conclude that they should be of no concern for our analysis.

[24] We group ages using three-year windows to identify cohort effects later on, given the four-year distance between the three survey waves.

[25] We provide results based on pooled worker-level OLS in Appendix G.4.

In the first step, we control for plant-year effects by demeaning all variables at the plant level (year by year):

$$\hat{w}_{it} := w_{ipt} - w_{.pt} = \hat{\gamma}_i + \beta_J \hat{J}_{it} + \beta_I \hat{I}_{it} + \hat{\epsilon}_{it}, \tag{26}$$

where $\hat{X}_{it}$ denotes the difference between variable X_{ipt} for worker i and its average $X_{.pt}$ at the plant where this worker is working. As a consequence, the plant component drops from the regression and we explain below how we construct the estimate of the plant component, $\tilde{\zeta}_{pt}$. We now define cohorts based on workers' sex, birth year, and regional information (north-south-east-west),[26] and we aggregate variables to the cohort level to obtain

$$\hat{\bar{w}}_{ct} = \hat{\bar{\gamma}}_c + \beta_J \hat{\bar{J}}_{ct} + \beta_I \hat{\bar{I}}_{ct} + \hat{\bar{\epsilon}}_{ct}, \tag{27}$$

where $\hat{\bar{X}}_{ct}$ denotes the average of $\hat{X}_{it}$ within cohort c. This means that we estimate the coefficients of interest, β, from aggregate cohort data instead of from individual data. This allows us to use fixed effects OLS to obtain unbiased estimates $\tilde{\beta}_J$ and $\tilde{\beta}_I$ from equation (27) even in the presence of unobserved heterogeneity at the individual level that might lead to cohorts differing in their unobserved average fixed effect $\hat{\bar{\gamma}}_c$. Hence, we rely on the key idea of Deaton's (1985) synthetic panel estimator and use between-group variation in outcomes and observables for identification of the coefficients of interest.

Using the coefficient estimates $\tilde{\beta}_J$ and $\tilde{\beta}_I$, we construct directly 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 worker and job observables. It is given by

$$\tilde{\zeta}_{pt} \quad = \quad w_{.pt} - \tilde{\beta}_J J_{.pt} - \tilde{\beta}_I I_{.pt}. \tag{28}$$

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 worker fixed effects so that the average worker effect at a plant is not separately identified from the plant effect, but the individual and job components are consistently estimated. If there is no assortative matching in unobserved plant and worker heterogeneity, then the plant component is consistently estimated. If matching is positively (negatively) assortative, the plant effect tends to be positively (negatively) biased.[27] If sorting takes place over the life cycle, then a trend towards positive assortative matching will show up as a rising average plant component over the life cycle.

[26] The annual gross migration rate between German states in the past 30 years is low and has been roughly 1.3% per year; see *Wanderungsstatistik* of the Statistisches Bundesamt. More than a third of this migration is between states of the same region.

[27] The estimate by Card et al. (2013) for Germany is based on the Abowd et al. (1999) approach that 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.

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.

To emphasize again the identifying variation, recall that we have first demeaned the data at the plant-year level and, hence, we have taken out region-year effects. Second, we have taken out cohort effects in the estimation. Therefore, we do not use differences across cohorts or common time trends of all cohorts in a region for identification but instead exploit different time variation across cohorts for identifying β_J and β_I. In other words, we exploit how wages and (job) characteristics evolve over time within a cohort while simultaneously controlling for variations that affect all cohorts in a region.

An example of the type of variation we use is the entry of a new plant into a region, for which this plant has an atypical organizational structure. If this has more of an effect on the job characteristics of worker cohorts that are young at the time of entry at that plant relative to 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 hence more likely to exploit new job opportunities. Another example would be (regional) business cycles with heterogeneous impacts on cohorts. More generally speaking, 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.

5.2 Results

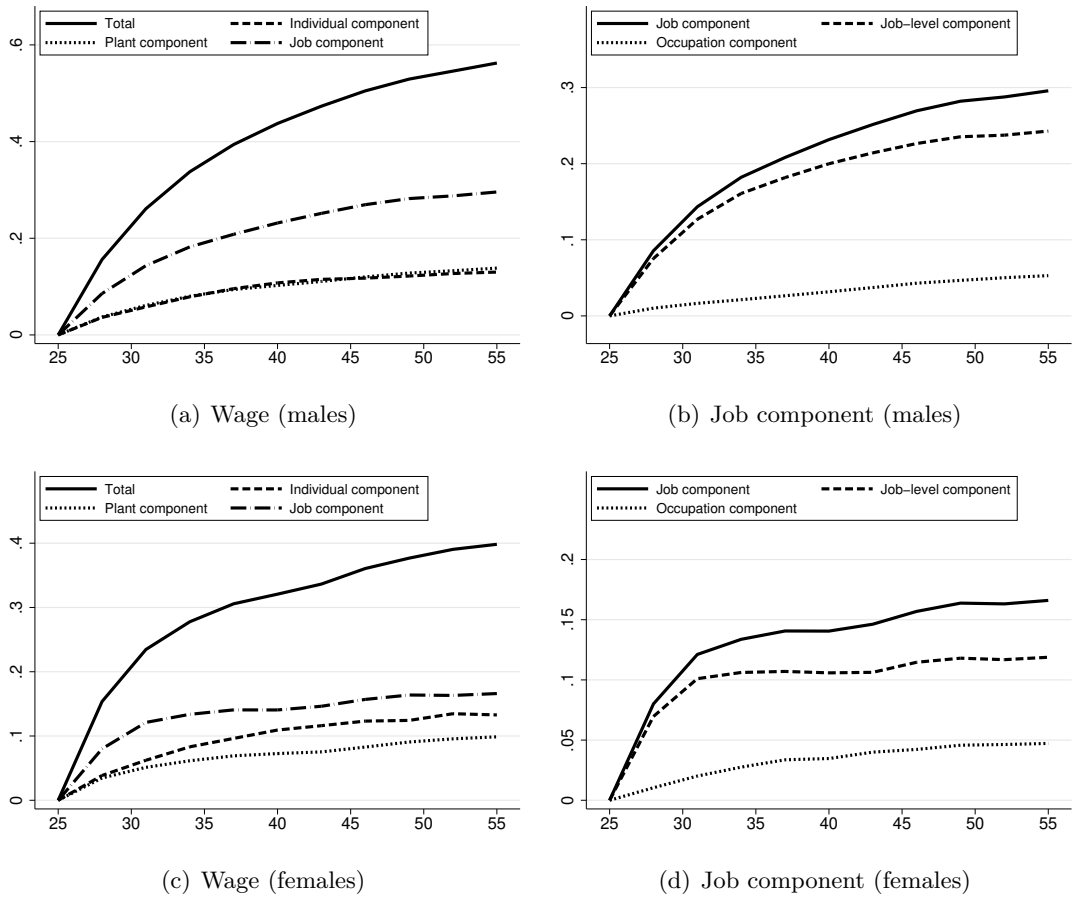
5.2.1 Wage growth

Based on our estimation results, we decompose 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. 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, we, therefore, 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 8(a) reports the decomposition of mean log wages for men. On average, wages grow by approximately 56 log points over the life cycle and we find that the job component accounts for more than 50% of this wage growth. Moving to better-paying plants over the life cycle, 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 fact that climbing the career ladder towards higher job levels is the most important component of wage growth can be seen when looking at the decomposition of the job component (Figure 8(b)). We find that an increase in the average job level accounts for most of the wage

Figure 8: Wage and job component decomposition



(a) Wage (males)

(b) Job component (males)

(c) Wage (females)

(d) Job component (females)

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. Right panel: Decomposition of the job component (solid line) into the contribution of occupations (dotted) and job levels (dashed).

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, meaning jobs with increasing CAR intensity over the course of their careers.

Figures 8(c) and (d) show the corresponding wage decomposition for females. Female wages grow by only 40 log points compared to 56 log points for males. Our decomposition in Figure 8(c) shows that a substantial part of this difference is accounted for by the smaller increase in the job component, in particular, a slower progression of women towards higher job levels. While the job component still accounts for the lion's share (17 log points), it increases substantially less than the one of males (30 log points). The reason is that between ages 30 and 45, there is hardly any growth in the job component for females. It starts to increase slightly again

only after age 45. As for men, we find for women that only a small part of the increase in the job component stems from the occupation component, which accounts for less than 5 log points of females' wage growth (Figure 8, d). The individual component for females accounts in relative terms for slightly more of the total growth than for men (33% versus 25%). The plant component for females shows a similar profile as for males but the increase slows down around the age of 30. One reason for the slowdown of wage growth in the plant component could be that the non-wage aspects of a plant, such as its location or working time arrangements, play a more important role for females than for males at this stage of the life cycle. In line with Morchio and Moser (2020) who find an important role of non-wage components for gender wage gaps in Brazil.

In summary, these results demonstrate that most of the life-cycle wage growth for males and females is accounted for by changes in how tasks are executed (job levels) rather than which tasks are executed (occupations). We find that most wage growth is accounted for by workers climbing the career ladder to high CAR-intensity jobs that are more complex and require jobholders to execute more autonomy and take on more responsibility. In short, CAR intensity drives wage growth.

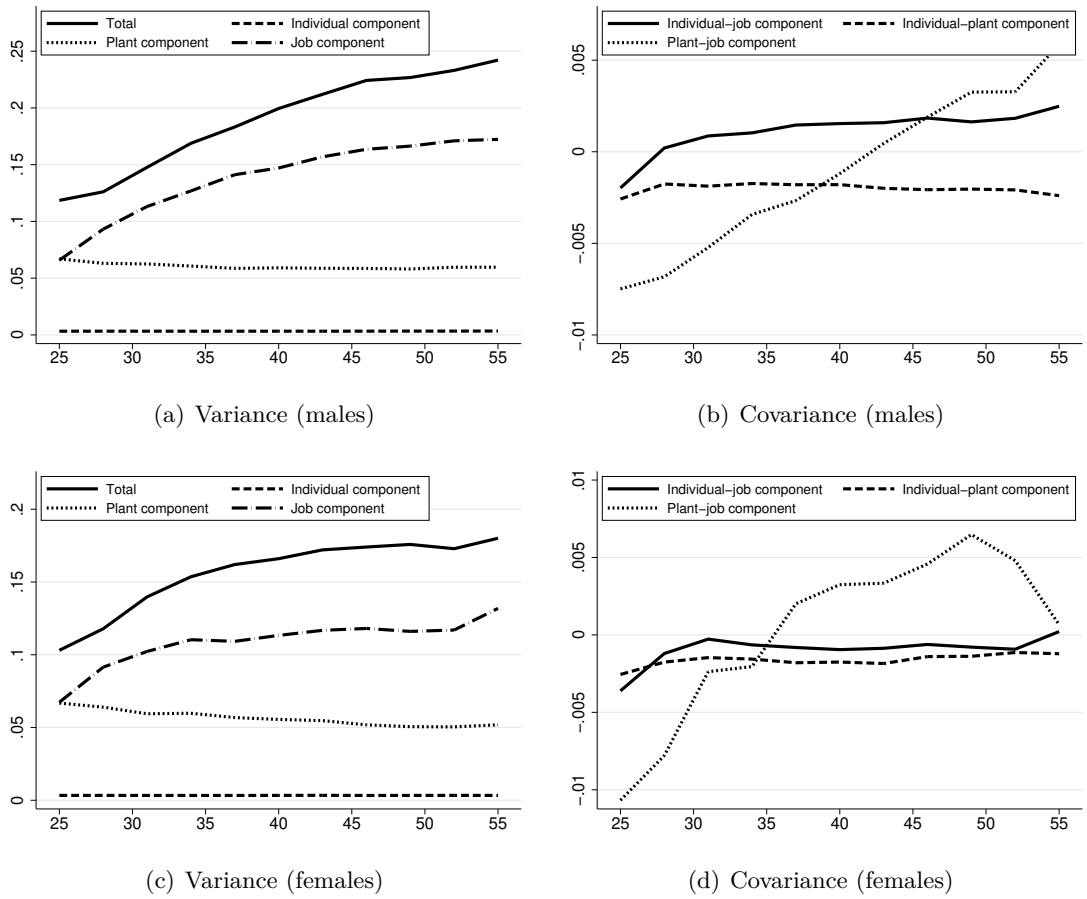
5.2.2 Wage inequality

Next, we show that not all workers follow the same career path so wage inequality rises over the life cycle. Hence, rising differences in CAR intensity also account for rising wage inequality. For this purpose, we decompose rising wage inequality over the life cycle. The high degree of statistical determination of 81% in our data (see Table 2) allows us to provide a much more fine-grained decomposition of the determinants of this rising wage inequality than is possible based on alternative data sources including administrative data.

Existing microdata based on cross-sectional regressions typically account for about 30% of wage inequality by observables and leave the largest part of wage inequality unexplained. Consequently, the literature interprets the largest part of the increase in wage inequality by age as the result of idiosyncratic risk captured by a stochastic process. This way of interpreting residual wage differences is the typical approach in a wide range of models including the large class of microfounded models of consumption-savings behavior (Huggett, 1993; Aiyagari, 1994). In our data, observable characteristics of workers and jobs statistically explain a much larger share of wages and wage growth. Hence, we can relate the idiosyncratic risk that remains as a residual in other data to concrete, observable, events in the labor market.

In Figure 9(a), we display the decomposition of life-cycle wage dispersion for male workers. We find that the variance of log wages increases from roughly 12 log points to 24 log points. The variance of the plant component contributes to the level of wage dispersion on average 6 log points but is virtually flat over the life cycle. The job component, by contrast, shows an 11 log point increase in its variance, from 6 to 17 log points. Put differently, almost the entire increase in wage variance is accounted for by workers becoming increasingly different in the type of jobs

Figure 9: Variance-covariance decomposition



(a) Variance (males)

(b) Covariance (males)

(c) Variance (females)

(d) Covariance (females)

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.

they hold. As for average wages, the job level is the main driving variable (not displayed). The variance of the individual component is virtually zero. Education itself has a negligible direct effect on wage differences across workers once we control for job levels (see Section 5.4 for further details).

Figure 9(b) complements these results by adding 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 on high job levels mostly at plants that do not pay well on average and as workers age, high-level jobs become more prevalent at well-paying plants. In other words, only when young there is 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 a systematic variation over the life cycle.

The additional covariance term between the plant and the job component increases from slightly less than -0.5 log points to slightly more than 0.5 log points over the life cycle. This means that the covariance terms contribute another 2 log points to the increase in the variance over the life cycle (twice the life-cycle increase of the covariance term). This additional covariance term accounts for the remaining part of the increase in wage dispersion not accounted for by the job component alone. Hence, the dispersion in the job component and the covariance of the job component with the plant component account for virtually all of the increase in wage dispersion over the life cycle.

Figures 9(c) and (d) show the decomposition results for the life-cycle wage dispersion of females. We have seen that women have a flatter average job-level component than men after age 30. This result also has implications for the evolution of life-cycle wage inequality among women. Their wage dispersion grows less by age (Figure 9, c). In particular, the increase accounted for by job-level dispersion is much smaller for women than for men and levels off after age 35. Still, the life-cycle profile of the job component accounts for 84% of the 8 log point increase in wage dispersion over the working life of females (compared to a 12 log point increase for the variance of males). For females, we find a slightly declining life-cycle profile in the plant component (Figure 9, c). At the same time, the job-plant covariance profile is even steeper for women than for men (Figure 9, d) meaning that those women who end up at high job levels at age 50 work in high-paying plants. Adding this covariance term to the job component as in the decomposition for males, we also find that virtually all of the life-cycle increase in wage inequality is accounted for by the job component.

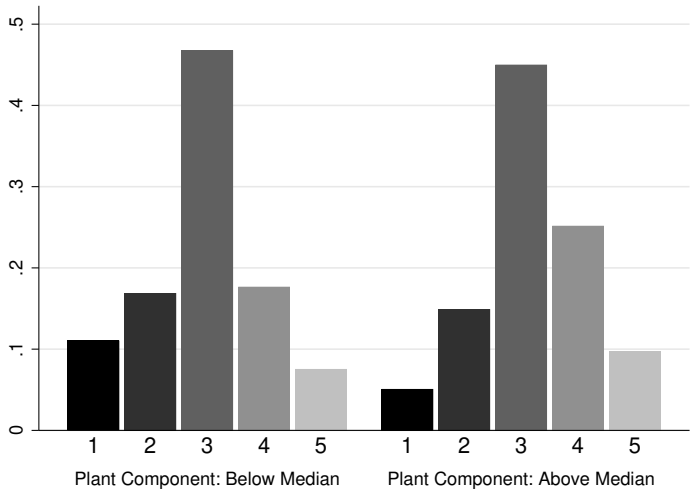
In summary, our decompositions of life-cycle wage growth and life-cycle wage inequality assign a key role to career ladder dynamics, i.e., workers progressing differentially across job levels as they age. We find a tight link between wages and changes in workers' job levels capturing the CAR intensity of task execution. We conceptualized these differences in job levels within and across occupations as variations in how tasks are executed and hence as distinct from occupations. Except for the average wage growth of females in the second half of their working life, we always find a dominant role for changing job levels in accounting for life-cycle wage dynamics.

5.3 Organizational structure and employer wage differences

Our decomposition shows that the plant component accounts for a significant part of the wage differences across workers. Importantly, we also 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. Recent evidence for Germany and the United States finds employer wage differences to be a key driver of increasing wage inequality over time (Card et al., 2013; Song et al., 2015). At first

glance, these two pieces of evidence seem to contradict each other, but they can be reconciled if we take the job-level perspective on wage dynamics that emerges from our novel empirical evidence. If plants differ in their organizational structure, plants with many high-level jobs will appear to be "high-wage plants", and any correlation between organizational structure and the plant component will reinforce this pattern.

Figure 10: 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, there is such a correlation between estimated plant component and organizational structure, as can be seen by looking at the distribution of workers across job levels for plants sorted by the estimated plant component $\tilde{\zeta}_p$ (Figure 10). 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. (2013) for Sweden.

These findings on the correlation between plant pay and organizational structure imply that missing information on organizational structure will bias conclusions about the importance of plant differences for the life-cycle dynamics of wages. We provide such a decomposition in Appendix F. When the job component is not observed, we find that moving to better-paying plants accounts for 38% of the life-cycle wage growth for male workers, but this figure is only 25% when the job component is observed. As the residual to the decomposition of wage growth, the individual component now absorbs most of the increase in wages, accounting for 62% of life-cycle wage growth. Similarly, we find that without considering the organizational structure of firms, 74% of female wage growth is now accounted for by the individual component. The increase in the individual component almost completely absorbs the contribution of the job

component from our full decomposition. For the increase in wage inequality, we find that when the job component is ignored, almost all of the increase in wage inequality is unaccounted for for both men and women. Only the individual component accounts for part of the increase if job information is missing.

5.4 Job levels as mediators of returns to education

The returns to education are a widely studied fact on wage differences between worker groups (e.g. Goldin and Katz, 2009; Katz and Murphy, 1992; Juhn et al., 1993). As our decomposition highlights the importance of differences in job levels, it raises the question of the relationship between differences in job levels and the returns to education. The natural idea is that the observed returns to education are mediated through a faster progression in terms of job levels. This idea arises directly from looking at the estimated coefficients on education from our baseline regression (27) where we get coefficients on education that are close to zero and insignificant. Table 4 exemplifies this for the coefficient on college education. It shows the estimated returns to education from three different specifications. Under our baseline, which controls for job characteristics, a college education yields a virtually zero wage premium over vocational training. Once we leave out job-level information, the returns to education go up to 31% and become highly statistically significant. If we drop all job component controls (job levels and occupations), the returns to education increase further to 54%. In words, the returns to education arise because education enables a worker to execute more CAR-intensive jobs and this is what increases wages.

Table 4: Transmission of returns to education through jobs

	data		
	baseline	w/o job levels	w/o all job info
College	-0.01	0.31***	0.54***

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 three different specifications: first our baseline, second a specification that leaves out job-level information, and third a specification that leaves out job information (levels and occupations) altogether. The baseline education category is vocational training. *, **, *** indicate significance at the 10%, 5%, and 1% levels, respectively.

As it was with occupations, however, education is neither a necessary nor sufficient condition to be on high job levels as human capital utilization can differ across workers. Table 5 shows how workers with different levels of education are distributed across job levels. 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 groups that each education group has significant shares of workers ($> 10\%$) across at least three job levels. Second, education is positively correlated with job levels. Workers with more education are found on higher job levels in line with the higher complexity of these jobs. Typically, 60% or more of workers with only secondary education are

Table 5: Share of job levels within formal education and age groups

	at ages 25-35 (in %)					at ages 35-45 (in %)				
Education	1	2	3	4	5	1	2	3	4	5
Males										
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
Females										
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 male workers, the bottom part female workers. 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.

at the two lowest job levels (levels 1 and 2). For workers with a college education, we find that typically 60% or more 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.

5.5 Returns to seniority

Buhai et al. (2014) have shown that also the seniority of workers within the firm is an important factor for wage growth beyond plant, occupation, and pure experience effects. Specifically, Buhai et al. establish that not only a worker's own tenure but also the relative ranking among coworkers determine workers' wages. Similarly, we know from Jäger and Heining (2022) that the wages of workers and the probability of moving within a plant to better-paid jobs increase if coworkers leave the plant (in their case, because of death).

These findings are important from a normative point of view because the effect of coworker characteristics adds an element of luck to wage dynamics. Although workers can change employers and coworkers over time, coworker characteristics can still be considered largely beyond a worker's control—in particular, it is beyond a worker's control whether other workers at the employer are more experienced. From the career ladder perspective, the natural question that arises is whether the returns to seniority that the literature finds are mediated through job characteristics or whether they show up as an independent (residual) factor.

To explore this question, we estimate the effect of the seniority ranking within a plant among a group of peers that might effectively be competitors for career progression. We consider two measures of wages. The first measure is the log wage as reported in the data. The second is the *job-level wage*, constructed as the wage that is predicted by the current job level of a worker using the coefficient estimates from equation (27). We also consider two measures for the seniority ranking. In the first case, we include a dummy only for the most experienced worker within each peer group (based on tenure with the firm). The estimated coefficient quantifies a *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 construct the distance between ranks as $\log(N_i+1-r_i)-\log(N_i)$ where r_i is the seniority rank of worker i within the worker's peer group and N_i is the number of members in worker i's peer group. For example, the most experienced worker within each peer group has seniority rank $r_i = 1$, and the least experienced worker has $r_i = N_i$. We get that within each peer group, the distance between seniority ranks varies between $[-\log(N_i), 0]$. We restrict the sample to male workers because of the different career dynamics for females after age 30. We define a worker's competitive peer group at a plant as the group of workers who are at most five years older than the respective worker and who have the same educational attainment. We construct within each age-education cell of the plant the silverback dummy and the distance of seniority ranks. We run three sets of regressions: First, we regress log wages on controls for the seniority ranking; second, we swap the actual wage with workers' job-level wage; and third, we use the difference between the two as a regressand and to determine the residual seniority premium. Table 6 shows the estimated coefficients.

Table 6: Being the silverback: the effect of experience ranking on job-level wages

	Relative experience concept					
	Silverback effect			Seniority rank		
Wage measure	Raw	Job level	Residual	Raw	Job level	Residual
More experienced	6.8***	4.7***	2.1***	4.7***	3.5***	1.2***
adj. R^2	0.70	0.52	0.63	0.70	0.53	0.63
N	370,792	370,792	370,792	370,792	370,792	370,792

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 three different wage concepts: first, the *raw* log hourly wage of the worker; second, the wage predicted by the worker's *Job level*; and third (*Residual*), the difference between the two (i.e., the wage controlling for a 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), and plant-fixed effects. *, **, *** indicate significance at the 10%, 5%, and 1% levels, respectively.

On average, we find the *silverback effect* and seniority rank distance to be statistically significant. The more experienced a worker the higher is his wage. In the first case, considering only the most experienced worker, we find that these workers obtain a statistically highly significant 6.8% wage premium for seniority based on their raw wages. Their job-level wages are also 4.7% higher, and consequently, there is only a small seniority premium of 2.1% 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 4.7% for raw wages (close to Buhai et al.'s estimates for Denmark and Portugal) and 3.5% for job-level wages and again a much smaller residual seniority 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.

These seniority effects are also economically significant. The coefficient for the *silverback effect* implies that the job-level wage is 4.7 log points higher if a worker is the most experienced worker within his peer group. To put this into perspective, the job-level component accounts for approximately 25 log points in wage growth for 45- to 50-year-old workers, such that being the silverback of a group increases the job-level wage by 19%. To quantify the effect of the seniority rank, note that the average number of members within a peer group is 11. Hence, the difference in the job-level wage between the least experienced member and the most experienced member of an average peer group is 8.1 log points, 32% of the average job-level component at that age. If one views the relative seniority rank in a group of peers as being largely outside the control of a worker, this result suggests an economically significant role of luck in a worker's life-cycle wage dynamics.

5.6 Labor market mobility and career dynamics

This effect of seniority on job-level wages also hints towards the importance of internal job markets for career progression toward higher CAR-levels of jobs. In this section, we corroborate this and explore individual panel data to trace the importance of labor market mobility and employer switching versus internal job markets for career progression. Importantly, we do not explore the complex reasons why workers move to different employers but only explore the consequences of such employer switching.

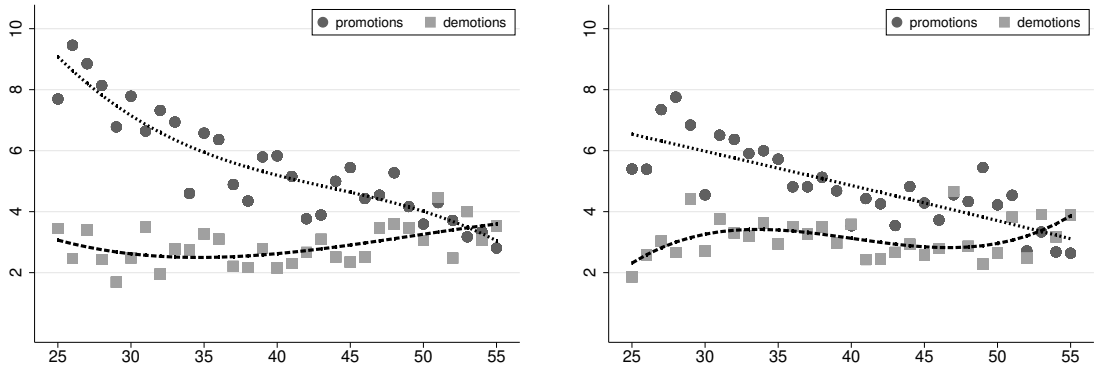
The SES data are limited in their potential to study career dynamics and labor market mobility.[28] The SOEP data provide information on individual labor market situations together with workers' demographics and income (Goebel et al., 2019). See Appendix H for further data details. The data cover the period from 1984 to 2015. As part of these data, the SOEP collects information similar to job levels with a coding that is based on ideas from the sociological literature (Hoffmeyer-Zlotnik, 2003). Compared to job levels encoding CAR intensity, the SOEP coding loads more heavily on education and therefore tends to bias downward worker mobility

[28] In Appendix I, we rely on the information on employer tenure from the SES and document that employment spells with the same employer increase further up on the career ladder in line with an important part of career progression happening at a single employer.

across job levels.[29] With this caveat in mind, we use the job level from the SOEP data to explore worker mobility and career progression.

To align the SOEP sample and the SES sample, we keep workers ages 25 to 55 working at employers with 10 or more employees. We drop self-employed workers, apprentices, military personnel, and public service workers. We drop all observations with missing information on job level, industry, education, occupation, or number of employees at their employer. Data are at an annual frequency, and we explain below how we define labor market mobility events.

Figure 11: Promotion and demotion rates by age



(a) Males (b) Females

Notes: Annual promotion and demotion rates by age for males and females based on SOEP data, years 1984-2015. All rates are shown in percentages. The left panel shows promotion and demotion rates for males, the right panel the promotion and demotion rates for females.

In the first step, we construct life-cycle profiles of promotion and demotion rates. Promotions (demotions) 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. Figure 11 reports estimated life-cycle profiles of annual promotion and demotion rates for males and females. We find declining promotion rates for both genders during working life, in line with a concave wage profile. Males show higher promotion rates in the first part of the life cycle. At age 55, the levels of promotion rates for males and females have converged. Demotion rates are strikingly constant over the entire working life, and levels are very close between males and females. For both genders, demotion rates are substantially below promotion rates at the beginning of working life. In the late 40s, the levels of promotion and demotion rates roughly converge, implying no further net career progression. In Appendix H, we compare net promotion rates, promotion rates minus demotion rates, for males and females. We find that net promotion rates between males and females diverge most strongly between ages 25 and 35 (Figure A16). We return to these differences in promotion patterns between males and females when discussing the gender wage gap through the lens of our model (Section 6).

[29] Conditional on the job level, the SOEP data show quantitatively similar wage differences between job-level age profiles, as found in the SES data. We provide details in Appendix H.

To explore how labor market transitions are associated with career dynamics, we distinguish between stayers, who are consecutively employed by the same employer for two full years and those who change employers. Workers changing employers are either persons who are employed at both survey dates but have been employed for less than one year with the current employer on the second date, or, a subgroup of workers with a nonemployment spell. Another subgroup of workers that we look at is those with an occupation change. These are workers who answer affirmatively to the question of whether "there has been a change in their job" and there is a recorded occupation change.[30] Using these mobility definitions, we ask whether promotions and demotions happen with the same employer or whether labor market mobility is a key driver of promotion and demotion dynamics.

Table 7: Promotions and demotions for stayers and movers

	employer change (%)	stayer (%)
promotion	28.6	71.4
no change	11.8	88.2
demotion	38.1	61.9

Notes: Shares of all promotions and demotions that happen for workers staying with the same employers during the year (column *stayer*) and workers changing employers (column *employer change*). Each row sums to 100%.

Table 7 shows the share of all promotions and all demotions accounted for by stayers and movers. We find that more than 70% of promotions happen for workers who stay with their employer, while less than 30% of all promotions are associated with an employer change. For demotions, we find a similar distribution: about 60% of demotions happen at the same employer, while 40% involve a change of employers. Finally, 88% of workers without a promotion or demotion also stay with their employer. Labor market mobility seems to be no necessary condition for mobility over job levels and most workers with mobility over job levels are stayers.

Table 8 changes perspectives and asks whether labor market mobility is associated with particular promotion patterns. It reports the distribution of promotions, demotions, and lateral moves conditional on employer changes, transitions through nonemployment, and occupation changes. We report stayers and the average across all workers as a reference. Labor market mobility implies more movement on the career ladder. All workers with job changes, be it changing employer, going through non-employment, or changing occupation, show more mobility on the career ladder compared to job stayers. We find that 9% of all employer changes involve a promotion, in line with the idea that workers move to another employer to climb the career ladder. Yet, we also find that 7% of employer changes are associated with a demotion. On net, workers with an employer change have a 20% higher than average probability of career progression. Perhaps surprisingly, we also find that 12% of nonemployment transitions involve a promotion.

[30] 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 8: Promotions and demotions for labor market transitions

	employer change (%)	non-employment (%)	occupation change (%)	stayer (%)	average (%)
demotion	6.6	10.7	11.0	2.2	3.0
no change	84.5	77.3	75.6	93.6	92.0
promotion	9.0	12.0	13.5	4.2	5.0
net promotion	2.4	1.3	2.5	2.0	2.0

Notes: Promotions and demotions 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%.

The promotion in this case is relative to the last job before nonemployment; that is, here we look for at least two-year changes in job levels. Since 11% of all nonemployment transitions involve a demotion, on net, workers after a nonemployment spell experience slower career progression than any other group. Their annualized net promotion rate is at least 70% lower than the rate of the average worker. We observe the strongest career progression for occupation changers, who have a 25% higher net promotion rate than the average worker. Notwithstanding, a change in occupation does not involve a promotion for 87% of all occupation changers (11% demotions, 76% lateral moves). Looking at job stayers, we find that there is substantially less mobility on the career ladder: only 4% of workers move up the career ladder each year, and 2% move down.

5.7 Sensitivity and extensions

We provide an extensive sensitivity analysis to our analysis in Appendix G. In particular, we explore several extensions to our baseline specification from equation (27). In the first step (Section G.1), we explore heterogeneity in the job component of wages across worker groups. We explore 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. Results for wage growth are very similar for full-time male workers, and 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. 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. We relegate further details and discussion to Appendix G.

In the second step, we explore more flexible specifications to equation (27) where we allow returns to experience to be education-specific (Section G.2) and occupation-specific (Section G.3). We find the key result of the importance of career ladder dynamics for wage dynamics to be 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.

Finally, we estimate in Section G.4 the regression in equation (27) 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 A model for a job level perspective on wage dynamics

Our empirical analysis provides evidence for the key role of career dynamics as a driver of wage dynamics. In this section, we develop a stylized model to study whether the empirical evidence is consistent with a quantitative theory of career dynamics. A positive answer will provide us with a framework to study the implications of differential career dynamics for life-cycle wage dynamics. In order to keep the model tractable, we rely on a reduced form for the job levels and abstain from a microfoundation of job differences. Career dynamics in the model are determined by the organizational structure of employers, and we study the model in partial equilibrium with an inelastic supply of jobs. Thus, we study the consequences of career ladder dynamics for wage dynamics rather than providing a theory of why career ladders and job levels exist.

6.1 Setup

We consider a model of ex-ante homogeneous workers and firms. Firms are multi-job establishments with jobs having different job levels labeled by $i = 1, 2, \ldots, n$. We assume that a production structure with n workers constitutes a firm and we will not provide a theory of firm size n but take the number of jobs n as given.[31] For simplicity, we focus on an organizational structure where there is exactly one job at every level at each firm.

We denote the job level of a worker by $e \in 1, 2, \ldots, n$. In addition, $e = 0$ denotes unemployment. Conversely, the employment state of a firm $f \in 0, 1, 2, \ldots, n$ describes whether all jobs in that firm are filled ($f = 0$) or whether the firm has an opening at level $f = i$. For tractability, we

[31] Calvo and Wellisz (1979) provide a theory of the firm size with hierarchical layers.

rule out that firms can have more than one open position in a period. The state of a worker in a firm is the tuple $\{e, f\}$ where we set $e = f = 0$ for unemployed workers. We assume that not only the unemployed search but also employed workers search for an opening at a higher level if they work in a firm that has all its jobs filled, i.e., with firm-employment state $f = 0$. Since f denotes the vacant job level for firms with vacancies, combinations with a vacancy ($f > 0$) but a worker at that job level $e = f$ cannot exist.

We normalize the mass of workers to unity. For firms, we abstract from entry and exit and normalize the mass of jobs to unity. The state of the economy determining worker mobility is the entire joint distribution over firm and worker types $\{\mu_{e,f}\}_{e=0,\dots,n,f=0,\dots,n}$, where $\mu_{e,f}$ is the share of workers on job level e in a firm with job level f being vacant. In line with the definition above, $\mu_{0,0}$ denotes the share of unemployed workers. We assume that a share of employed workers always searches for a better job and firms always try to fill vacant positions. Next, we describe how the distribution evolves.

The labor market is frictional. A firm with a vacant position randomly meets searching workers. We assume that firms cannot downgrade an incumbent worker in order to hire another worker. After meeting a worker, the firm can hire the worker only if the current level of that worker e is smaller than their vacancy level f. What is more, we assume, in line with our results on seniority, that firms will offer a worker they meet the lowest position necessary to poach this worker and promote their more senior employees instead to fill the actual vacant position.

Regarding the mobility of workers, we assume that in each firm with full employment ($f = 0$), nature draws which worker searches in a given period. In firms with a vacancy ($f \neq 0$), workers cannot search. The draw is i.i.d. and the probability for a worker in position e is $\nu = 1/n$. The selected worker is forced (by nature) to separate into unemployment (with probability δ) and if there is no separation into unemployment searches for alternative jobs. She meets firms with vacancies with probability λ that results from a standard matching function described below. Conditional on contacting a searching firm, $\phi(f)$ denotes the probability to meet a firm with an unfilled position at level f. The probability equals the share of these positions among all vacant ones (which have mass $1 - \mu(0, 0)$):

$$\phi(f) = \frac{\sum_{e \geq 1} \mu(e, f)}{1 - \sum_e \mu(e, 0)}.$$

Conversely, the probability from the firm's perspective that a searching worker is currently employed at job level e is given by

$$\omega(e) = \frac{\mu(e, 0)\nu}{\mu(0, 0) + \sum_{e>0} \mu(e, 0)\nu}.$$

With these definitions, the probability of a worker moving from state $\{e, f\}$ to $\{e', f'\}$ is given

by

$$\pi(\{e,f\},\{e',f'\}) = \begin{cases} 1-\lambda & \text{for } e=0, f=0, e'=0, f'=0 \\ \lambda & \text{for } e=0, f=0, e'=1, f'=0 \\ \nu\delta & \text{for } e>0, f=0, e'=0, f'=0 \\ 1-\sum_h \nu[\delta+(1-\delta)\lambda\sum_{k>h}\phi(k)] & \text{for } e>0, f=0, e'=e, f'=0 \\ \nu(1-\delta)\lambda\phi(e') & \text{for } e>0, f=0, e'>e, f'=0 \\ \nu[\delta+(1-\delta)\lambda\sum_{k>f'}\phi(k)] & \text{for } e>0, f=0, e'=e, f'>0, f'\neq e \\ \chi\sum_{f>k\geq e}\omega(k) & \text{for } e>0, f>e, e'=e, f'=0 \\ \chi\sum_{k<e}\omega(k) & \text{for } e>0, f>e, e'=e+1, f'=0 \\ \chi\sum_{f>k}\omega(k) & \text{for } e>0, f<e, e'=e, f'=0 \\ 1-\chi[1-\sum_{k\geq f}\omega(k)] & \text{for } e>0, f>0, e'=e, f'=f \\ 0 & \text{for all other cases} \end{cases} \tag{29}$$

where χ is the contact rate, i.e., the probability for a firm with a vacant job to meet a searching worker. The steady-state distribution of workers over jobs μ^* is a fixed point of the mapping induced by $\Pi(\mu)\mu$, where Π is the stacked transition matrix generated by equation (29).

The first case in (29) shows the probability of an unemployed worker remaining unemployed. Since unemployed workers will always enter the job ladder at the lowest level, the second case gives the probability of an unemployed worker finding a job, which is the probability to meet any searching firm. The third case reflects the probability of transitions into unemployment. The fourth case is the probability that no worker leaves a full-employment firm, where the term in the summation of the probabilities of worker at level h receiving the mobility shock with probability ν and either leaving into unemployment (δ) or finding a better job elsewhere ($\sum_{k>h}\lambda\phi(k)$). The next line is the probability that the worker at level e finds a better job elsewhere and leaves the firm, the sixth case is the worker on level f' leaving the firm with the coworker on level e staying. The next four cases describe workers in firms with unfilled positions. Line seven is the probability that the firm with a vacant position at a level higher than worker e becomes a full-employment firm ($f'=0$) by hiring a worker who is currently employed at a level higher than e, such that the incumbent worker e is not promoted. The next line gives the probability that the worker is promoted because the firm fills the vacancy by hiring a lower-level worker. Line nine is the probability of a firm with an opening below the worker of level e to fill that position, which requires meeting a worker currently working on a sufficiently low level. Line ten is the probability of a firm with a vacancy being unable to fill this vacancy. The last line covers all other cases. This description of the different cases of worker mobility and career dynamics shows that our stylized model offers already very rich dynamics that require tracking the within-firm distributions and current job distribution of searchers. and we explore if a calibrated version of the model is able to account for the empirical career-ladder wage dynamics documented in our empirical analysis.

We use the matching function from Den Haan et al. (2000) that is

$$M = \frac{SV}{(V^{\rho} + S^{\rho})^{\frac{1}{\rho}}} \tag{30}$$

where M denotes the number of matches, $S = \sum_e \mu(e, 0)$ the mass of searching workers on and off the job, and $V = \sum_e \sum_{f>0} \mu(e, f)/(n-1)$ denotes the number of vacant positions. The contact rates in (29) are then $\lambda = \frac{M}{S}$ and $\chi = \frac{M}{V}$. A match in our case does not necessarily lead to a worker transition as employed workers reject a job offer if the offered position does not yield an improvement relative to their current job level.

We calibrate the model at a monthly frequency for prime-age males in the German labor market. We set $n = 5$ to align with the five encoded job levels in the SES data. We take wages as exogenous and calibrate wages to the estimated job-level wages from our empirical analysis.[32] The model then has only two free parameters that need to be calibrated internally. We calibrate the separation probability δ and the parameter of the matching function ρ such that the model matches the average transition rates into and out of unemployment for prime-age males over the time period from 2007 to 2018 based on German social security data. Specifically, we match monthly transition rates into unemployment of 0.63% and out of unemployment of 7.98%. The calibrated parameters are $\delta = 0.0466$ and $\rho = 0.3451$.

We solve the model by iterating on the transition matrix $\Pi(\mu)$ of the joint distribution until convergence, updating at each iteration step contact rates λ and χ and job-offer distributions according to the prevailing distribution across worker and firm types.

To obtain life-cycle implications of the model, we use the stationary transition matrix and simulate a cross-section of workers (with mass zero) that enters the labor market as unemployed at age 18.[33] We drop the first seven years and follow workers in the model for 30 years so that we only consider prime-age working life as in the data.

6.2 Implications for average career progression and wage dispersion

The model abstracts from any other wage dynamics than the one from job levels. In particular, it does not contain any worker or firm component. We, therefore, extract the job-level component from the empirical results and contrast the model prediction with the estimated life-cycle pattern of the job-level wage component. We first compare the model outcomes to the empirical findings for males as the model has been calibrated to the labor market dynamics of males.

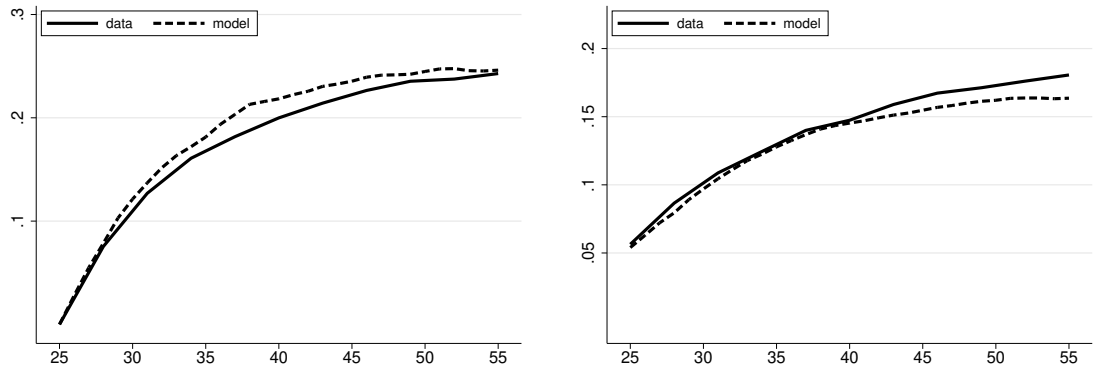
Figure 12(a) shows the average male job component from the data and the model prediction for average job-level wages. We find that model and data align closely, lending support to the model's career-ladder dynamics. Importantly, wages in the model only match average job-level differences but we do not target life-cycle dynamics. The arising life-cycle patterns of wage

[32] We provide in Appendix J a microfoundation for the calibrated wage differences based on Winter (2004) and Winter (2006).

[33] We start at age 18 as 86% of workers in our SES sample are non-college workers who typically enter the labor market at ages between 16 and 18 years.

dynamics are endogenous and untargeted. Figure 12(b) compares the variance of model wages to the data counterpart.[34] The figure shows that the model matches closely the heterogeneity in career progression and the resulting increase in wage inequality, too.

Figure 12: Wage dynamics from model and data



(a) Mean job component

(b) Variance of job component

Notes: Life cycle wage growth and wage inequality for prime-age males from model and data. The model simulates a cohort (of mass zero) that enters the economy as unemployed at age 18 using the stationary transition matrix implied by the model. Workers are on 5 job levels and wages for these levels are taken from the estimation in Section 5.

In summary, we find that our stylized model of career dynamics is qualitatively and quantitatively consistent with the empirically observed life-cycle pattern of job-level wage growth and inequality. Thereby it supports our interpretation of the empirical wage dynamics as career ladder dynamics. We will now use this theory to provide a new perspective on wage facts.

6.3 The gender wage gap through the lens of our model

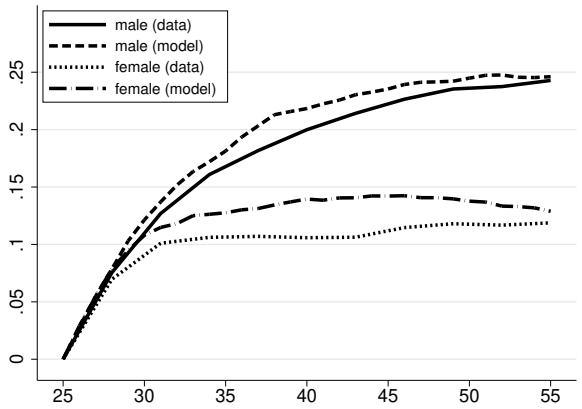
The fact that our new model framework matches qualitatively and quantitatively life-cycle wage dynamics offers us the opportunity to understand other labor market phenomena through its lens.

We first revisit the gender wage gap as a widely studied empirical wage phenomenon. At age 25, females in our sample receive a roughly 7% lower wage than males. At age 50, females earn wages that are more than 30% lower than wages for males. As raw averages, these differences still contain occupational and employer differences. Our empirical analysis shows that more than half of the increase in the difference comes from the job component (cf. Figures 8 (a) and (c)).

Figure 13 directly compares the development of the job component from Figure 8 across genders. In the data, males and females experience a virtually identical increase in the job-level wage component until age 28, but after age 28, the career progression of females comes to a halt,

[34]The variance also includes the covariance terms of the job component with the individual and plant components.

Figure 13: Life-cycle profile of gender wage gap



Notes: Gender wage gap in the model and the data. The model simulation shown as blue lines and red dots show estimated job level components in the data for males and females. In the model, female workers are hit by non-mobility shocks starting at age 28 that render them non-mobile ($\nu_i = 0$) for three years. Male workers follow the model description from before.

while males keep on climbing the career ladder for an additional 15 to 20 years. The differential progression in terms of job levels leads to male wages exceeding female ones by more than 10 log points at the age of 50.[35] These results support the idea that the gender-wage gap largely stems from a gender-promotion gap and differences in career ladder dynamics (see also Bronson and Thoursie (2018)).

Our model allows us to understand if and to what extent differences in mobility affect job-level progression. To quantify the effect, we simulate a cohort of female workers where we model the difference to male workers by introducing an immobility shock. Male and female workers are otherwise identical. The immobility shock starts to hit female workers starting at age 28. If a female worker is hit by the immobility shock, she becomes immobile for three years. A worker can be affected repeatedly by immobility. If immobile, the worker will not change jobs, but may still move into non-employment, i.e. the transition matrix Π for immobile workers is an identity matrix with a positive probability of separating into unemployment that we will calibrate to changing female employment shares.

To have a data counterpart to target the immobility shock, we identify immobility as part-time work. We match the difference in the share of males and females in part-time at the age of 40 (a 51.5 percentage point difference). This gives us a monthly probability of $p_n = 2.25\%$ of being hit by the immobility shock.[36] In addition to the sharp increase in part-time work with age for women, we also observe a decline in employment after age 28. We target the decline in the female employment share between age 25 and age 40 (a decline of 1.5 percentage points) by

[35] The data span twelve years, so the estimated life-cycle pattern also comes from comparisons across cohorts. Yet, in Section 5.6 we documented career ladder dynamics between males and females in SOEP data that support the idea that women do not climb the career ladder as much as men. The SOEP data have the advantage over the SES data that they offer panel data for more than 30 years.

[36] The share of immobile workers in all workers is roughly $\frac{p_n}{p_n + \frac{1}{\text{length of spell}}}$.

allowing for a positive separation rate into unemployment for immobile workers. The calibrated probability separation probability is 61.9% of the separation rate of typical workers ($0.619\nu\delta$).

Figure 13 demonstrates that such a dynamic view on the gender wage gap is consistent with the empirical evidence on wage dynamics. The immobility shock in the model represents in a reduced form that, in the German context, it is typically women who reduce labor supply in their mid-career by going on maternity leave, working part-time, or changing to jobs with more family-friendly work requirements. Through the lens of our career ladder model, this will lead to a dynamically arising gender promotion gap that materializes in a gender wage gap in the later part of working life. Our results suggest that the gender wage gap is dynamic with its roots early in working life when a gender wage gap might still be small but differences in career progression will leave their long shadow on the future of female careers.[37]

6.4 Returns to education

Since returns to education are empirically mediated through job levels, our model provides a laboratory for this phenomenon, too. We introduce a notion of education by extending the model to have low- and high-skill worker types. The organization of the production process is such that high-skill workers are always put on higher job levels than low-skill workers when both are present in a firm. Put differently, we assume perfect assortative matching within the firm between job levels and education. Note however that this within firm sorting does not rule out having low-skill workers on all job levels in some firms.

Specifically, we assume that within a firm workers are ordered lexicographically with skill level being the most important dimension.[38] This implies that low-skill workers are demoted if a high-skill worker is hired on a lower job level. Yet, within each education group, the career dynamics apply as described for the baseline model. As a consequence, we get that high-skill/college workers will climb the career ladder more quickly. The assumption of a lexigraphic ordering leads to a tractable within-firm worker-job distribution as it is summarized by the highest job level of a low-skill worker within each firm.

For the calibration, we match that 86% of workers do not have a college degree in our SES sample. We proceed otherwise with the same parameters as before. Table 9 adds to our regression results from Table 4 the corresponding results for the simulated model data. The estimated return to college education is a 64% higher wage in the model. This return is close to the one in the data of 54%. In the model, returns will tend to fall once we introduce deviations from the perfect assortative matching of education types and job levels within firms. By construction, there are never returns to education in the model once job levels are controlled

[37] Note that our model likely still underestimates the consequences of career interruption if there is additional congestion by males on the career ladder so that women have a harder time catching up after a non-mobility period. Our current simulations assume for tractability separate labor markets for males and females as otherwise, the entire job allocation within a firm, i.e. which position is filled by a male or female becomes a state variable of the problem that would render the model intractable.

[38] One way to microfound such an ordering is to assume the organization structure of production as in Winter (2006) with high-skill workers having lower effort costs.

for a finding in line with the data.

Table 9: Education returns in model and data

	data		model	
	baseline	w/o all job info	baseline	w/o job levels
College	-0.01	0.54***	0.00	0.64***

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 three different specifications: first our baseline, second a specification that leaves out job-level information, and third a specification that leaves out job information (levels and occupations) altogether. The baseline education category is vocational training. *, **, *** indicate significance at the 10%, 5%, and 1% levels, respectively.

In other words, we provide an interpretation of the returns to education as a reduced-form wage fact as differences in career progression between workers with different educational backgrounds. Importantly, this means that returns to education are not independent of the typical organizational structure in an economy and will change if firms reorganize their production processes toward production structures that provide more jobs at higher job levels (Krusell et al., 2000). Importantly, high-level jobs are not identical to management jobs but simply involve more complexity, autonomy, and responsibility (CAR) in task execution (higher CAR intensity) which connects our analysis to the underlying ideas in Autor et al. (2003) and Acemoglu and Autor (2011).

6.5 Returns to seniority

Similar to returns to education, we found that returns to seniority are largely mediated through job levels (Table 6). Seniority at an employer and wage differences across similar workers suggest a role for differences in career progression. Again, our model provides a laboratory to study also this phenomenon. For this purpose, we simulate firms rather than worker careers as returns to seniority require capturing coworker dynamics. In the data, the average peer group in our regression sample consists of 11 workers. The size of the peer group determines the size of the returns to seniority and we therefore group workers of 5-worker firms in the model together into larger firms. We do this by combining several simulated firms targeting the average peer group size from the data. The model does not explicitly include age as a state variable. We, therefore, use tenure as a proxy instead. We restrict the simulation sample to observations with 4 to 21 years of tenure in line with the interquartile range of the regression sample (male workers, age 45-50). The average tenure of the simulated workers is 10 years compared to 13 years in the data. Finally, to match the organizational structure of firms and the distribution of workers in the sample across job levels, we construct weights for the simulated model data to be in line with the empirical estimates.[39] We use the model with two education types as before and

[39]Returns to seniority are a function of the coworker characteristics and the sample composition determines the size of the estimated returns.

construct the peer group as in the data by taking all workers within the firm with the same educational attainment.

Table 10: Experience rank and job levels

Data		Model	
Silverback effect	Seniority rank	Silverback effect	Seniority rank
4.7	3.5	6.4	6.0

For the simulated model sample, we then regress the log wages on the seniority rank and the silverback dummy.[40] Table 10 shows the empirical estimates for job-level wages together with the model counterpart. We find that as in the data the silverback effect is slightly larger than the effect of the seniority rank and although the model-based returns are slightly higher they are in a very similar order of magnitude. Hence, the model accounts for the observed returns to seniority in the data. The causal explanation of the model for these returns is luck as workers cannot affect coworker mobility at the firm. However, the slightly higher model returns also point to a non-luck component in career progression. In summary, these results suggest that exploring career dynamics is likely a fruitful avenue for future research.

7 Conclusions

This paper provides a new answer to the long-standing question of what determines a worker's wage. We document that the job level of a worker's job accounts for a large fraction of the observed wage differences across workers and over the life cycle. We explain that a job's level describes the complexity, autonomy, and responsibility (CAR intensity) of task execution. While occupations and the derived task-based approach describe *what* tasks a job holder executes, job levels capture *how* tasks are executed. Thus, job levels capture differences both within and across occupations. Our work thus builds on and refines the influential task-based view of wage determination (Autor et al., 2003).

Wages by job level have been an important part of labor market statistics since the 1950s. They continue to form the basis of collective bargaining by unions and job-based compensation schemes in firms. Conceptually, job levels are related to the organization of production, and we explain and document the difference between job levels and occupations, education, and task content of jobs. Using direct information on task execution from survey data, we demonstrate the explanatory power of job-leveling factors for wages and thereby establish the economic content of the statistical concept of job levels. Using high-quality microdata from Germany, we decompose life-cycle wage dynamics and show the key role of changing job levels, climbing the career ladder, in accounting for life-cycle wage growth and rising wage inequality. We find

[40] As in the data, we multiply log wages by 100.

that career ladder dynamics account for about half of wage growth and virtually all of the increase in wage dispersion over the life cycle. We also document that labor market mobility is associated with career progression, but that most moves up and down the career ladder occur with the same employer. Although we focus on the German labor market, we also document the importance of differences in job levels in accounting for wage differentials in the U.S. labor market.

We develop a structural labor market model of career dynamics and show that the calibrated model is consistent with documented empirical life-cycle wage dynamics. The model provides us with a laboratory to elaborate our new perspective on wage dynamics as career ladder dynamics. We use the model to study widely documented wage phenomena such as the gender wage gap, the returns to education, or the returns to seniority. We find that the model, despite its simplicity, can explain all three phenomena. Importantly for future work, our results suggest that wage dynamics are closely related to the organization of production, so that the organization of production and its changes ultimately determine the wage dynamics of workers in the macroeconomy.

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.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. The Quarterly Journal of Economics, 109(3):659–684.

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.

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 Jr, R. E., and Rossi-Hansberg, E. (2018). Learning, career paths, and the distribution of wages. American Economic Journal: Macroeconomics, forthcoming.

Caliendo, L., Monte, F., and Rossi-Hansberg, E. (2015). The Anatomy of French Production Hierarchies. Journal of Political Economy, 123(4):809–852.

Calvo, G. A. and Wellisz, S. (1979). Hierarchy, ability, and income distribution. Journal of political Economy, 87(5, Part 1):991–1010.

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., Gurun, U., and Ozel, N. B. (2023). Too many managers: The strategic use of titles to avoid overtime payments. Working Paper 30826, National Bureau of Economic Research.

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.

Den Haan, W. J., Ramey, G., and Watson, J. (2000). Job destruction and propagation of shocks. American economic review, 90(3):482–498.

Dengler, K., Matthes, B., and Paulus, W. (2014). Occupational tasks in the German labour market. FDZ Methodenreport, 12.

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.

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.

Gibbs, M., Ierulli, K., and Milgrom, E. M. (2003). Careers in firm and occupational labor markets. Technical report.

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. (2014). The glass ceiling and the paper floor: Gender differences among top earners, 1981-2012. Technical report, National Bureau of Economic Research.

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. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. Journal of Economic Dynamics and Control, 17(5-6):953–969.

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., Qiu, X., et al. (2022). Coordinated firm-level work processes and macroeconomic resilience. Technical report, University of Bonn and University of Cologne, Germany.

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.

Morchio, I. and Moser, C. (2020). The gender pay gap: Micro sources and macro consequences. Available at SSRN 3176868.

Pastorino, E. (2019). Careers in firms: the role of learning and human capital. Technical report, mimeo.

Pham-Dao, L. (2019). Public insurance and wealth inequality—a euro area analysis.

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. (2015). Firming up inequality. Technical report, National Bureau of Economic Research.

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 Lohnstrukturerhebungen 1998 bis 2006. Technical report, Schweizer Bundesamt für Statistik.

Tåg, J., Åstebro, T., and Thompson, P. (2013). Hierarchies, the small firm effect, and entrepreneurship: Evidence from Swedish microdata. Technical report, IFN working paper.

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.

US 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 Job levels and occupations in 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 look at wages of assemblers and fabricators in the US and Germany as a case study.

A.1 Evidence from the NCS

The NCS is a nationally representative employer survey conducted by the Bureau of Labor Statistics (BLS) that collects information from private industry as well as state and local government establishments. The survey collects detailed job characteristics that are encoded as job levels using the BLS job-leveling system.[41] 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.[42] 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. (2023) 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 4.3 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

[41] See 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 C.

[42] We provide a case study for assemblers and fabricators in Appendix A.2 below to demonstrate that the BLS job levels summarize job differences that are similar to the job levels in the German data.

Table A1: Mean wages in 2015 by job level and occupational group

	Occupational groups (SOC)					
Level	11-29	31-39	41-43	45-49	51-53	All
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.

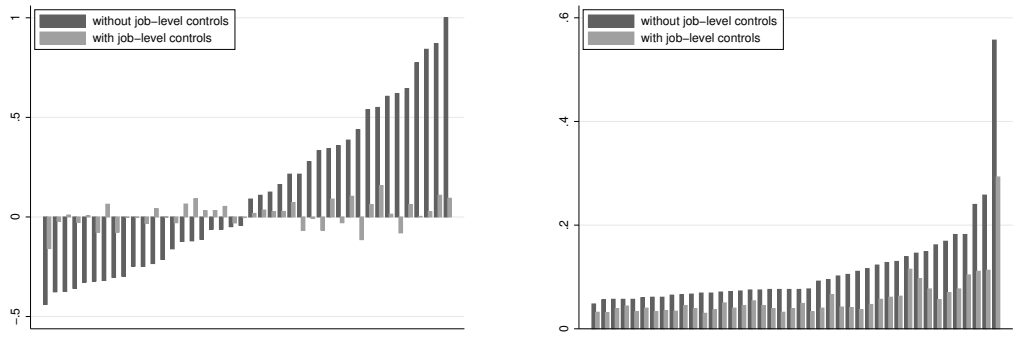
dispersion compared to the SES data. When all employer and job information is included, observables account for 85% of cross-sectional wage dispersion ($R^2 = 0.847$, Pierce (1999), Table 4), and job-leveling factors alone account for 75% of wage variation. These results corroborate key findings from 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 high explanatory power of job levels as an additional dimension of task execution accounts also in US data for a large part of wage dispersion, so this finding is not a particularity of the German labor market and its institutions.

Next, we explore similar to our analysis in Section 4.2 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 A1 shows mean wages by job level and occupational group from the 2015 NCS.[43] We see that within coarse occupational groups, there is a wide variation of wages across job levels. 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

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

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), 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 A1: Occupation wage premia and within-occupation wage dispersion



(a) occupation wage premia

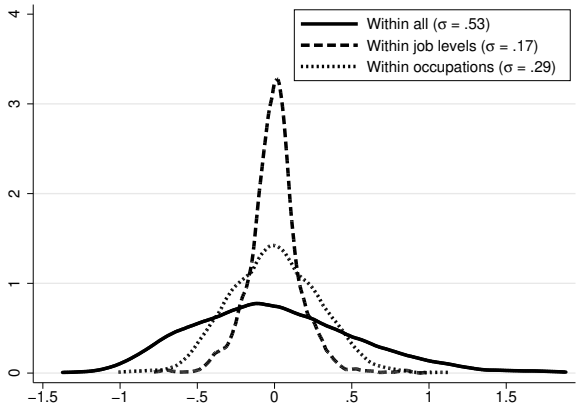
(b) within-occupation wage dispersion

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 of the 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 4. Figure A1 visualizes results from Table 7 in Pierce (1999). Figure A1(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 A1(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 A1). 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 A1(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 4 on the distinction between job levels and occupations.

Figure A2: US 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 4.2, 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 now use equivalent occupation-job-level cell data from the NCS to conduct the same analysis on U.S. data. Figure A2 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 than job levels in the decomposition. Figure A2 shows density estimates for residual (log) wages for the equivalent three cases from Figure 6. 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.[44] 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.

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

A.2 Case study of within-occupation job-level differences

To further substantiate the differences between occupations and job levels and to highlight that these differences also apply beyond the German case, we consider a case study for a narrowly defined occupation group: *assemblers and fabricators in production*. For our case study, we start with the German union bargaining agreement for metal- and steelworkers in North-Rhine-Westphalia. This union bargaining agreement has at its core an analytic job-leveling scheme to assign workers to wage scales; it is closely comparable to the BLS job-leveling scheme. Together with the job level, we observe the bargained wage for each job level.[45] For assemblers (*Montierer*) and fabricators (*Maschinen- und Anlagenbauer*), we have job-leveling information that distinguishes these occupations at six different job levels: four job levels for the occupation group assemblers and two for fabricators.[46] We start from the German job-leveling information (i.e., specific job descriptions regarding tasks and duties of the jobholder) and assign job levels based on the BLS job-leveling guide. Using the resulting U.S. job levels, we assign wages for full-time workers from the tabulations for production occupations from the NCS in 2010.[47] In the NCS data, we stay within a single occupation group according to the classification in the 2000 SOC System and use only wages at different job levels. After leveling the German jobs using the BLS procedure, we remove mean wage differences between Germany and the United States so that the average across the assigned wages is one in both countries. Hence, we classify German workers as if they worked in the U.S. labor market and compare their German pay to their U.S. counterparts in identical occupations and on the same U.S. job level.

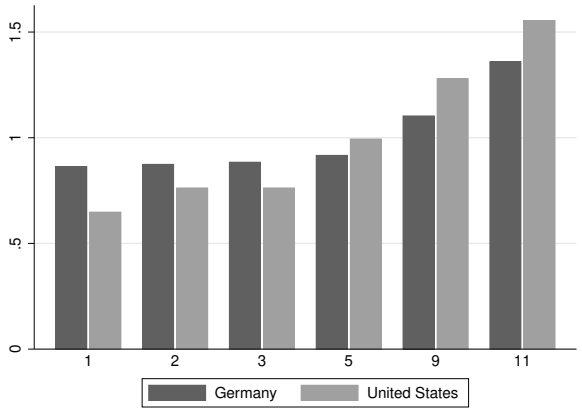
Figure A3 shows the standardized wage differences across job levels for Germany and the United States. We find that wage structures show a very similar shape across countries, with the key difference being that the German wage structure shows more wage compression in the lower part which is typically 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 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

[45] These bargained wages are lower bounds for wages and are typically supplemented by performance components that are worker- and firm-specific.

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

[47] United 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.

Figure A3: 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.

defined occupational group as further evidence for the importance of job levels for determining wages and wage differences in Germany and the United States.

B Job levels and task-based classification of jobs

Section 4 documents that job levels capture the CAR intensity of a job's task execution and that they have strong explanatory power for wages. This explanatory power of job characteristics complements the idea from 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 task-based approach to job leveling. First, the task-based approach is derived from occupation-level information (not as the Cartesian product of occupations and job levels) so that it does not differentiate within occupations, while job levels provide within-occupation differentiation (Section 4.2 and Appendix A). Second, one way to

interpret the task-based approach is that it projects occupational tasks on their amenability to being executed by a computer.[48] This projection aligns most closely with autonomy that enters the CAR intensity measurement of the job level but it does not relate directly to responsibility and complexity. Even for autonomy, there would be no distinction between the bakers if one baker decides about the amount of the ingredients and the baking time and the other baker mixes the ingredients following closely the recipe of the former.

In Section 4.2, we document 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 other occupations.[49] This variation in average job level allows us to shed some more light on the relationship of job levels and the cognitive/non-routine classification of occupations.

We rely for our analysis on previous work that has implemented the task-based approach (Spitz-Oener, 2006; Dengler et al., 2014). For the task-based classification, Dengler et al. (2014) follow closely the original approach by Autor et al. (2003) by relying on expert assessments of job task contents. Spitz-Oener (2006) classifies occupations based on BIBB/BAuA survey data on workers to assign tasks to occupations. We use the classification by Dengler et al. (2014) based on 2013 occupational tasks to the 2018 SES data aggregated to the three-digit occupation level. Before aggregating the SES microdata, we apply the sample selection as described in Section 4. 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 A2: 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).

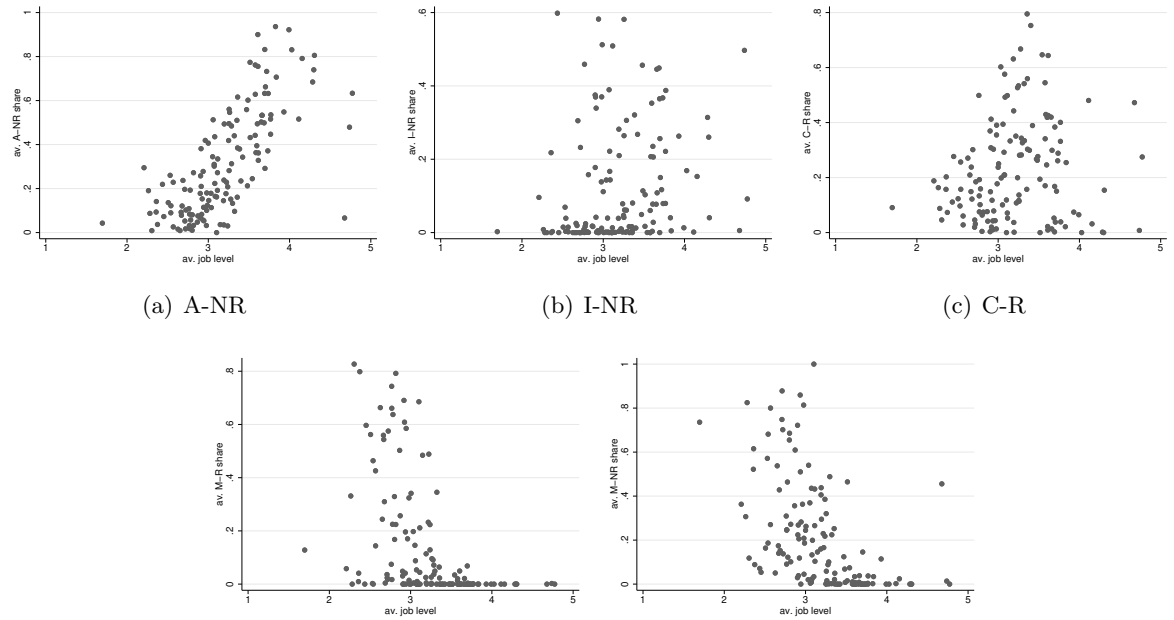
In Table A2, we look at correlations between the average job level and task shares. The key conclusion of the task-based approach is that routine tasks can be replaced by computers and that non-routine tasks are relative complements to computer capital. In line with the fact that autonomy is one of the 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

[48] Examples of tasks from Appendix Table 1 in Autor et al. (2003) are "computes discount, interest, profit, and loss," "mixes and bakes ingredients according to recipes."

[49] See Table A1 for this fact based on U.S. data.

most positively with the average job level. Manual routine (M-R) correlates the most negatively with the job level but also for non-routine 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 A4: Tasks and job levels



(a) A-NR (b) I-NR (c) C-R

(d) M-R (e) M-NR

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 in more detail, Figure A4 shows scatter plots of the average job level and the shares of the different task components across the 140 occupations. Looking at Figure A4(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 A4(b), the data are much more dispersed and a positive relationship is less striking. The cognitive routine tasks in Figure A4(c) show a positive relationship, yet again, there is also substantial dispersion. For the manual routine tasks (M-R) in Figure A4(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 A4(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

Table A3: 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***	0.54***						
	(0.00)	(0.00)						
A-NR		-0.23**	0.25	0.63***				
		(0.01)	(0.06)	(0.00)				
I-NR		-0.23*	-0.35*		0.11			
		(0.02)	(0.03)		(0.50)			
M-R		0.08	-0.27			-0.32**		
		(0.35)	(0.05)			(0.00)		
M-NR		-0.16*	-0.51***				-0.58***	
		(0.04)	(0.00)				(0.00)	
C-R								0.34*
								(0.01)
N	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.

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.

Given the observed correlation between occupational task contents and job levels from Figure A4, we next ask how much each component contributes to occupational wage differences. 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 A3 shows 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 4.2 and Appendix A of the large explanatory power of job levels on between-occupation wage differences.

C Fifth occupation digit and job levels

Table A4: 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.

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 job-level information in these data. Table A4 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 over the very fine-grained occupational codes.

D 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 4.2. 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 4.2 focuses on white-collar workers. Finally, we compare the wages by job level constructed from the survey data to the actual bargained wages by job level.

D.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.[50] The job-leveling system has four components: required skills and knowledge, autonomy, cooperation 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)

[50] See 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/Tarifkarte_ERA_2010-2012_englisch_01.pdf (accessed May 22, 2019).

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?

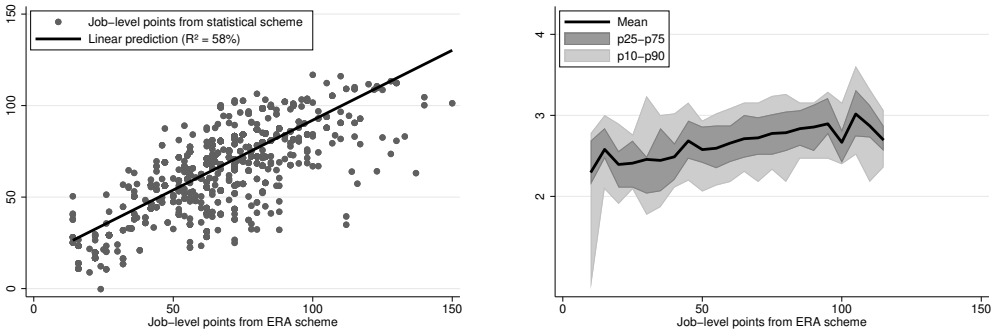
To apply the job-leveling scheme, we have to assign job-level points to answers from the BIBB/BAuA survey. The point range of the job-leveling system is from 10 to 170 points and we apply the following assignment of points. For the skills part, we assign 10 points if a quick briefing is sufficient and 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*.

D.2 Results for blue-collar workers

In Section 4.3, we restricted the sample to white-collar workers, Figure A5(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 4.3) 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 A5(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 dom-

Figure A5: Job-level points and average wages by job-level points (blue-collar workers)



(a) Statistical scheme

(b) Distribution (by groups of 5 points)

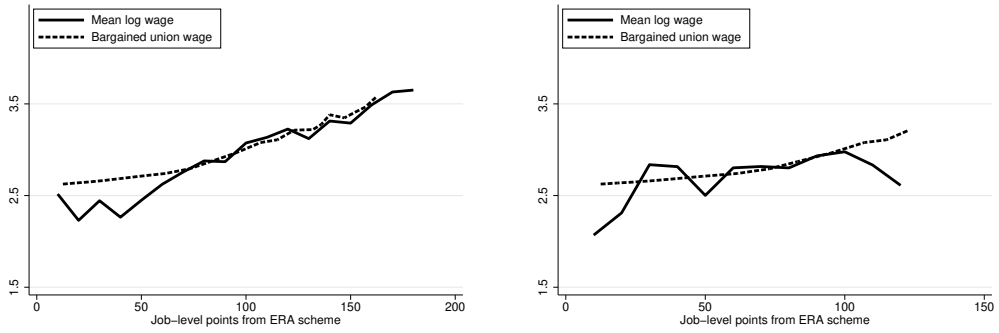
Notes: Left: Scatter plot of a worker's implied job-level points from statistical job-leveling scheme against the worker's job-level points from union bargaining scheme (ERA scheme). The statistical job-level scheme is based on the regression of wages on survey answers. The solid line shows the linear 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).

inates 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.

D.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 A6 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 A6: Average and bargained wages by job-level points for North-Rhine-Westphalia (blue-collar workers)



(a) White-collar workers

(b) 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).

E Identification and instrumental variable regression

Our analysis addresses two key identification challenges that are motivated by theoretical 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 then the means to prevent other employers from poaching 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. Under this view, all jobs are individually set up for the individual worker skills. 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 G.4, 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

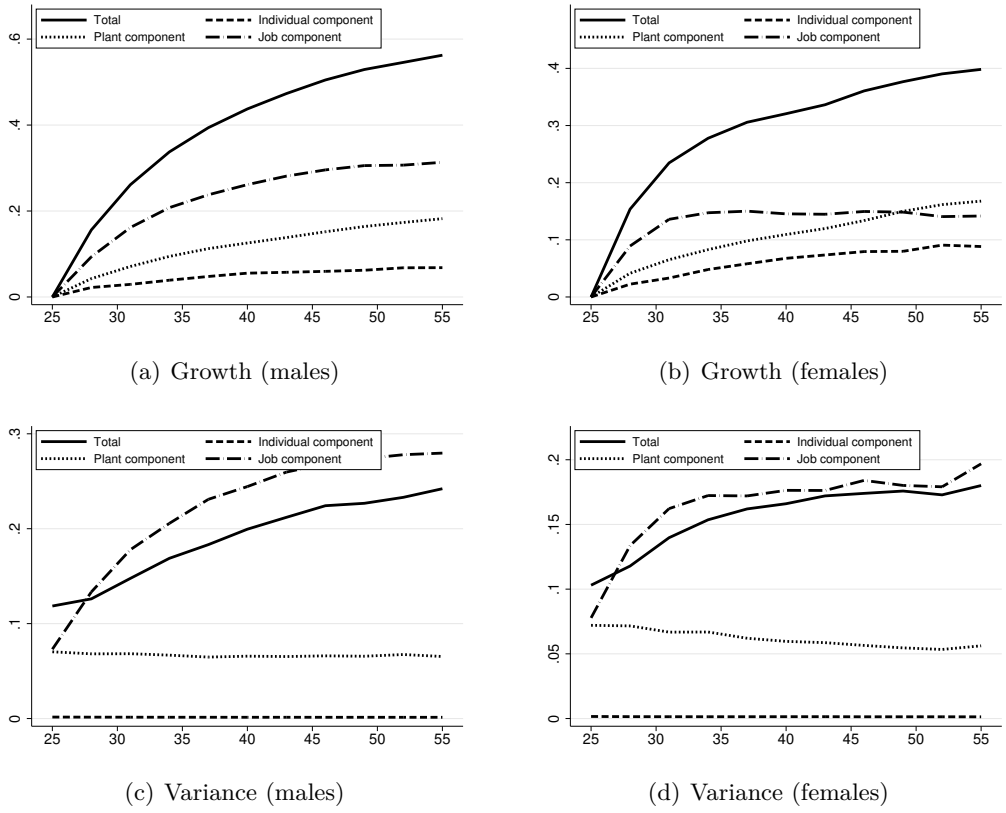
the associated wages as prizes implies that wages only represent a prize for previous performance but not remuneration for task execution on the current job. Unlike in the task-based approach, it is not the executed tasks on the current job that determine the wage but a worker's past performance. If wages are prizes, differences in a job's tasks will not be systematically related to wages. In Section 4.3, 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 such information on 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.

E.1 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 A7 shows the resulting decomposition results for wage growth and wage inequality for males (Figures A7(a) and A7(c)) and for females (Figures A7(b) and A7(d)).

In the decomposition of wage growth, we find that for both males (Figure A7(a)) and females (Figure A7(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 males (Figure A7(c)) and females (Figure A7(d)), the relative importance of the job component increases. For males, the job component even exceeds the total increase of the variance when we do not account for the covariances. In the case of females, 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.

Figure A7: Decomposition of wage growth and wage dispersion over the life cycle using IV approach



(a) Growth (males)

(b) Growth (females)

(c) Variance (males)

(d) Variance (females)

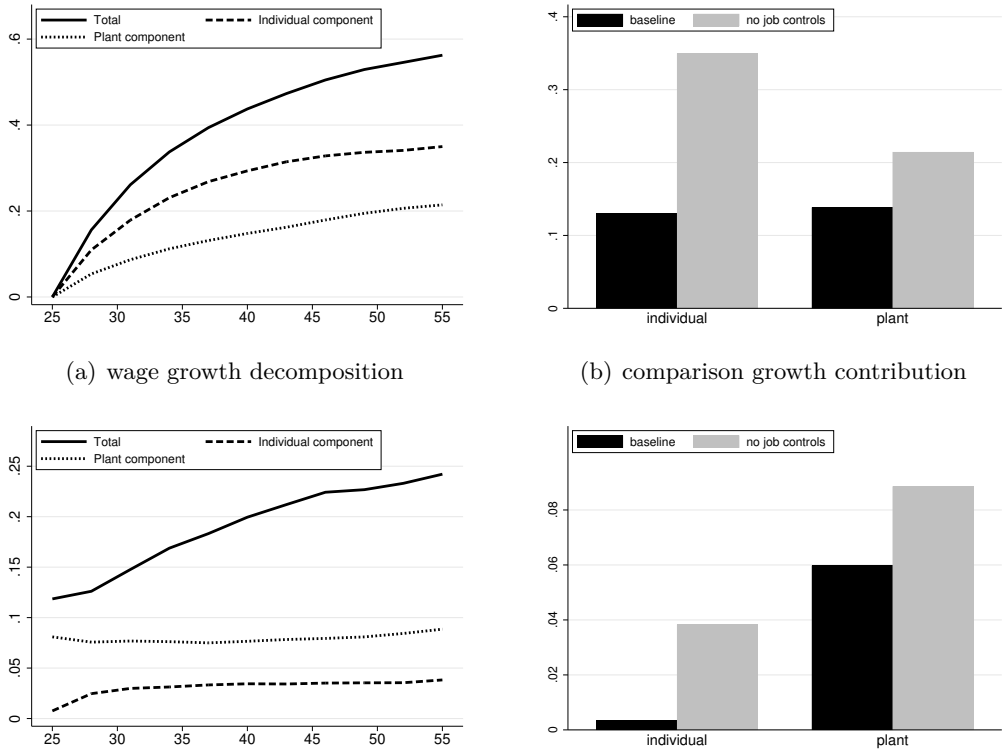
Notes: Contribution of the job component to wage growth (top row) and wage dispersion (bottom row) for males (left panels) and females (right panels). The solid line shows the job component for the baseline from the main part of the paper; the short dashed line shows the case with no collective bargaining 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).

F Wage decomposition without job component

The plant component in our 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 the plant. To explore the importance of the job component and job composition for the wage decomposition, we repeat the wage decomposition from equation (27) but drop the job component. We then compare the resulting plant and individual components to those from our baseline decomposition.

Figure A8 shows the decomposition of life-cycle wage dynamics for males. Comparing the decomposition of wage growth in panel (a) to our baseline in Figure 8, we draw qualitatively very different conclusions about the sources of life-cycle wage growth and hence about the sources of wage growth that a theory of wage dynamics should entail. We find that the experience

Figure A8: Decomposition of wage growth and variance of wages by age (males), ignoring job controls



(a) wage growth decomposition

(b) comparison growth contribution

(c) wage variance decomposition

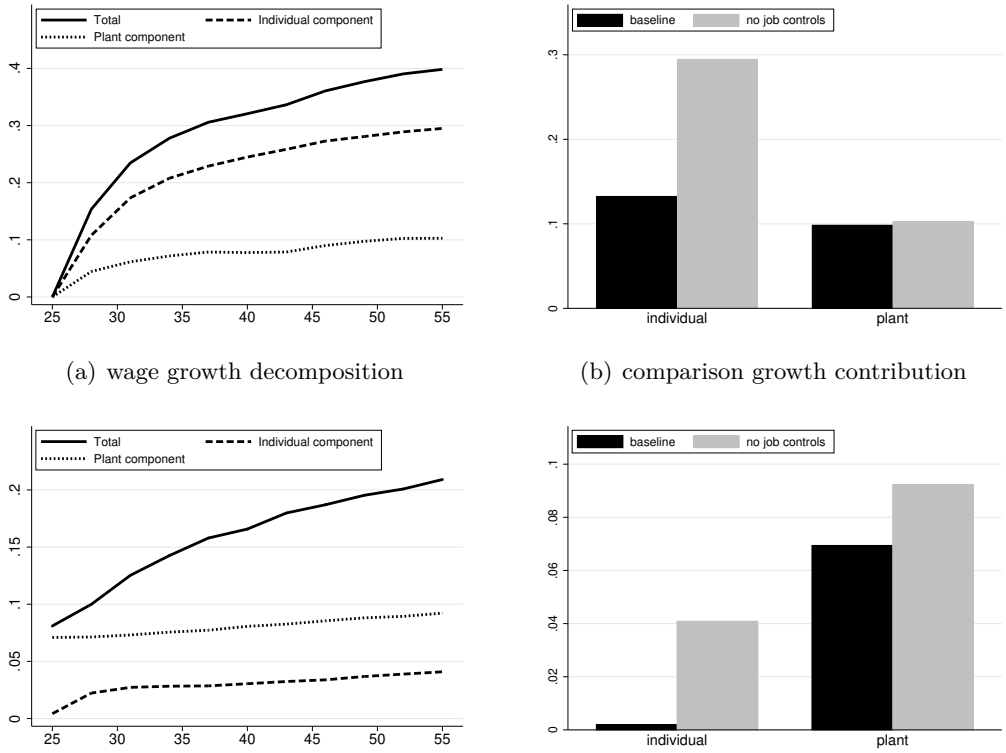
(d) comparison variance contribution

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 males. 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).

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 now remains unexplained. Figure A8 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 organizational structure and occupational composition of plants also leads to an increase in the plant component but this increase is comparatively modest.

Figure A8(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 this fraction grows over the working life. The implied rising dispersion of the residual wage component has been traditionally interpreted as persistent

Figure A9: Decomposition of wage growth and variance of wages by age (females), ignoring job controls



(a) wage growth decomposition

(b) comparison growth contribution

(c) wage variance decomposition

(d) comparison variance contribution

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 females. 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. The finding that most of this dispersion stems from the career ladder does not invalidate the interpretation as risk but asks for further investigation of the reasons for career progression and provides a new view on how changes in the macroeconomy might lead to changing labor market risk. Being able to point to the changing job levels as a source of wage dynamics opens new opportunities to understand labor market risk or to even scrutinize the interpretation of rising inequality as risk. Finally, Figure A8(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 A9 reports the results for the life-cycle wage dynamics for females and the changes in the individual and plant component compared to the baseline decomposition. Looking at the decomposition results for wage growth in Figures A9(a) and A9(b), we draw generally the same conclusions as from the corresponding decomposition for males. The individual component for

wage growth picks up almost all wage growth. For the variance in Figures A9(c) and A9(d), we also get that as for males the individual component increases and accounts now for a sizable fraction of the life-cycle increase in wage inequality.

Contrasting these results with our baseline offers an important insight for why employers are important for life-cycle wage dynamics. Differences in employers 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 as our baseline decomposition shows.

G 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 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 G.1. As extensions to our baseline results, Sections G.2 and G.3 explore more flexible specifications for the wage equation. Section G.4 reports results if instead of a synthetic cohort panel approach, we rely on a pooled OLS regression when decomposing wages.

G.1 Heterogeneous returns to job and individual characteristics

For the first set of sensitivity checks, we interact variables from the baseline regression in equation (27) 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 A5, 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 A5 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 females is large: 60% of employees at public employers are female. 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 A6 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 A5: 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 A6 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.[51] 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

[51]This 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 A6: Test statistics for coefficient tests

	no collective bargaining		only full-time		large plants	
	p-value	*F-stat*	*p-value*	*F-stat*	*p-value*	*F-stat*
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.

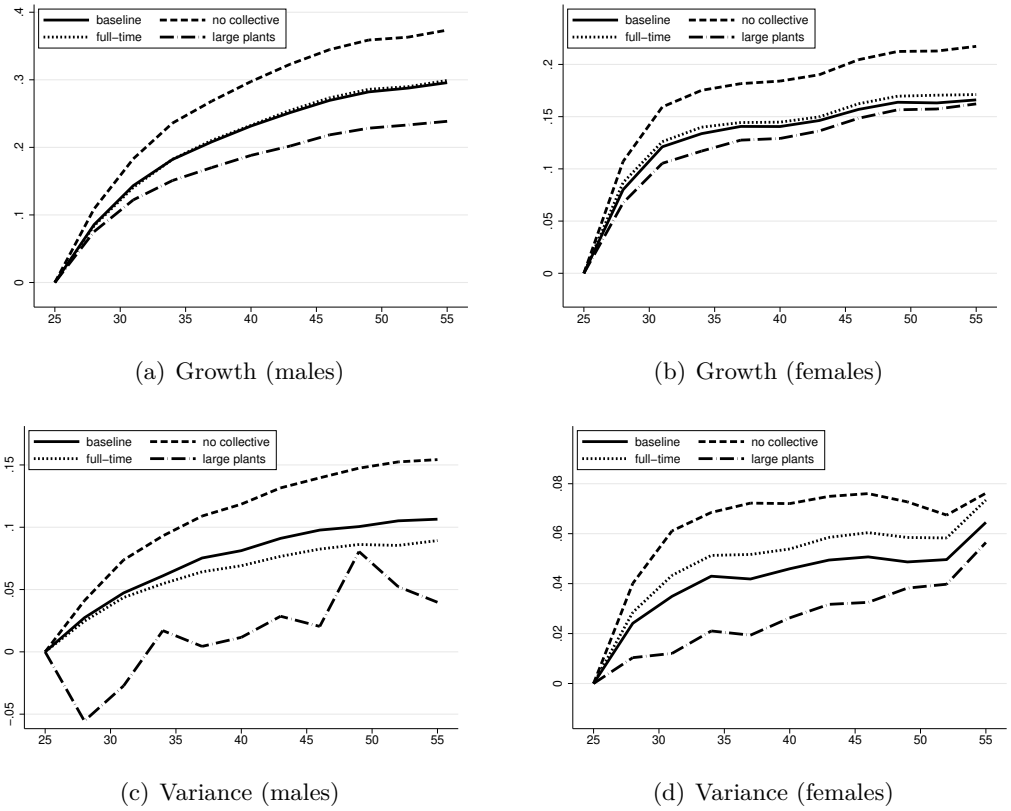
presentation in the main part of the paper, we removed level differences at age 25 for easier comparison.

Looking first at the case of no collective bargaining, we find the age-wage profile (for the job 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 A.2). 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 females and the same increase for males. 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. In our model analysis of the gender wage gap (Section 6.3), we keep the job-level wage unchanged but non-mobility changes the distribution over job levels. The results for full-time workers here are therefore consistent with the result from the model that the difference in the job component stems from differences in the distribution across job levels rather than different wages.

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 males are negligible. The more notable effect is for females. Including public employers adds slightly less than a third to the job component for female wage growth. This finding

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



(a) Growth (males)

(b) Growth (females)

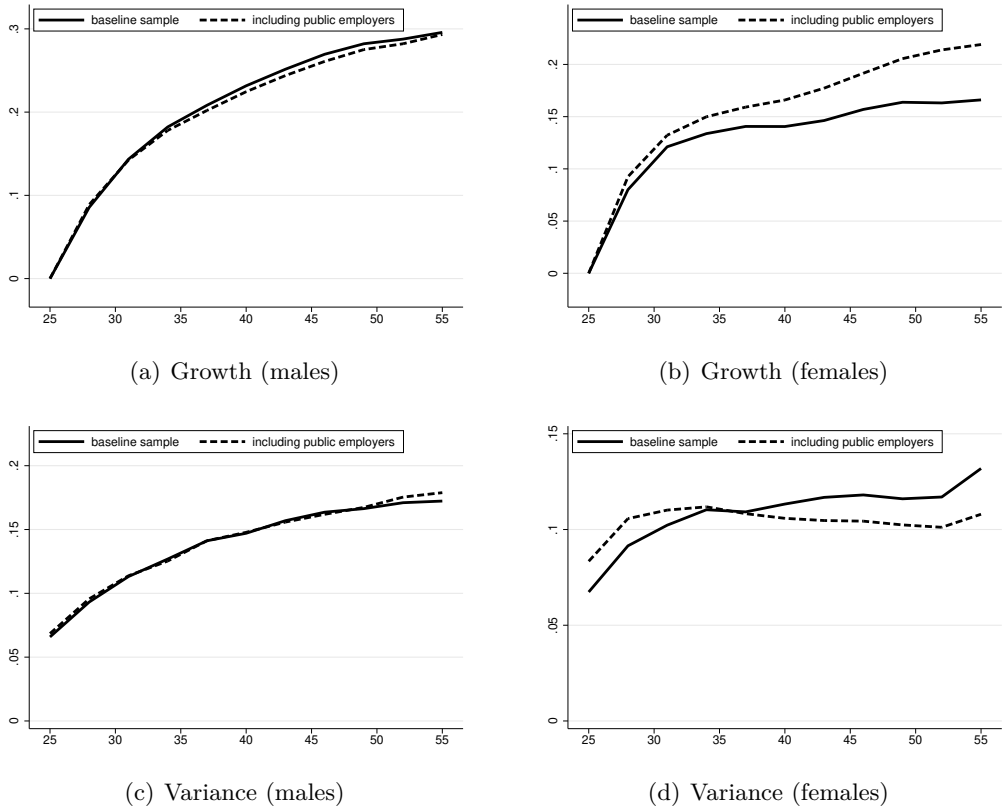
(c) Variance (males)

(d) Variance (females)

Notes: Contribution of the job component to wage growth (top row) and wage dispersion (bottom row) for males (left panels) and females (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).

suggests that public employers are an important contributor to female career progression after age 35 and that females 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 females is substantially smaller once we include public employers in our sample.

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



(a) Growth (males)

(b) Growth (females)

(c) Variance (males)

(d) Variance (females)

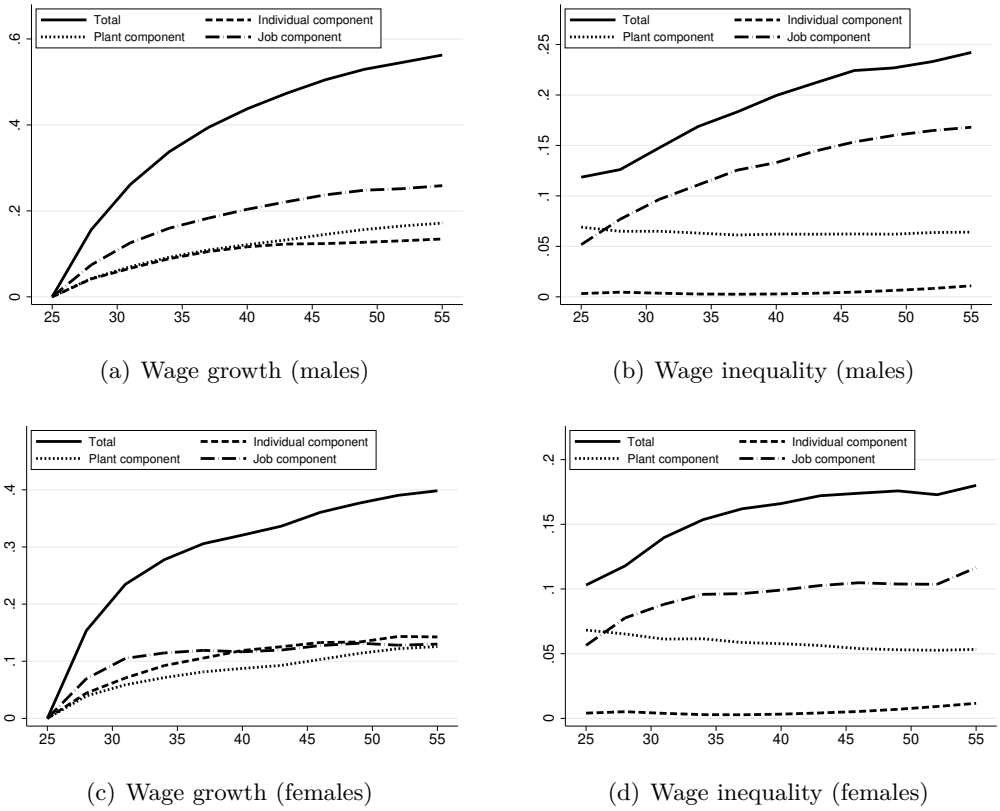
Notes: Contribution of the job component to wage growth (top row) and wage dispersion (bottom row) for males (left panels) and females (right panels). The solid line shows the job component for the baseline from the main part of the paper; the dashed line shows results for a sample including public employers 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).

G.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 males and females 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). 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 males and females for this extended wage regression.

We find our decomposition results to be very similar under this extended wage specification.

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



(a) Wage growth (males)

(b) Wage inequality (males)

(c) Wage growth (females)

(d) Wage inequality (females)

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.

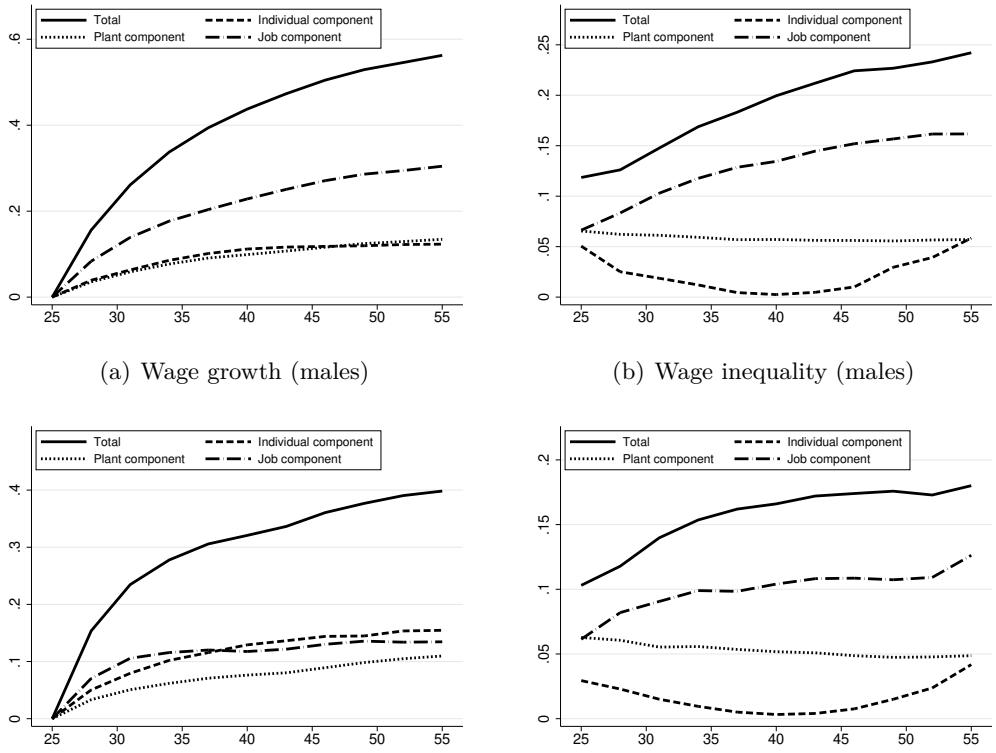
For wage growth in Figures 13(a) and 13(c), the job component declines slightly for males and females but remains for males by far the most important driver of wage growth. For females, all three components account now for a third of wage growth at the end of working life. For males and females, the plant component gains slightly in importance. For males, it becomes more important than the individual component. For females, 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 males and females. Hence, we do not find evidence that the job component is inflated by picking up an education-specific skill accumulation effect.

Our analysis in Section 6.4 also explores the relationship between job levels and education. Instead of augmenting the baseline regression from equation (27), we drop the job component

to explore how the omission of the job component affects estimated returns to education. We find that in this case the effects are large and provide the interpretation that returns to education arise from career ladder dynamics.

G.3 Occupation-specific returns to experience

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



(a) Wage growth (males)

(b) Wage inequality (males)

(c) Wage growth (females)

(d) Wage inequality (females)

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 males and females 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 males and a third for females. The contribution of the individual and plant components remains largely unaffected for males. For females, 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 component for males and females becomes U-shaped over the life cycle with a countervailing effect (not shown) from the individual-job covariance term.

G.4 Pooled regression without individual fixed effects

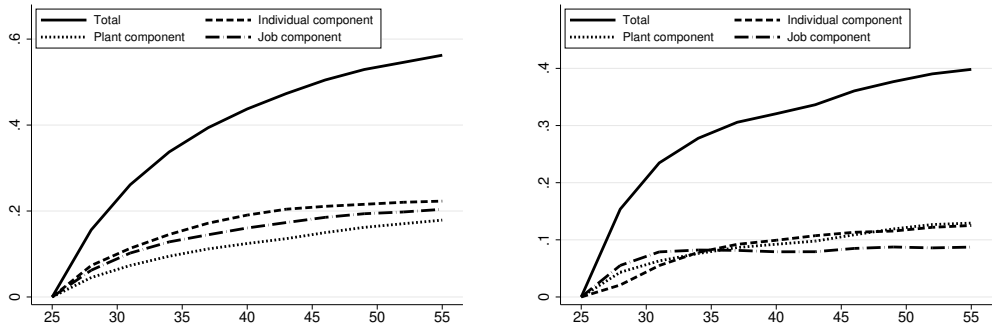
The main part of the paper uses synthetic cohorts to control for individual fixed effects that are arguably correlated with education, career progression, and potentially employer types. In this section, we run as an alternative specification a pooled OLS regression controlling for cohort effects but not controlling for individual fixed effects. Specifically, we set $\hat{\gamma}_i = \gamma_c$ in equation (26) and instead run the following regression on the pooled data:

$$\hat{w}_{it} = \gamma_c + \beta_J \hat{J}_{it} + \beta_I \hat{I}_{it} + \hat{\epsilon}_{it}. \tag{31}$$

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 A14 shows the decomposition of wage growth in the individual, plant, and job components if we do not control for individual fixed effects.

Comparing the decomposition results for wage growth to the baseline results in Figure 8 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 males and females, 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.

Figure A14: Wage decomposition for males and females without controlling for individual fixed effects



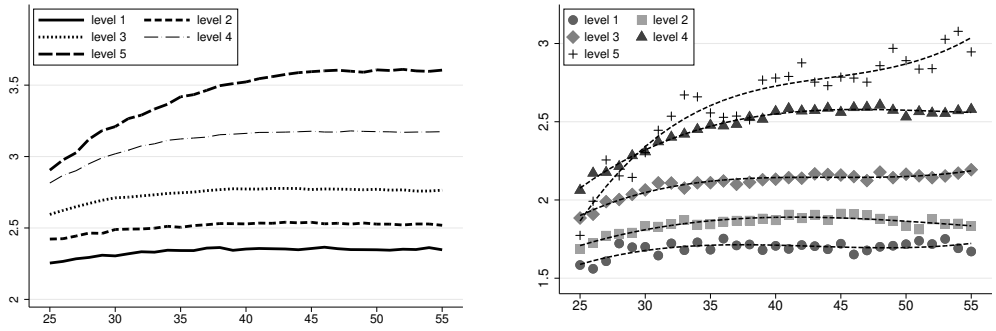
(a) Males (b) Females

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).

H Socio-Economic Panel (SOEP) data

The German Socio-Economic Panel (SOEP) data are the equivalent to the US Panel Study of Income Dynamics (PSID) data. The SOEP data provide individual-level panel data that cover the period from 1984 to 2015. This section provides additional information on wages, job levels, and career progression of males and females from the SOEP data.

Figure A15: Wage by age and job level



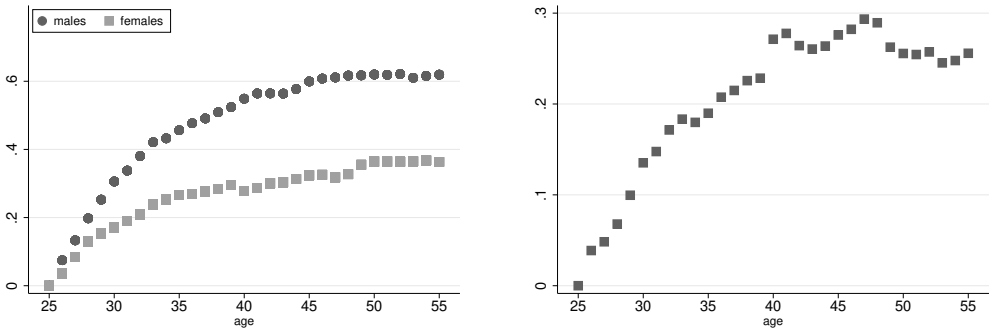
(a) SES (b) SOEP

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 A15 compares (log) wages by age and job level from the SES and SOEP data. The data span different time periods

so that average levels of wages differ and job levels are not directly comparable because of different coding approaches (see Section 5.6). 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.[52] 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. A key difference that is related to the different coding approaches is the strong increase in wages on job level 5 in the first part of the working life. This finding reflects that compared to the SES data, the SOEP job-level data have a smaller top group with less mobility between groups.

Figure A16: Cumulated net promotions and demotions by age



(a) Cumulated net promotions

(b) Differential cumulated net promotions

Notes: The figures display the accumulated net promotion rates for male and female workers from the SOEP data (1984-2015) and their differences across gender by age.

Figure A16 complements the findings from Figure 13 in the main part of the paper. In Figure 13 of the main text, we document the differences in the job-level component for males and females by age and observe a widening around age 28 when female careers slow down considerably. Figure A16 uses the information on promotions and demotions from the SOEP data. We exploit the panel dimension of the data to accumulate promotions and demotions at the individual level. We summarize the life-cycle promotion dynamics by net promotions where we sum over all promotions up to a certain age and subtract all demotions. Figure A16(a) shows accumulated net promotion profiles for males and females. The vertical axis shows the average net promotion so that a number of 0.5 means that up to this age, every second worker moved up one job level on net. For males, we find cumulative net promotions of 0.6 at age 55, and for females less than 0.4 net promotions. The net promotion profiles trace the dynamics of the job-level components from Figure 13 in the main part of the paper; in particular, we observe a strong slowdown of promotion dynamics for females after age 30. Figure A16(b) shows the difference in net promotions for males and females. Unlike for the job-level component from Figure 13, we already see a widening of promotion dynamics at age 25 that continues up to age 40 when the difference in net promotion stabilizes at about 0.25. This implies that every fourth net promotion for males is not taking place for females and that this difference arises in the first

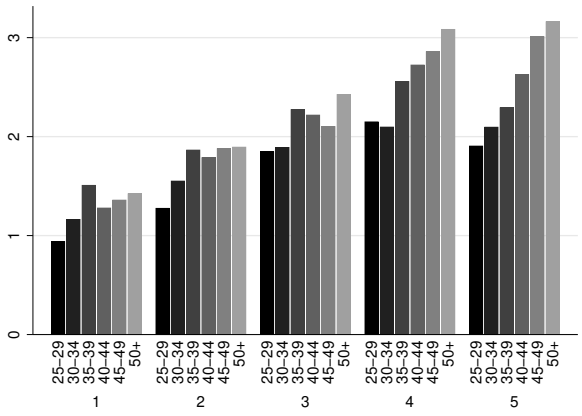
[52] Because of missing hours information, we construct wage data only from 1990 onward.

half of the working life. However, we need to take into account that the SES data and the SOEP data cover different survey years when comparing the results.

I Career ladder dynamics in SES data

The SES data come as repeated cross-sections and do not allow tracking workers and their career ladder dynamics over time. We therefore rely on panel data from the SOEP to trace out career dynamics of workers and their relationship to employer changes. The SES data provide however high-quality data on employer tenure that provide indirect information on employer changes and career progression. Every employer change will end a worker's tenure with the current employer so low tenure is an indirect measure of higher mobility rates. We rely on these data from the SES to explore in Figure A17 how long workers stay with their employers at different stages of the career ladder. Specifically, the figure shows by how much tenure increases (in years) from one five-year age group to the next five-year age group, at different job levels. If all workers stayed with their employer all the time, the increase between age groups would be five. We find that the tenure increase by age is larger at higher job levels and that the increase accelerates over workers' careers. The steeper increase across job levels and age supports our finding based on SOEP data that many workers climb the career ladder while staying with their employer.

Figure A17: Tenure increase by age and job level



Notes: The figure displays the average additional years of tenure of an age group relative to the preceding one by job level. Averages over all sample years are shown for both males and females. For 25- to 29-year-olds, the figure shows the average number of years of tenure in the group.

J Model extension

The model in the main part of the paper sets wages across jobs to the empirical evidence on job-level wages. In this section, we provide a microfoundation of wage differences across otherwise identical workers that has been developed in Winter (2004) and Winter (2006).

The production in each firm is as described in Winter (2006) where workers work together to produce one unit of a final good. Each worker can decide to provide effort to increase the probability of success of her own task. The effort provision is $d \in \{0, 1\}$ and if workers exert effort ($d = 1$) their probability of success is 1. If workers do not exert effort ($d = 0$), their probability of success is $\alpha_i \in (0, 1)$. Exerting effort is associated with cost c. The success probability of the final product is the product of the individual task outcomes. Workers execute tasks sequentially so that workers towards the end of the production process can observe the effort of their upstream co-workers. For the firm, effort is not verifiable and contractible but only the final output. For the simplest case with $\alpha_i = \alpha$ for all jobs $i = 1, 2, \ldots, n$, Winter (2006) shows that the sequential production structure leads to an optimal incentive inducing wage scheme, i.e. all workers choose $d = 1$, with

$$w_i = \frac{c}{1 - \alpha^{n-i+1}}. \tag{32}$$

Winter (2004) derives a similar result where identical workers are paid differently depending on their assigned position in the production process for the case of simultaneous effort choice. As the wage scheme is always incentive inducing, so that $d_i = 1$ for all $f = 0$ firms, the dynamics of the economy can be captured by a law of motion of workers and firms across states.

If we set $n = 5$, we have c and α as the two parameters to be calibrated. When we calibrate these parameters to the estimated job-level wages that are currently set exogenously, we get $c = 2.51$ and $\alpha = 0.79$ as calibrated parameters and a very close fit between the model and data. Table A7 shows model wages and job-level wages from the data. We adjust the lowest job level to match exactly.

Table A7: Calibrated wages for microfounded wage setting

job level	1	2	3	4	5
data	2.34	2.37	2.66	3.02	3.43
model	2.34	2.47	2.65	2.94	3.48

Notes: Calibrated (log) wages from the model by Winter (2004) with $n = 5$ job levels and calibrated parameters $c = 2.51$ and $\alpha = 0.79$.

The model is stylized and does not capture the full richness of task execution encoded in job levels. Its underlying organizational structure of the production process allows however to relate it to the ideas of the paper and it can provide a basis to develop a microfounded model of wage setting in models of job levels.