Superstar Teams

Additional Material (Not for Publication)

Lukas B. Freund§

This document contains additional material for the paper "Superstar Teams." The indexing of this Supplemental Appendix continues that of the Online Appendix.

E Theoretical appendix II

E.1 Brown-Resnick model for skills

The main text uses Fréchet marginals together with a Gumbel-Hougaard copula with distance-dependent parametrization to capture correlation of skills across pairs of workers. This approach is highly tractable but leaves open the question how to generalize it, in the sense of constructing the joint distribution for $n > 2$ Fréchet random variables with distance-dependent association. This section sketches an approach that relies on constructing a spatial, max-stable model for the skills process. See Davison *et al.* (2019) for a survey of models of spatial extremes. To introduce the ideas, and without loss of generality, I work with a unit Fréchet shape parameter (Resnick, 1987, Proposition 5.11).

Suppose skills are the realizations of a non-negative random field $\{Z(s) : s \in \mathcal{C}\}$, where $\mathcal{C} = \mathcal{R} \times \mathbb{R}^+$ is a cylinder, with $\mu \in \mathcal{R} = [0, 2)$ representing the circular coordinate and $h \in \mathbb{R}^+$ denoting height. We interpret $s = (\mu, h) \in \mathcal{C}$ as the latent position of an individual worker. As in the main text, the distance function is defined on the circle as $d : \mathcal{R} \times \mathcal{R} \to \mathbb{R}^+$

$$d_{il} = d(\mu_i, \mu_l) = \min\{|\mu_i - \mu_l|, 2 - |\mu_i - \mu_l|\}.$$

The field $Z(s)$ is constructed as

$$Z(s) = x(h)\tilde{Z}(\mu), \quad s = (\mu, h) \in \mathcal{C}.$$

where $\tilde{Z}(\mu)$ is a Brown-Resnick max-stable process (Brown and Resnick, 1977) and $x(h) > 0$ the height-dependent Fréchet scale parameter.

Under the de Haan spectral representation (De Haan (1984), also see Kabluchko *et al.* (2009)), this process can be written as

$$\tilde{Z}(\mu) = \max_{i \geq 1} \zeta_i W_i(\mu), \quad \mu \in \mathcal{R},$$

§*Email*: lukas.beat.freund@gmail.com. *Web*: www.lukasfreund.com.

where $\{\zeta_i\}$ are points of a Poisson process on $(0, \infty)$ with intensity $\zeta^{-2}d\zeta$; $W_i(\mu)$ are independent copies of the spectral function $W(r) = \exp\{\varepsilon(\mu) - \varrho(\mu)\}$, where $\varepsilon(\mu)$ is a stationary Gaussian process with mean zero and stationary increments; and the semi-variogram $\varrho(\mu_1, \mu_2)$ determines dependence and is isotropic, i.e., $\varrho(\mu_1, \mu_2) = \varrho(d(\mu_1, \mu_2))$.

By construction, the process $\tilde{Z}(r)$ has unit Fréchet marginals:

$$P(\tilde{Z}(\mu) \leq z) = \exp(-z^{-1}), \quad z > 0.$$

Furthermore, for a max-stable random vector with unit Fréchet margins, the d-dimensional joint distribution can be written as

$$\Pr(Z_1 \leq z_1, \ldots, Z_d \leq z_d) = \exp\{-V(z_1, \ldots, z_d)\}, \quad z_i > 0,$$

where $V(z_1, \ldots, z_d)$ is called the exponent function and is homogeneous of order -1.[E.1]

In the Brown-Resnick model, for any *pair* of locations $\{s_1, s_2\} \subset C$ the bivariate exponent function is

$$V(z_1, z_2) = \frac{1}{z_1}\Phi\left(\frac{1}{a}\log\frac{z_2}{z_1} + \frac{a}{2}\right) + \frac{1}{z_2}\Phi\left(\frac{1}{a}\log\frac{z_1}{z_2} + \frac{a}{2}\right),$$

where $a^2 = \varrho(d(\mu_1, \mu_2))$. Huser and Davison (2013) provide expressions for V in the higher-dimensional case.

Writing down an explicit model for the max-stable skills *process* ensure that the finite-dimensional marginal are max-stable with coherent dependence structure. Moreover, the aggregation properties summarized in Lemma A.3 obtain for a max-stable skills process under mild restrictions, yielding a system of equations that can be solved for output Y produced by a team of n workers given knowledge of their locations $(s_1, \ldots, s_n)$. The Gumbel copula distinctively has an L^p-norm like structure, so the partial derivatives are power functions, which leads to factorable forms that yield a closed-form expression for $f(\cdot)$. While deriving such an explicit expression is not feasible more generally, as already noted in the Appendix O, in finite samples the Hüsler–Reiss copula is statistically indistinguishable from the Gumbel copula.

[E.1] Relating this to the notion of a copula, an extreme-value copula C can be written as $C(u_1, \ldots, u_d) = \exp\{-G(-\log u_1, \ldots, -\log u_d)\}, ; u_i \in (0, 1)$, where G is the tail dependence function and homogeneous of degree 1. By the probability integral transform, for Fréchet margins, we have $V(z_1, \ldots, z_d) = G(\frac{1}{z_1}, \ldots, \frac{1}{z_d})$. See Appendix O A.1.2 for an expression for G in the case of the Gumbel copula.

F Empirical appendix II

F.1 Constructing worker types

My baseline measure of worker talent types is based on their position in the economy-wide lifetime earnings distribution, which I recover from the individual fixed effect (FE) in a two-way fixed-effects wage regression à la Abowd *et al.* (1999, "AKM" henceforth). It bears emphasis that in this paper I am not concerned with the estimation of employer FEs or an AKM-based variance decomposition, so prominent debates in the literature about the structural interpretation of these terms can be side-stepped. Instead, I simply adopt this approach as a statistical tool to recover the time-invariant component of an individual's earnings ability that conveniently allows controlling for the effects of unmodelled person-level observable characteristics and employer heterogeneity.

Estimation concerns that relate to limited mobility bias affecting pertain to the estimation of *employer* FEs, and specifically their variance (as well as a the covariance term) are less relevant for person FEs. Nonetheless, to follow best practices, I initially cluster similar firms through a weighted k-means problem, similar to Bonhomme *et al.* (2019),

$$\min_{k(1),\dots,k(J),H_1,\dots,H_K} \sum_{j=1}^{J} n_j \int (\hat{F}_j(w) - H_{K_j}(w))^2 d\mu(w), \tag{F.1}$$

where $k(1), \dots, k(J)$ constitutes a partition of firms into K known classes.[F.1]

After imputing a cluster to each worker-year observation, I estimate the regression

$$\tilde{w}_{it} = \alpha_i + \sum_{k=1}^{K} \psi_k \mathbf{1}(j(i,t) = k) + X'_{it}\beta + \epsilon_{it} \tag{F.2}$$

where α_i is the individual fixed effect, $\mathbf{1}(J(i,t) = k)$ are dummies indicating which cluster k firm the employer of i in period t, $j(i,t)$ has been assigned to, and X_{it} is the same vector of time-varying controls as in the preceding section B.1.1.[F.2] The estimation is

[F.1] Here, $\hat{F}_j$ is the empirical cdf of log-wages in firm j; n_j is the average number of workers of firm j over the sample period; and $H_1, \dots, H_K$ are generic cdf's. I use a baseline value of $K = 20$ but have experimented with $K = 10$ and $K = 100$ as well; the choice makes little practical difference. I use firms' wage distributions over the entire sample period on a grid of 20 percentiles for clustering.

[F.2] While residualizing wages for observables aligns with the model's exclusion of life-cycle and on-the-job learning effects, it is not self-evident that worker types should likewise be computed from residualized wages. It could be argued, for instance, that for the interpretation of the production function it matters, for example, whether a worker is good," not whether they are good for their age." I include controls to maintain maximum consistency, both internally and with respect to the existing literature. A previous version (Freund, 2023) reported results when worker types are constructed without controlling for observables;

implemented in Stata using the *reghdfe* package (Correia, 2017).

I estimate this regression separately for each of the 5 sample periods, then equate worker i's type $\hat{x}_i$ with their decile rank in the distribution of α_i. Note that I thus allow for the possibility of human capital accumulation across but not within sample periods. In the data, the auto-correlation structure for period-specific α_i, exhibits a strongly positive though less than unit correlation.

F.2 Labor market transition rates from the LIAB

This section describes how the empirical labor market transition rates that discipline the job arrival and destruction rates in the quantitative model are computed. As a data source, I supplement the SIEED with the Linked Employer Employee Data longitudinal model (LIAB LM7519), which, unlike the SIEED, contains information on non-employment spells.

I proceed in four main steps. First, I convert the spell-level data into a monthly panel. Second, I restrict the sample to approximate the selection criteria used in the other empirical analysis but without being limited to employed persons, i.e., I select individuals aged 20-60 who only ever worked for establishments in West Germany. Third, in the construction of transition rates I largely follow Jarosch (2023). Employment refers to full-time employment subject to Social Security. The job finding rate is computed as the rate at which currently non-employed workers who are receiving unemployment insurance (UI) transition into employment. For the job destruction rate, I compute the frequency with which a worker is employed in one month but not in the month thereafter. Note that here I do not condition on receiving UI after separation, as the model does not distinguish between unemployment and non-employment. I instead define the job finding rate based on unemployment to employment transitions, since the model does assume search effort conditional on non-employment. Finally, for job-to-job transitions I compute the rate at which currently employed workers are employed at another establishment the following month. In step four, I compute averages of these different transition rates across months and for different sample periods.

---

differences to the baseline were minimal.

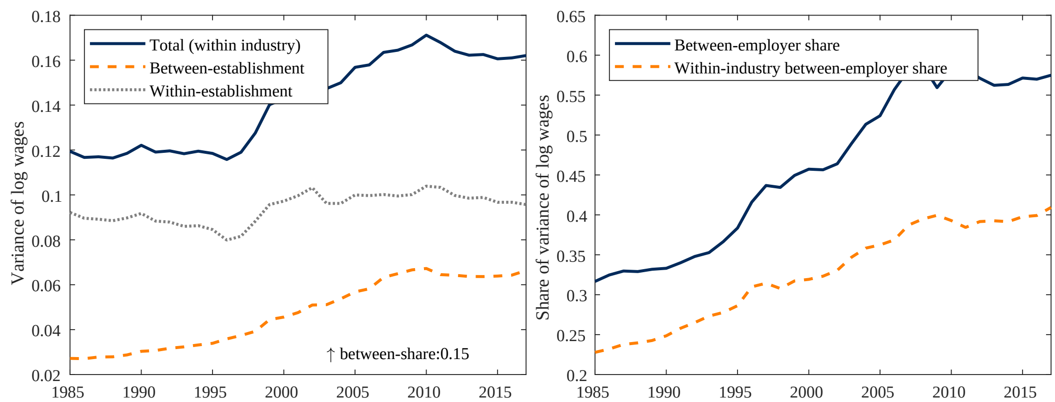

(a) Within-industry decomposition

(b) Between-share: total vs. within-industry

Figure F.1: Industry-level log wage variance decomposition

Notes. All results are based on the baseline measure of residualized log wages and are weighted by person-years.

F.3 Further industry-level results

Log wage variance decomposition. Figure 4a decomposes the yearly total log wage variance into between-establishment and within-establishment components. Figure F.1 supplements this analysis by considering the role of industry differences. I decompose the variance of log wages into three components: between-industry, within-firm within-industry, and between-firm within-industry.

Figure F.1a shows the year-by-year decomposition of the within-industry variance into between- and within components. Figure F.1b depicts the total between-firm share of the variance of log wages (i.e., the fraction of total variance due to both between-industry and between-firm within-industry components), while the dashed line shows the share of the within-industry wage variance due to the between-firm within-industry component. Unsurprisingly, between-firm wage differences play a somewhat smaller role when controlling for industry (dashed line) compared to the economy-wide analysis (solid line). The rise captured by the dashed line clearly indicates, however, that the between-employer share has risen also within industry.

Industry-level correlations. The binscatter plots in Figures 1b and 1c are constructed controlling for period FEs, i.e., relying on cross-industry within-period variation. Figure F.2 reports the relationship between the task-complexity proxy for χ and β_c, as well as between β_c and coworker sorting ρ_{xx}, when including industry FEs instead, i.e., relying on within-industry variation over time. As can be seen, the implied positive relationships

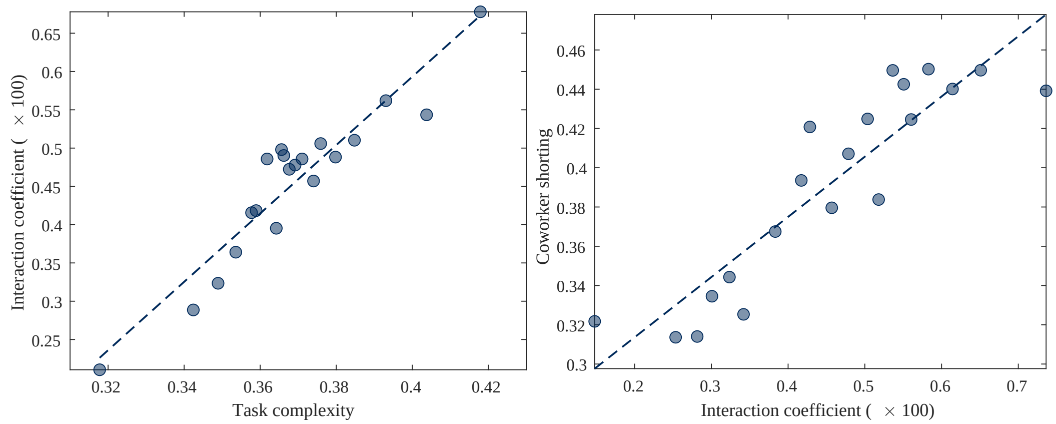

(a) Task complexity vs. β_c

(b) β_c vs. ρ_{xx}

Figure F.2: Industry-level evidence: within-industry variation

Notes. This figure contains binscatter plots controlling for industry fixed effects.

are comparable to those reported in the main text.

Further, Figure F.3 demonstrates that these positive relationships can be observed also when worker types are constructed based on within-occupation rankings.

Industry-level statistics. Lastly, Table F.1 reports industry-level statistics averaged across sample periods.

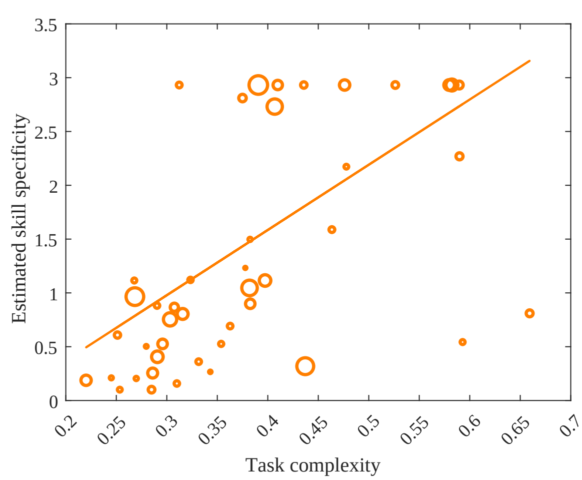

Figure F.4: Task complexity and estimated skill specificity χ

Notes. This figure shows the estimated value of χ for each industry, plotted against task complexity. In a first step, I repeat the online estimation of the parameter vector ψ, targeting average within-industry moments for 2010-2017. In a second step, I estimate industry-specific values of χ, denoted χ_s, by keeping all other parameters fixed and letting only the targeted moment β_c vary.

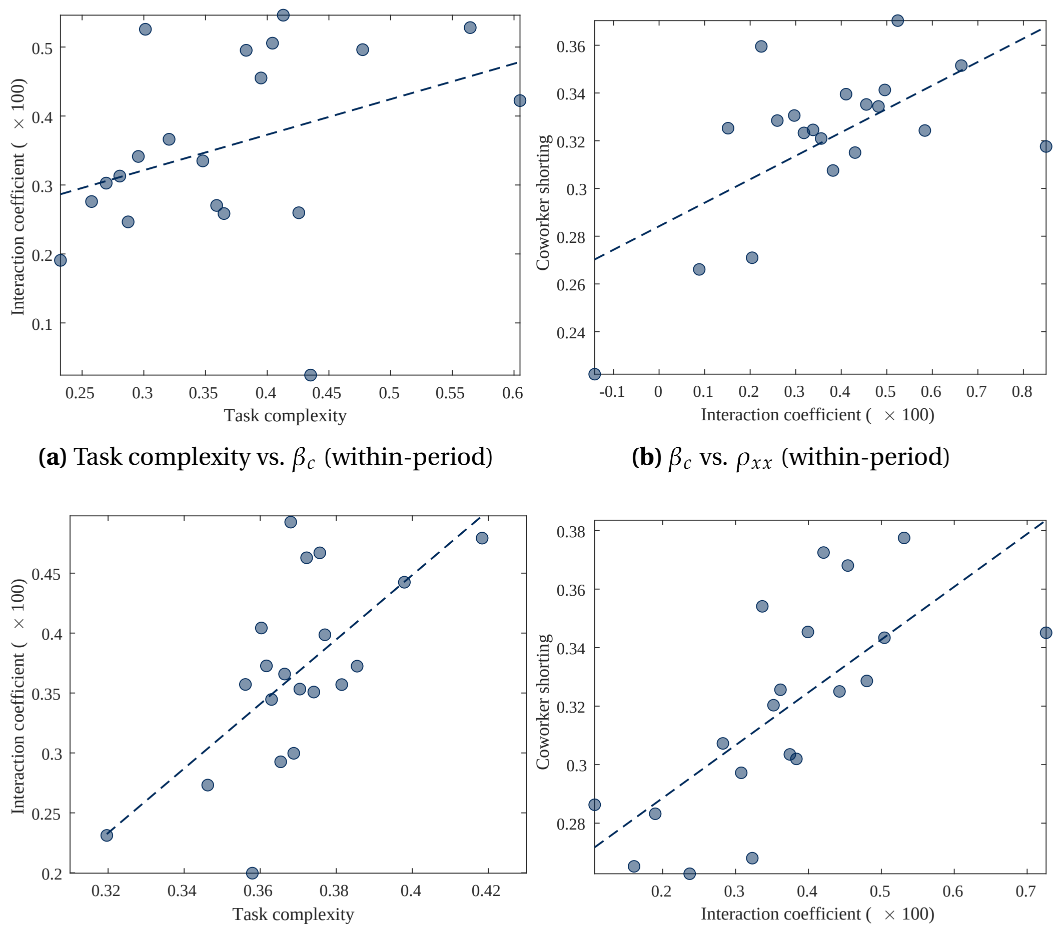

(a) Task complexity vs. β_c (within-period)

(b) β_c vs. ρ_{xx} (within-period)

(c) Task complexity vs. β_c (within-industry)

(d) β_c vs. ρ_{xx} (within-industry)

Figure F.3: Industry-level evidence: within-occupation ranking

Notes. This figure contains binscatter plots where worker talent types are based on within-occupation rankings. The top two panels include sample-period FEs; the bottom two instead include industry FEs.

Industry	Obs.	Wage var.	Between-share	$\rho_{xx'}$	$\beta_c(\times 100)$	Task complexity
Manufacture of food products	495128	0.15	0.40	0.31	0.50	0.30
Manufacture of textiles	22770	0.12	0.34	0.25	0.43	0.22
Printing and reproduction of r...	27501	0.13	0.35	0.36	0.35	0.35
Manufacture of chemicals and c...	187130	0.17	0.40	0.40	0.41	0.40
Manufacture of basic pharmaceu...	35111	0.19	0.30	0.32	0.54	0.46
Manufacture of rubber and plas...	194399	0.12	0.30	0.30	0.17	0.27
Manufacture of other non-metal...	58795	0.14	0.38	0.30	0.60	0.31
Manufacture of basic metals	290696	0.09	0.22	0.28	0.17	0.29
Manufacture of fabricated meta...	489184	0.11	0.31	0.31	0.17	0.28
Manufacture of computer, elect...	499144	0.20	0.36	0.30	0.47	0.43
Manufacture of electrical equi...	212468	0.15	0.31	0.28	0.26	0.34
Manufacture of machinery and e...	1085632	0.12	0.29	0.32	0.27	0.35
Manufacture of other transport...	104808	0.14	0.38	0.30	0.68	0.42
Manufacture of furniture	79884	0.11	0.32	0.35	0.26	0.26
Other manufacturing	32926	0.18	0.44	0.37	0.35	0.36
Repair and installation of mac...	62597	0.16	0.43	0.35	0.51	0.37
Construction of buildings	295761	0.12	0.45	0.35	0.61	0.28
Civil engineering	117780	0.11	0.39	0.32	0.55	0.27
Specialised construction activ...	371579	0.12	0.38	0.33	0.33	0.29
Wholesale and retail trade and...	241568	0.16	0.35	0.35	0.32	0.37
Wholesale trade, except of mot...	990941	0.18	0.45	0.36	0.54	0.41
Retail trade, except of motor ...	676353	0.15	0.37	0.29	0.51	0.39
Land transport and transport v...	88234	0.09	0.55	0.52	0.26	0.21
Warehousing and support activi...	573032	0.11	0.35	0.35	0.33	0.31
Food and beverage service acti...	90756	0.17	0.34	0.21	0.36	0.25
Publishing activities	89965	0.13	0.24	0.21	0.28	0.56
Telecommunications	69529	0.13	0.37	0.34	0.64	0.49
Computer programming, consulta...	284500	0.19	0.38	0.38	0.49	0.59
Legal and accounting activitie...	145238	0.23	0.38	0.39	0.45	0.63
Activities of head offices; ma...	437119	0.20	0.41	0.32	0.60	0.51
Architectural and engineering ...	212346	0.18	0.38	0.41	0.39	0.59
Scientific research and develo...	23883	0.16	0.19	0.20	-0.40	0.53
Rental and leasing activities	10346	0.18	0.47	0.28	0.69	0.38
Employment activities	577282	0.13	0.43	0.24	0.22	0.26
Security and investigation act...	41017	0.10	0.31	0.21	0.21	0.29
Services to buildings and land...	146923	0.21	0.49	0.34	0.32	0.20
Office administrative, office ...	72014	0.30	0.61	0.32	0.54	0.48
Human health activities	520082	0.17	0.15	0.19	-0.21	0.42

Table F.1: Industry statistics

Notes. This table reports lists industry-level statistics averaged across all sample periods. stands for the between-establishment share of the total log wage variance. Only industries for which the β_c estimate is statistically significant at 5% are shown.

F.4 The task content of production in Germany

This section provides further details on the construction, patterns, and trends in task complexity, based on the Employment Surveys (ES) carried out by the German Federal Institute for Vocational Training (Bundesinstitut fuer Berufsbildung, BIBB; Hall and et al. (2018)).[F.3]

The BIBB surveys have several attractive features. They provide detailed information on tasks performed at work. The survey has been fielded, in repeated waves, since 1985 (1985/86, 1991/92, 1998/99, 2006, 2012, and 2018), facilitating time series analyses. Each wave has a large sample size of between 20,000 and over 30,000 respondents per wave, facilitating between-group comparisons. Responses are at the worker-level, and consistent occupation codes can be used across multiple waves, making it possible to capture changes in the nature of work not only associated with employment shifts across occupations but also within-occupation changes (see, e.g., Spitz-Oener (2006); Atalay *et al.* (2020)). Moreover, a supplemental survey in 2012 allows for enriching binary task indicators with information on the actual shares of time spent by employees in different occupations on various tasks.

F.4.1 Methodology

Sample restrictions. As detailed in Rohrbach-Schmidt and Tiemann (2013) and Hall and Rohrbach-Schmidt (2020), time comparisons with the BIBB/IAB surveys require standardization of the sample basis. I follow the steps detailed in those reports and focus on employees from West Germany, aged 15 to 65, who belonged to the labor force (defined as having a paid employment situation) with a regular working time of at least ten hours per week.[F.4] The final sample comprises 91,152 worker-year observations.

I omit the 1998/99 wave from my analysis because the number of activities queried in that wave is substantially lower than in the other surveys. While this reduces the overall sample size, it avoids bias in the results due to limited comparability in tasks. For example, none of the activities "accommodating," "caring," "storing," "protecting," "programming," and "cleaning" were queried in the 1998/99 survey.[F.5]

[F.3] Access was provided by the Research Data Center of the BIBB through scientific use files.

[F.4] In addition, I drop observations for workers who report having performed none of the activities queried in at least two waves. Given the extensive use of occupational codes, I also drop any occupations with fewer than thirty observations across all waves.

[F.5] I thank Daniela Rohrbach-Schmidt for her generous advice on how to handle the older waves and for sharing useful programs.

Task classification. As Rohrbach-Schmidt and Tiemann (2013) make clear, comparisons of task intensities using the BIBB ES over time need to be implemented carefully and must account for variation over time in what tasks are queried and whether their content has changed in meaning. In the context of typical studies that compare task items in the categories non-routine analytical, non-routine interactive, non-routine manual, routine-cognitive, and routine-manual, the authors highlight that routine-cognitive tasks are particularly difficult to classify (e.g., "measuring" may be routine-cognitive or routine-manual; see also Antonczyk *et al.* (2009) compared to Spitz-Oener (2006)).[F.6] Given my focus on complex tasks, these classification problems are less severe. As Rohrbach-Schmidt and Tiemann (2013) note, "these items are regularly observable throughout the cross-sections, their content did not change significantly from year to year, and measurement validity is comparatively strong."

Table B.1 summarizes which tasks are classified as "complex tasks," following Spitz-Oener (2006) and Rohrbach-Schmidt and Tiemann (2013), and compare those with all other tasks.[F.7]

Task index. I define an index capturing the usage of abstract/complex tasks for worker i in period p, following Antonczyk *et al.* (2009):

$$T_{ip}^{\text{abstract}} = \frac{\text{number of activities performed by } i \text{ in task category "abstract" in sample period } p}{\text{total number of activities performed by } i \text{ in sample period } p}$$

To illustrate, if worker i performs five distinct activities in p and two of those belong to the category of abstract/complex tasks, then the complexity index for her work is 0.4.

Occupational classification. To ensure a consistent classification of occupations when using information from multiple waves, I use the German Classification of Occupations 1988 (KldB88). As the oldest classification available in the two most recent waves (2012 and 2018) is the KldB92 classification (Hall and Rohrbach-Schmidt, 2020, cf. Table 9), in processing these two waves I rely on a conversion table KldB92→KldB92; the conversion quality is high as the two classifications are very similar.[F.8]

[F.6] Autor and Handel (2013) also treat the "physical" dimension of tasks as a combined measure of physical and routine tasks. Meanwhile, Acemoglu and Autor (2011) subsume non-routine analytical and non-routine interactive into "abstract," while routine-cognitive and routine-manual tasks are subsumed into "routine."

[F.7] I do not use the task items "managing," "applying law," and "negotiating," because they are only measured in the early waves. Moreover, I associate buying/selling with "other," since even though these tasks may be hard to automate, they are arguably not among the most complex activities. This decision makes no practical difference to the results.

[F.8] This crosswalk is based on the Klassifikationsserver der Statistischen Ämter des Bundes und der Länder, current occupations coded in the 2006 wave in which both KldB88 and KldB92 are available as well as

	Total	Between	Within	Within-share
1986 level	0.252			
1986-1992	0.025	0.002	0.022	0.906
1992-2006	0.298	0.057	0.241	0.809
2006-2012	0.019	0.002	0.017	0.890
2012-2018	0.053	0.028	0.025	0.476
Total change	0.395	0.089	0.306	0.775

Table F.2: The evolving task content of production in Germany

Notes. This table reports the within-between occupation decomposition of the change in the share of complex tasks over time. The "Total" column aggregates across all individuals. The decomposition is performed at the level of KldB-1988 2-digit occupations.

F.4.2 Results

Time trends. The first column in Table F.2 indicates that the aggregate usage share of complex tasks in workers' activities has monotonically increased from 1986 to 2018, with the increase being particularly pronounced in the first half of the time period.

The second to fourth columns decompose the period-by-period change in the importance of complex tasks into two components: a "between" component that captures shifts in occupational employment shares and a "within" component that measures changes in the task content within occupations. Formally, as in Atalay *et al.* (2020), I decompose changes in the usage of abstract tasks between periods t and $t-1$ according to the equation

$$\Delta \bar{T}_t^{\text{complex}} = \sum_o \omega_{o,t-1} (\bar{T}_{t,o}^{\text{abstract}} - \bar{T}_{t-1,o}^{\text{complex}}) + \sum_o (\omega_{o,t} - \omega_{o,t-1}) \bar{T}_{t,o}^{\text{abstract}}$$

where $\bar{T}_{t,o}^{\text{complex}}$ measures the average usage of complex tasks by members of occupation o in period t and $\omega_{o,t}$ is the period- t employment share of occupation o.

Consistent with the findings of **?** for the US, this decomposition reveals that about three quarters of the increase in complex tasks over the sample period has occurred *within* occupations.

Education offers an alternative lens through which to view the changing task content. Figure 3ashowed that the share of complex tasks in the portfolio of university-educated individuals is substantially greater than that of persons with less formal education. The *increase* over time takes place across the board, however.

personal judgements. I thank Anett Friedrich for creating and sharing the crosswalk.

Occupation	$\bar{T}_o^{\text{complex}}$	MT-NRA
Business and administration professionals	0.84	0.47
Legal, social and cultural professionals	0.83	0.67
Business and administration associate professionals	0.82	0.29
Teaching professionals	0.81	0.57
Administrative and commercial managers	0.81	0.58
...	...	...
Drivers and mobile plant operators	0.2	0
Agricultural, forestry and fishery labourers	0.14	0
Food preparation assistants	0.14	0
Market-oriented skilled forestry, fishery and hunting workers	0.12	0
Cleaners and helpers	0.12	0

Table F.3: Top- and bottom-5 occupations in terms of task complexity

Notes. This table reports the top-5 and bottom-5 ISCO-08 2-digit occupations when ranked by $\bar{T}_o^{\text{complex}}$ in pooled 2012 and 2018 waves. The column "MT-NRA" shows the non-routine abstract score taken from Mihaylov and Tijdens (2019) after collapsing to the 2d level using occupational employment shares.

Cross-sectional patterns. The share of complex tasks also varies substantially by occupation. I compute the average complex-task shares at the ISCO-08 2-digit level in the waves 2012 and 2018. Table F.3 lists the bottom 5 and top 5 occupations. I also show the non-routine abstract score from Mihaylov and Tijdens (2019). The comparison reveals large and consistent variation in task shares across occupations according to either measure.

Time usage. One concern is that the analysis thus far considered whether a given task represents an important activity in the respondent's job, rather than measuring how important that activity is relative to others. To address this concern, I draw on a supplemental survey from 2012 that details the amount of time a subset of workers spent o different tasks on a given day. Figure F.5 charts the shares of time spent on the seven complex tasks for different occupational groups. Specifically, I rank occupations according to their task index and group them into 4 equally sized groups. Drawing on the supplemental survey, I then compute the average share of time members of these occupational groups spent on the various tasks.

Figure F.5 shows that each occupational group spends some time on such tasks as "organizing" or "consulting." Crucially, though, the fraction of time spent on each of these

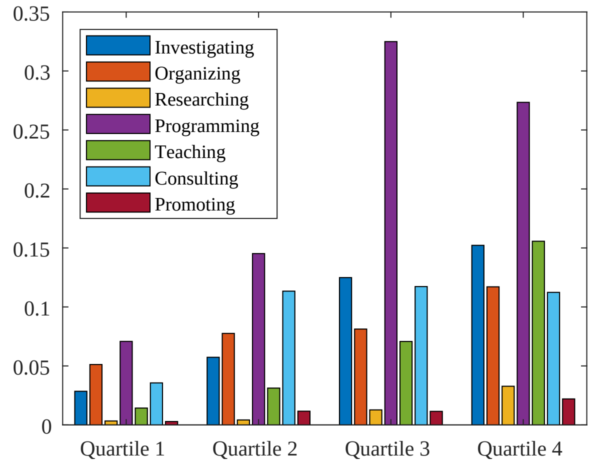

Figure F.5: Allocation of time to complex tasks by occupational groups

Notes. This table reports the share of time spent on complex tasks by different occupational groups. Occupations are first ranked according to their task complexity index and grouped into four groups of approximately equal size. Then the average share of time members of these occupational groups spend on the various tasks labelled as "complex" in Table B.1 is computed.

tasks is several times greater for the top quartile than for the bottom quartile No single task drives this result.

F.5 Cross-country evidence on the "firming up" of inequality

Figure F.6 illustrates cross-country trends for the between-firm share of wage inequality, drawing on aggregated statistics kindly made available by Tomaskovic-Devey *et al.* (2020). While the levels cannot be straightforwardly compared across countries due to variation in the measure of earnings used (e.g., hourly vs. daily vs. monthly earnings), one can observe a consistent upward trend for almost all countries.

Interestingly, even in a country like France, where total wage inequality has broadly flatlined over the past few decades, the between-firm component tends to have increased due to rising sorting and segregation, whereas within-firm inequality has declined for a variety of reasons (Babet *et al.*, 2022).

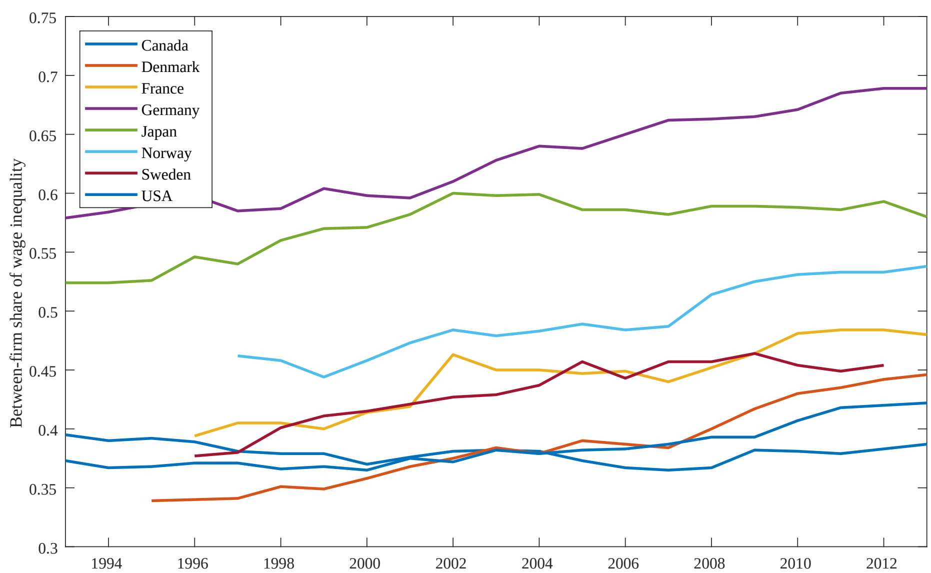

Figure F.6: Cross-country evidence on the between-firm share of wage inequality

Notes. This figure reports the evolution of the between-firm share of wage inequality for a set of OECD economies. Data source: Tomaskovic-Devey *et al.* (2020).

G Quantitative appendix II

G.1 The role of endogenous complementarities in shaping productivity

Figure 7 shows how productivity varies with χ for different labor market allocations. The amended production function used in the quantitative analysis, equation (21), provides for an alternative way to compare outcomes with endogenous talent complementarities to a counterfactual production function keeping them fixed. Suppose we generalize this production function to read

$$\hat{f}(\hat{x}, \hat{x}', \xi) = 2 \times \left(\frac{\bar{n}}{\bar{n}-1}\right)^{\chi_1 \xi} \times \left(a_0 + a_1 \left(\frac{1}{2}\left(\hat{x}\right)^{\frac{1}{\chi_2+1}} + \frac{1}{2}\left(\hat{x}'\right)^{\frac{1}{\chi_2+1}}\right)^{\chi_2+1}\right).$$

Now suppose that, starting from $\chi_1 = \chi_2 = \chi_0$, we let χ_1 increase. The baseline with endogenous talent complementarities corresponds to $\chi_2 = \chi_1$; the fixed-complementarities counterfactual keeps $\chi_2 = \chi_0$.

Figure G.1 illustrates the productivity implications. Throughout, to isolate the effects of talent complementarities in the production function, we set $\xi = \frac{1}{2}$ and instead of re-solving the model for different values of χ, the joint distribution of talent types corresponds to either (i) the equilibrium distribution in the baseline economy (2010-2010), (ii) PAM, or (iii) random matching. As a reference point, and by construction, under PAM – as in Kremer (1993), for instance – talent complementarities have no effect on output. The solid orange line and diamond markers show that *if* complementarities are held fixed, productivity moves with χ "as if" matching was perfectly assortative, even if labor markets exhibit mismatch. This specification misses that greater specificity amplifies the cost to output of coworker mismatch. By contrast, when complementarities are allowed to vary endogenously with χ, productivity increases by less (blue vs. orange lines); and this differential is larger the farther the labor market allocation is from the PAM benchmark (blue solid vs. blue dotted line).

G.2 Validation of identification approach

To validate the identification of jointly inferred parameters, ψ, I conduct three exercises. First, to support the argument that each element of the parameter vector ψ is closely linked to a particular moment, Figure G.2 plots the relevant moment against the respective parameter. As required for local identification, the relationships are monotonic and exhibit substantial variation.

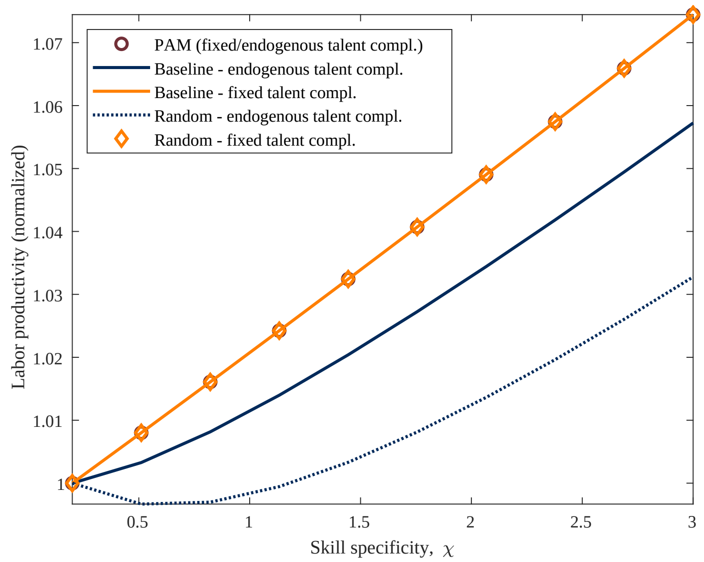

Figure G.1: Productivity gains from rising specialization with fixed vs. endogenous talent complementarities

For the second exercise, let an individual parameter ψ_i vary around the estimated value ψ_i^* and plot the distance function $\mathcal{G}(\psi_i, \psi_{-i}^*)$, where the remaining parameters are allowed to adjust to minimize $\mathcal{G}$. Figure G.3 indicates that $\mathcal{G}$ has a steep U-shape for $\psi_i \in \{a_0, a_1, b_1, \lambda_u\}$. The two plots for χ – the first of which considers the 2010-2017 targets while the second examines an earlier period in which the value of χ is lower – highlight a nuance worth noting: For very high values of χ, inference becomes less reliable, with the second half of the U not "closing."

To further examine this issue, I perform a third, Monte-Carlo style exercise that works as follows: let χ vary over the interval $[\hat{\chi}^{\text{min. across periods}}, \hat{\chi}^{\text{max. across periods}}]$; for each of ten equally spaced values in this interval, simulate data from the model holding all parameters fixed at their average value across sample periods; and apply the indirect inference method to recover the vector of estimated parameters ψ. Figure G.4 shows that the inferred parameter values align well with those used to generate the data used in the estimation, though for high values of χ the estimate lies below the true value. The reason is that as χ increases, variation in ξ – and, hence, systematic heterogeneity in selection on ξ, where the *average* value of ξ has to be larger for teams with bigger talent gaps (cf. Borovickova and Shimer, 2024) – exerts a stronger effect on output. This means the direction of the bias in the estimation of χ is downward, so our estimates represent a conservative lower bound.

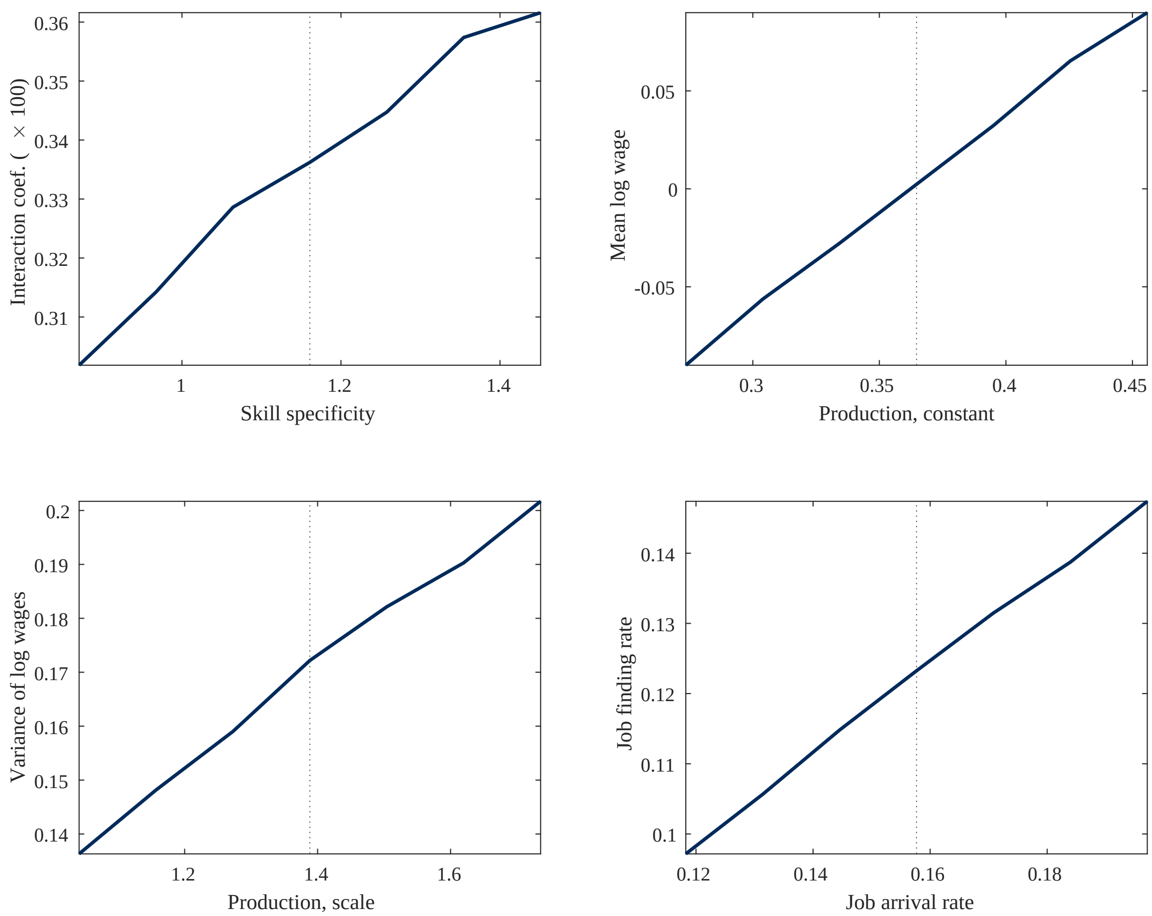

Figure G.2: Validation of identification: moment against parameter

Notes. This figure plots the targeted moment against the relevant parameter, holding constant all other parameters at their average value across sample periods.

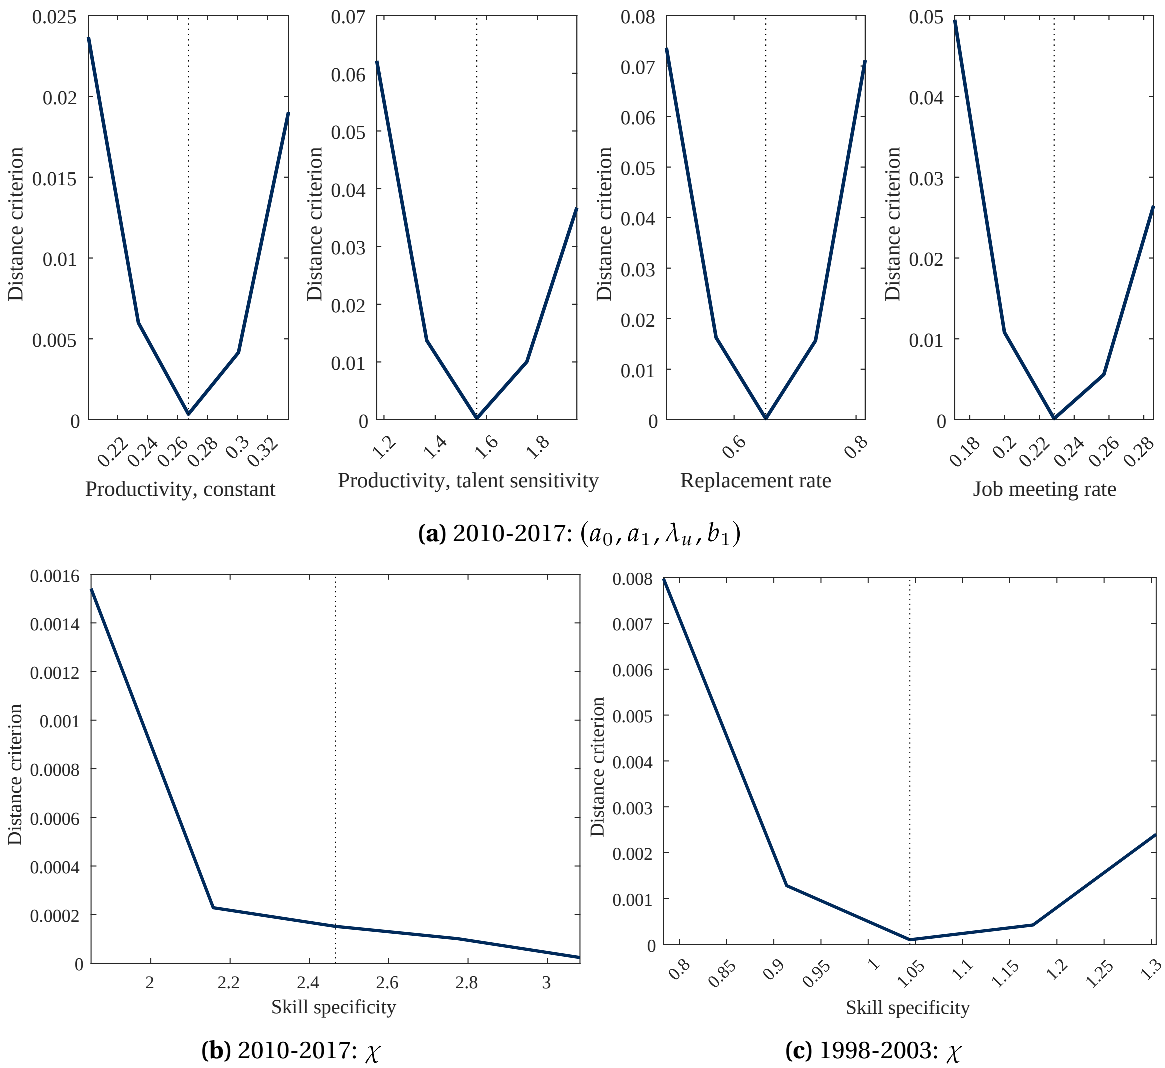

(a) 2010-2017: $(a_0, a_1, \lambda_u, b_1)$

(b) 2010-2017: χ

(c) 1998-2003: χ

Figure G.3: Validation of identification method: distance criterion

Notes. This figure plots the distance function $\mathcal{G}(\psi_i, \psi^*_{-i})$ when varying a given parameter ψ_i around the estimated value ψ^*_i. The remaining parameters are allowed to adjust to minimize $\mathcal{G}$.

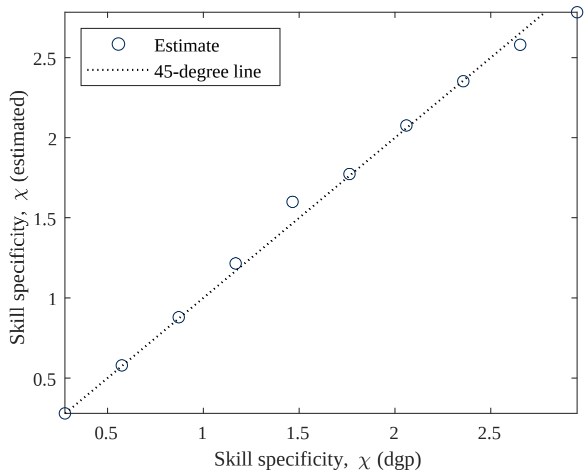

Figure G.4: Validation of identification method: Monte-Carlo experiment

Notes. This figure shows the value of χ used to generate simulated data on the horizontal axis and the corresponding estimated value on the vertical axis.

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. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Antonczyk, D., Fitzenberger, B., and Leuschner, U. (2009). Can a Task-Based Approach Explain the Recent Changes in the German Wage Structure? *Jahrbücher für Nationalökonomie und Statistik*, **229**(2-3), 214–238.

Atalay, E., Phongthiengtham, P., Sotelo, S., and Tannenbaum, D. (2020). The Evolution of Work in the United States. *American Economic Journal: Applied Economics*, **12**(2), 1–34.

Autor, D. H. and Handel, M. J. (2013). Putting Tasks to the Test: Human Capital, Job Tasks, and Wages. *Journal of Labor Economics*, **31**(S1), S59–S96.

Babet, D., Godechot, O., and Palladino, M. G. (2022). In the Land of AKM: Explaining the Dynamics of Wage Inequality in France. *mimeo*.

Bonhomme, S., Lamadon, T., and Manresa, E. (2019). A Distributional Framework for Matched Employer Employee Data. *Econometrica*, **87**(3), 699–739.

Borovickova, K. and Shimer, R. (2024). Assortative Matching and Wages: The Role of Selection. *mimeo*.

Brown, B. M. and Resnick, S. I. (1977). Extreme values of independent stochastic processes. *Journal of Applied Probability*, **14**(4), 732–739.

Correia, S. (2017). Linear models with high-dimensional fixed effects: An efficient and feasible estimator. *Unpublished manuscript, http://scorreia. com/research/hdfe. pdf*.

Davison, A. C., Huser, R., and Thibaud, E. (2019). Spatial extremes. In *Handbook of Environmental and Ecological Statistics*, pages 711–744. Chapman and Hall/CRC.

De Haan, L. (1984). A Spectral Representation for Max-stable Processes. *The Annals of Probability*, **12**(4), 1194–1204.

Freund, L. B. (2023). Superstar Teams: The Micro Origins and Macro Implications of Coworker Complementarities. *Available at SSRN 4312245*.

Hall, A. and et al. (2018). BIBB/IAB / BIBB/BAuA Employment Survey 2018 (2012, 2006, 1998/99, 1991/92, 1985/86). Technical Report doi:10.7803/501.18.1.1.10, doi:10.7803/501.12.1.1.60, doi:10.7803/501.06.1.1.20, doi:10.4232/1.12247, doi:10.4232/1.2565, 10.4232/1.12563, Research Data Centre at BIBB.

Hall, A. and Rohrbach-Schmidt (2020). BIBB/BAuA Employment Survey 2018. BIBB-FDZ Data and Methodological Report 1/2020. *BIBB Report*.

Huser, R. and Davison, A. C. (2013). Composite likelihood estimation for the Brown–Resnick process. *Biometrika*, **100**(2), 511–518.

Jarosch, G. (2023). Searching for Job Security and the Consequences of Job Loss. *Econometrica*, **91**(3), 903–942.

Kabluchko, Z., Schlather, M., and de Haan, L. (2009). Stationary max-stable fields associated to negative definite functions. *The Annals of Probability*, **37**(5), 2042–2065.

Kremer, M. (1993). The O-Ring Theory of Economic Development. *The Quarterly Journal of Economics*, **108**(3), 551–575.

Mihaylov, E. and Tijdens, K. G. (2019). Measuring the Routine and Non-Routine Task Content of 427 Four-Digit ISCO-08 Occupations.

Resnick, S. I. (1987). *Extreme Values, Regular Variation and Point Processes*. Springer Series in Operations Research and Financial Engineering. Springer, New York, NY.

Rohrbach-Schmidt, D. and Tiemann, M. (2013). Changes in workplace tasks in Germany—evaluating skill and task measures. *Journal for Labour Market Research*, **46**(3), 215–237.

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.

Tomaskovic-Devey, D., Rainey, A., Avent-Holt, D., Bandelj, N., Boza, I., Cort, D., Godechot, O., Hajdu, G., Hällsten, M., Henriksen, L. F., Hermansen, A. S., Hou, F., Jung, J., Kanjuo-Mrčela, A., King, J., Kodama, N., Kristal, T., Křížková, A., Lippényi, Z., Melzer, S. M., Mun, E., Penner, A., Petersen, T., Poje, A., Safi, M., Thaning, M., and Tufail, Z. (2020).

Rising between-workplace inequalities in high-income countries. *Proceedings of the National Academy of Sciences*, **117**(17), 9277–9283.