

Essays on Technical Change and Labor Markets

A Dissertation Submitted in Partial Fulfillment

Of the Requirements for the Degree of

Doctor of Philosophy in Economics

Tepper School of Business

Carnegie Mellon University

Pittsburgh, PA

Andre Mouton

April 2022

Dissertation Committee:

Ali Shourideh (chair)

Laurence Ales

Brian Kovak

Rebecca Lessem

Copyright © Andre Mouton

All rights reserved.

Essays on Technical Change and Labor Markets

Abstract

This dissertation consists of three chapters exploring the impact of technical change on labor markets. In the first chapter I conduct a forensic analysis of wage sorting: an observed tendency for high-earning workers to match with high-paying employers, which over time has contributed to rising wage inequality. I evaluate competing causal mechanisms using matched employer-employee data from Germany and find that (1) wage sorting is entirely a between-industry/-occupation phenomenon, and inconsistent with models of assortative matching within markets; (2) increased wage sorting over time can be fully accounted for by declining employment in low-skill, high-paying manufacturing sectors, and rising wages in skill-intensive occupations more common in high-paying firms; and (3) wage sorting is poorly predicted by measures of anti-competitive rents in output and labor markets, but strongly associated with job and workplace characteristics related to technology, which proxy for two well-known dimensions of wage variation: worker skill and employer scale.

In chapter 2, I study the quantitative implications of firm-wage premia for theories of skill-biased technical change. I develop a search- and assignment-based model of labor markets that accounts for equilibrium interactions between demand and supply, skill premia, and firm premia, while remaining sufficiently tractable that the key distributional parameters can be non-parametrically identified from empirical wage effects. I structurally estimate the model using matched data from West Germany, and find that more than half of the rise in wage variance associated with industry and occupation demand is the result of interactions with firm premia. I show that the “firm-bias” of demand confounds the relationship between skill-biased shocks and wages, but can provide insight into regional variation in wage trends. I show in addition that because skill and firm premia are highly correlated across labor markets, policies that seek to reduce wage dispersion by targeting firm premia are generally skill-biased, partially offsetting their aggregate impact on the wage distribution.

In the final chapter I study the effects of task-level automation when jobs consist of multiple tasks. I consider an environment with endogenous assignment of workers to occupations, and of worker time to different tasks within jobs, and I characterize the aggregate effects of a technology that replaces labor at a low-skill task. The model predicts a reverse pattern of automation: the low-skill task is first automated in high-skill occupations, where labor costs are higher. In the short-run this creates wage and employment polarization. In the long-run, automation has ambiguous implications for wage inequality and employment. I use panel survey data on occupational tasks and computerization to test the model's short-run predictions, and I estimate a structural version of the model in order to obtain long-run labor market predictions. Further declines in IT-related costs are predicted to have little effect on wages, but to increase employment in low- and middle-skill occupations.

Acknowledgements

I am deeply grateful for the time, effort, and guidance provided by the members of my dissertation committee: Ali Shourideh, Laurence Ales, Brian Kovak, and Rebecca Lessem. Many members of the Tepper community contributed to this work by way of discussions (and occasionally debates), and I am indebted to all of the workshop participants who provided continual feedback over the years. I owe a special thanks to Xuege Zhang and Santiago Tellez, for their spiritual and intellectual support. In this dissertation I draw heavily on data that was generously provided by the Institute for Employment Research, whom I cannot recommend too highly. And above all I am thankful for my family, whose support and patience were necessary and, as it turns out, sufficient.

Contents

List of Tables	ix
List of Figures	xi
Introduction	1
1. What Drives West German Wage Sorting?	3
1.1 Introduction	3
1.1.1 Labor Sorting and Wage Sorting	8
1.2 West German Wage Sorting, 1993-2017	11
1.2.1 Data and Decomposition	11
1.2.2 West German Wage Sorting	14
1.2.3 Limited-Mobility Bias	17
1.3 Is Wage Sorting Within or Between Markets?	20
1.3.1 Between-Market Wage Sorting	20
1.3.2 Wage Sorting and Match Selection	24
1.3.3 Decomposing the Between-Group Trend	26
1.3.4 Discussion	30
1.4 Descriptive Evidence: Skill, Scale, and Technology	30
1.4.1 Wage Sorting and Establishment Characteristics	31
1.4.2 Wage Sorting and Job Characteristics	36
1.5 Conclusion	40
2. Skill-Bias, Firm-Bias, and Wage Inequality	42
2.1 Introduction	42
2.1.1 Motivation	47
2.2 A Model of Firm-Bias and Skill-Bias	49
2.2.1 Environment	50
2.2.2 Equilibrium	53
2.2.3 The Restricted Model	54
2.2.4 Demand, Supply, and Wages	57

2.2.5	Discussion	60
2.3	Structural Estimation	62
2.3.1	Data	62
2.3.2	Type Distributions	64
2.3.3	Match Productivity	65
2.3.4	Demand, Amenities, and Costs	66
2.3.5	Aggregate Parameters	67
2.4	Quantitative Results	69
2.4.1	Historical Firm-Bias	69
2.4.2	Firm-Bias and the Effects of a Skill-Biased Shock	73
2.4.3	Firm-Bias and Regional Wage Trends	75
2.4.4	Firm Premia and Wage Policies	78
2.5	Conclusion	82
3.	Task Automation and Labor Polarization	84
3.1	Introduction	84
3.2	Descriptive Analysis	87
3.2.1	Occupational Routineness in the BIBB Surveys	88
3.2.2	Task Variety Within Jobs	90
3.2.3	Computerization and Task Content	92
3.3	A Model of Partial Automation	95
3.3.1	Environment	95
3.3.2	Equilibrium	98
3.3.3	Automation, Wages, and Employment	101
3.3.4	Continuous Model	105
3.4	Qualitative Evaluation	106
3.4.1	Time Path of Automation	107
3.4.2	Employment Effects of Automation	109
3.5	Quantitative Analysis	111
3.5.1	Estimation	111

3.5.2	Employment, Wages, and Technology Adoption	113
3.5.3	Capital-Labor Substitution	115
3.5.4	Supply and Demand Elasticities	117
3.6	Conclusion	119
References		121
A. Supplementary Materials For Chapter 1		132
A.1	Data Preparation	132
A.2	Variance Decompositions	138
A.3	Additional Between-Group Results	140
A.4	Additional Results, Person/Establishment Characteristics	147
B. Supplementary Materials For Chapter 2		153
B.1	Proofs of Main Theoretical Results	153
B.2	Equilibrium Characterization	159
B.3	Structural Estimation	161
B.4	Quantitative Experiments	165
B.5	Additional Results	168
B.6	Alternative Specifications	176
C. Supplementary Materials For Chapter 3		177
C.1	Empirical Appendix	177
C.2	Assignment and Wages	180
C.3	Proofs of Theoretical Results	185
C.4	Continuous Model	187
C.5	Additional Qualitative and Quantitative Results	188

List of Tables

1.1	AKM Variance Decomposition, 1993-2017	14
1.2	Wage Sorting By Group	16
1.3	AKM Identification Statistics	18
1.4	Between-Group Wage Sorting (% Total)	21
1.5	Counterfactual Between-Group Wage Sorting (% Trend)	27
1.6	Wage Effects and Establishment Characteristics, 2003-2010	33
1.7	Wage Effects and Job Characteristics, 2003-2010	37
2.1	Aggregate Parameters	68
3.1	Summary Statistics for BIBB Employment Surveys, 1979-2018	88
3.2	Routine Risk Content in the BIBB Surveys	89
3.3	Task Characteristics and PC Use, 1979-2018	93
3.4	Occupation Mean Task Characteristics and PC Use (Two-Way FE)	94
3.5	Regression: Change in Log Occupational Employment Share	110
3.6	Regression: Change in Log Occupational Employment Share, By Year	110
3.7	Aggregate Parameter Estimates	113
A.1	LIAB Summary Statistics	133
A.2	Aggregate Industry Classifications	135
A.3	Aggregate Occupation Classifications	136
A.4	AKM Identification Statistics	137
A.5	Wage Sorting By Demographic (Correlation)	139
A.6	AKM Variance Decomposition, East Germany 1996-2017	140
A.7	AKM Variance Decomposition, Germany 1996-2017	140
A.8	Between-Group Wage Moments (% Total)	141
A.9	Counterfactual Between-Group Wage Sorting (Correlation)	143
A.10	Counterfactual Between-Group Wage Sorting (% Trend)	143
A.11	Between-Group Variances	144
A.12	Cross-Correlations, Establishment Characteristics	148
A.13	Wage Effects and Establishment Characteristics, Confidence Intervals	149

A.14 BIBB Task Descriptions	149
A.15 Cross-Correlations, Person/Job Characteristics	150
A.16 Wage Effects and Person Characteristics, Confidence Intervals	151
A.17 Wage Effects and Job Characteristics, Additional Tasks	151
A.18 Wage Effects and Job Characteristics, Non-Imputed	152
B.1 Empirical and Predicted Wage Moments	164
B.2 Predicted and Counterfactual Wage Moments	168
B.3 Predicted and Counterfactual Wage Moments, By Parameter	170
B.4 Aggregate Parameters, East Germany	172
C.1 Task Characteristics and Task Frequencies, 2006-2018	179
C.2 Task Frequencies and PC Use, 2006-2018	180
C.3 Occupation Mean Task Frequencies and PC Use (D-in-D)	180
C.4 Regression: Change in Log Occupation Employment Share, High-Skill	190
C.5 Regression: Change in Log Occupation Employment Share, Low-Skill	190

List of Figures

1.1	AKM Decomposition of West German Wage Variance, 1993-2017	4
1.2	Labor Market Segmentation	9
1.3	Technical and Match-Based Wage Sorting	10
1.4	West German Wage Sorting, 1993-2017	15
1.5	Wage Sorting By Establishment Size	19
1.6	Between-Group Wage Sorting, Industry and Occupation	22
1.7	Wage Effects by Occupation and Industry	24
1.8	Wage Sorting and Match Selection, 2010-2017	25
1.9	Evolution of Mean Person Effects and Shares, 1993-2017	29
1.10	Establishment Scale, ICT Investment, and Wage Effects	34
1.11	Worker Skill, PC Use, and Wage Effects	39
2.1	Wage Effects by Occupation and Industry, 2010-2017	48
2.2	Simulated Effects of a Skill-Biased Demand Shock	59
2.3	Estimated Type Distributions	64
2.4	Estimated Match Production Parameters	66
2.5	Estimated Technical Parameters	67
2.6	Predicted and Counterfactual Wage Trend, 1993-2017	71
2.7	Contribution to Change in Wage Variance, 1993-2017	72
2.8	Effect of a Skill-Biased Demand Shock	74
2.9	Predicted and Counterfactual Wage Variance, East Germany	76
2.10	East German Wage Trend, Alternative Entry Costs	77
2.11	Simulated Effect of Equal Pay for Temp Workers	80
2.12	Simulated Effect of Entry Cost Compression	81
3.1	Routine and Non-Routine Activity by Wage Percentile, 2006	91
3.2	Task Variety by Wage Percentile, 2006	92
3.3	Short-Run Effects of Automation	102
3.4	Long-Run Effects of Automation, $\rho = 0$	103
3.5	Long-Run Effects of Automation, $\rho = 1$	104

3.6	Short-Run Effects of Automation, Continuous Model	106
3.7	Computer Use by Occupation Wage Percentile	108
3.8	Technology Penetration by Wage Percentile	108
3.9	Model Goodness-of-Fit	112
3.10	Technology Adoption and Labor Outcome	114
3.11	Short-Run and Long-Run Distributional Effects	115
3.12	Cross-Price Elasticity of Labor Demand: Rental Rate	116
3.13	Varying the Elasticity of Labor Demand	117
3.14	Varying the Elasticity of Labor Supply	119
A.1	West German Wage Variance, 1993-2017	138
A.2	West German Wage Sorting (Correlation), 1993-2017	138
A.3	East German Wage Variance, 1993-2017	139
A.4	Annual Between-Group Wage Sorting, 1993-2017	141
A.5	Wage Effects by Detailed Occupation and Industry	142
A.6	Wage Sorting and Match Selection (Correlations), 2010-2017	142
A.7	Establishment Effects by Industry and Panel	144
A.8	Changes to Industry-Occupation Mean Person Effects and Shares	145
A.9	Establishment Effects by Industry and Size	145
A.10	Establishment Effects by Industry and Bargaining Agreement	146
B.1	Empirical and Predicted Wage Distribution	165
B.2	Predicted and Counterfactual Wage Trend, 1993-2017	169
B.3	Group Contribution to Change in Wage Variance, 1993-2017	171
B.4	Estimated Type Distributions, East Germany	171
B.5	Estimated Match Production Parameters, East Germany	172
B.6	Estimated Technical Parameters, East Germany	172
B.7	Industry Average Firm Premia, 1993-1999	173
B.8	Temp Agency Employment, West Germany	173
B.9	Simulated Effect of Equal Pay for Temp Workers	174
B.10	Simulated Effect of Entry Cost Compression	175
B.11	Counterfactual Wage Trend, Constant Skill Distribution	176

B.12 Counterfactual Wage Trend, Non-Unit Elasticity of Demand	176
C.1 Computer Use By Occupation Wage Percentile, 1979-1999	189
C.2 Computer Use by Occupation Wage Percentile, 1999-2018	189
C.3 Distributional Parameter Estimates	191
C.4 Automation of Skilled Task	191

Introduction

This dissertation consists of three chapters exploring the impact of technical change on labor markets. The first chapter is a forensic analysis of sorting patterns in labor markets, and how these patterns have contributed to the post-1980's rise in wage inequality. Past studies find that much of observed wage dispersion is due to differences in pay across employers, and that over time there is an increasingly strong association between high-paying firms and high-earning workers. I utilize matched German employer-employee data to characterize wage sorting over the 1993-2017 period, and to evaluate alternative causal mechanisms. Group decompositions indicate that wage sorting is entirely a between-industry/-occupation phenomenon, and I find little evidence to support assortative matching within narrowly-defined labor markets. The greater contribution from wage sorting over time is not driven by changes to the wages firms pay, but is instead accounted for by a shift in employment from low-skill, high-paying manufacturing firms to low-skill, low-paying service establishments; and by rising wages in skill-intensive jobs, which in all periods are more common in high-paying firms. To uncover the dimensions of worker and firm heterogeneity most closely associated with wage sorting, I turn to establishment and occupational survey data, which I use to conduct a partial correlation analysis. I find that wage sorting is poorly predicted by measures of anti-competitive rents in output and labor markets, but strongly associated with technology-related job characteristics, which are highly effective at capturing two well-known dimensions of wage variation: worker skill and employer scale.

In chapter 2, I study *firm-bias*: how changes to industry and occupation labor demand interact with the firm-specific component of wages, impacting the distribution of firm premia and their relationship with skill premia. I develop a search- and assignment-based model of labor markets that accounts for equilibrium interactions between demand, supply, skill premia, and firm premia, and that captures the behavioral response of agents to changes in wage premia. The model nests as a reduced form the canonical AKM wage regression,

allowing the key distributional parameters to be non-parametrically identified from empirical wage effects. I structurally estimate the model using matched German employer-employee data, and conduct counterfactual experiments quantifying (1) the contribution of firm-bias to historical wage trends, (2) the implications of firm-bias for the study of skill-bias and of regional wage trends, and (3) the importance of interactions between firm- and skill-bias for policies that target firm rents and rent-sharing. I find that in the absence of firm-bias, historical demand shocks would have increased wage inequality by less than half as much. By itself, the skill-bias of demand is an unreliable predictor of wage outcomes, as an increase in demand for low-skill (high-skill) labor can nevertheless increase (reduce) wage inequality, depending upon the associated industry. In practice this can help to explain regional differences in wage trends, and I show that a smaller increase in East German wage variance is to a large extent the result of weaker firm-bias. Finally I show that, because skill and firm premia are highly correlated, policies that target firm-side wage gaps are generally skill-biased, partially offsetting their aggregate impact on wage inequality.

In the final chapter, I study the short-term and long-run effects of task automation when jobs consist of multiple tasks. Leveraging German survey data, I show that task variety is ubiquitous at the job level, and that computerization over the 1979-2018 period is associated with intra-occupational shifts away from lower-skill and routine task content. I explore the implications of task automation in a model that combines occupational assignment with a time allocation problem where workers must perform multiple tasks. The model predicts a reverse pattern of automation: low-skill tasks are automated first in high-skill occupations, where labor costs are higher. In the short-term this creates wage and employment polarization around the automation “frontier.” In the long-run, once the technology has been fully adopted, further declines in technology costs tend to improve wage and employment outcomes at low-skill jobs. I show that the model’s short-run predictions are consistent with the historical time paths of computerization and occupational employment in West Germany. I then estimate a structural version of the model to obtain long-run predictions for German labor markets. Further declines in computer-related costs are expected to have little effect on wages, but to substantially increase employment in low- and middle-skill occupations.

Chapter 1

What Drives West German Wage Sorting?

1.1 Introduction

Two stylized facts emerge from recent studies of the wage distribution. First, much observed wage dispersion is due, not to differences between workers, but to differences between firms.¹ Second, in many countries high-earning individuals are more likely to work for high-paying employers. This phenomenon, known as *wage sorting*,² represents an increasingly important source of inequality. Wage sorting has contributed substantially to rising dispersion in Germany, the United States, and the Scandinavian countries,³ and is related to the more general finding that higher OECD wage inequality is predominantly due to wage gaps between firms, rather than within them.⁴ Wage sorting nevertheless remains a puzzle. We understand little about its causal origins, the reason for its growing significance, or its relationship with observable characteristics of markets and agents. Such an understanding is important because wage sorting represents a challenge to conventional explanations of wage inequality, that focus on one side of the labor market and abstract from the question of “Who works where?”

In this chapter I characterize West German wage sorting, and evaluate the consistency of German wage trends with various proposed explanations. I draw on a unique dataset that pairs administrative records for 1993-2017 with an annual survey of business estab-

¹Following [Abowd, Kramarz, and Margolis \(AKM, 1999\)](#), many papers have replicated this result; see [Card, Cardoso, Heining, and Kline \(2018\)](#) for a survey.

²After [Bagger, Sorensen, and Vejlin \(2013\)](#).

³See [Card, Heining, and Kline \(2013\)](#), [Song, Price, Guvenen, Bloom, and von Wachter \(2019\)](#), [Bagger, Sorensen, and Vejlin \(2013\)](#), and [Hakanson, Lindqvist, and Vlachos \(2021\)](#).

⁴For example [Dunne, Foster, Haltiwanger, and Troske \(2004\)](#), [Simon \(2010\)](#), [Barth, Bryson, Davis, and Freeman \(2014\)](#), and [Tomaskovic-Devey et al. \(2020\)](#).

lishments, offering an unusual level of detail on both workers and their employers. Wage sorting is measured as the covariance of person and employer wage effects, estimated from a standard AKM wage regression. I use decomposition methods to answer the question: to what extent is wage sorting occurring *within* labor markets versus *between* them? The purpose of this exercise is to differentiate between two main classes of theories: those that study worker-firm sorting in the context of match-level complementarities and unobservable heterogeneity, and those concerned with industry or geographic variation in the technical demand for skill. Focusing on the between-market component of wage sorting, I then conduct a series of reduced-form experiments to uncover how changes to market composition and wage premia have contributed to wage sorting over time, and to assess the causal role of several well-studied macroeconomic trends: rising *technical* demand for skilled labor, *structural* shifts in the distribution of labor across industries, and changes to *market structure* following reunification and the labor reforms that followed. Finally, I utilize establishment- and person-level survey data to conduct a partial correlation analysis that studies the statistical relationship between wage sorting, market characteristics, and observable dimensions of agent heterogeneity.

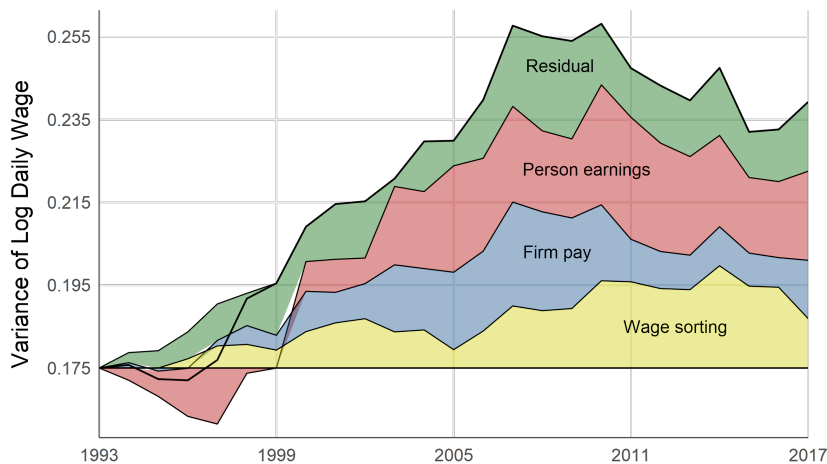


Figure 1.1: AKM Decomposition of West German Wage Variance, 1993-2017

SOURCE: German linked employer-employee dataset (LIAB). NOTE: AKM wage effects estimated following [Card, Heining, and Kline \(2013\)](#) in four panels: 1993-99, 1998-04, 2003-10, and 2010-17. From bottom to top, regions indicate changes to the covariance of the person and firm wage effects, the variances of each effect, and the regression residual. Results averaged in years with overlapping panels.

I find that wage sorting in West Germany is *entirely* a between-market phenomenon, apparently unrelated to geographic region, but fully accounted for by aggregated industry and occupation groups. Conditional on the broad industry of the employer and the occupation associated with the job spell, the covariance of the person and firm wage effects is small and negative over the duration of the sample. If wage sorting is between-market, then the relevant dimensions of person and firm heterogeneity should be observable and not require learning by agents; consistent with this prediction, I find that wage sorting just as pronounced, if not more so, among new job spells, new establishments, and workers entering the labor force. Over the years 1993-2017, inter-industry differences in the employer component of pay have been largely stable, indicating that changes to market structure are unlikely to account for increased wage sorting during this period. Counterfactual experiments indicate that, instead, the increase in dispersion from wage sorting can be almost entirely accounted for by (1) changes to industry employment shares and (2) rising wages (i.e. person effects) in skill-intensive industries and occupations. The first of these channels is associated with a shift in employment away from commodities and crafts manufacturing, and towards personal services, food and accommodation, and temp agencies - a dynamic similar to that studied by [Autor and Dorn \(2013\)](#). While each of these sectors is similar in terms of employing a mostly low-earning workforce, service firms are lower-paying than those in manufacturing, and so the net effect has been to increase wage sorting in the aggregate. The second channel arises because, in all periods, skill-intensive jobs are concentrated in high-paying employers. As wages in these jobs increase, for example due to increased skill demand as in [Acemoglu and Autor \(2011\)](#), so too does the overall covariance of the person and firm effects. This results in a larger impact on wage variance than would be the case were these sorting patterns absent.

At the industry-occupation level, wage effects are strongly predicted by worker skill and employer scale, while wage sorting is associated with technological skill demand. Correlations between mean establishment size and the firm wage effect lie between .8 and .9, depending on the measure; and I find similarly high correlations between the person wage effect and variables capturing the skill-intensity of the job, such as educational attainment and frequency of analytical tasks. Skill and scale are nevertheless poor at predicting wage sorting, as con-

trolling for these variables does little to reduce the statistical association between worker earnings and employer pay. On the other hand, the wage effects correlation falls by one-half when controlling for job tasks related to PC use and R&D, or for establishment-level product and process innovation; and close to zero when controlling for establishment IT investment. The explanatory power of these variables is apparently due to their ability to proxy for both scale *and* skill. I find that measures of competitiveness in output markets are only weakly informative about either wage effect; while collective bargaining agreements, though predictive of employer pay, are orthogonal to worker skill and consequently uninformative with respect to wage sorting. These results suggest that the relevant worker-firm sorting patterns are technical in nature, and not directly motivated by e.g. rent-seeking behavior as in [Lentz \(2010\)](#).

This first main contribution of this chapter is to the empirical literature on wage sorting. In recent years, researchers have studied the mechanical origins of wage trends by performing panel implementations of the regression decomposition proposed by [Abowd, Kramarz, and Margolis \(AKM, 1999\)](#). In this approach, wages are regressed on person and employer fixed effects, and the wage variance is decomposed into a sum of the variances and the covariance of the effects. [Card, Heining, and Kline \(2013\)](#) estimate that one-third of the post-1980's rise in West German wage variance is due to increased covariance of the person and firm effects. [Song et al. \(2019\)](#) find a comparable rise in the covariance for the United States, and similar results have been shown for Denmark ([Bagger, Sorensen, and Vejlin 2013](#)) and Sweden ([Hakanson, Lindqvist, and Vlachos 2021](#), appendix D). At the same time, wage sorting has been found to be positive but declining in several countries where wage inequality is either stable or falling ([Torres et al. 2018](#); [Alvarez, Benguria, Engbom, and Moser 2018](#)). A key challenge for these studies lies in their interpretation: wage effects are endogenous objects that tell us little about the causal mechanism(s) that generated them, and consequently the implications of wage sorting for quantitative and theoretical studies of the wage distribution are unclear. I show that in Germany, wage sorting can be closely tied to observable characteristics of agents, and that this is sufficient to rule out several competing explanations. [Card et al. \(2013\)](#) suggest that a decline in collective bargaining coverage may explain the contribution of firm wage effects to West German wage trends.

Song et al. (2019) review different mechanisms and argue that match complementarities (e.g. assortative matching or peer effects) are consistent with wage sorting in the United States. In this study I present evidence from Germany that is inconsistent with either of these mechanisms.

A second contribution is to the macroeconomic literature on rising wage inequality, which can be divided into three strands. The skill-bias literature focuses on the role of relative demand for heterogeneously skilled workers, which is seen as increasingly important in light of technological changes affecting labor markets and structural shifts including outsourcing and offshoring. Representative work in this area includes [Acemoglu and Autor \(2011\)](#), [Autor and Dorn \(2013\)](#), and [Goos, Manning, and Salomons \(2014\)](#).⁵ I find that wage sorting is well-explained by technological differences between jobs and workplaces, while the mechanical sources of increased wage sorting are consistent with trends studied in the literatures on technical and structural change. Other papers consider the impact of institutional (i.e. firm-side) features affecting rents and rent-sharing, for example [Krueger and Summers \(1988\)](#), [Fortin and Lemieux \(1997\)](#), and more recently [Card, Cardoso, Heining, and Kline \(2018\)](#). It is often argued that institutional changes can account for regional differences in wage trends, a frequently-studied example being the decline in Brazilian wage inequality. I find that German wage sorting is generally unrelated to market institutions like collective bargaining, and that employer pay is stable at an aggregate level. Finally, several papers study the rise in frictional or “residual” wage dispersion through the lens of match complementarities, either between workers and their peers (e.g. [Kremer and Maskin \(1996\)](#)), or between workers and firms as in [Helpman, Itskhoki, and Redding \(2010\)](#).⁶ If heterogeneity is unobserved then match rents will not be priced out and may pass through to wages; and an increase in the strength of the complementarities or a decline in search frictions will tend to increase wage dispersion. I find that wage sorting is absent within narrow labor markets where such frictions should be strongest, and that the relevant variation in person and firm wage effects is strongly predicted by observable characteristics. Moreover, wage sorting is *negatively* related

⁵See also [Goos and Manning \(2007\)](#), [Autor, Dorn, and Hanson \(2013\)](#), [Michaels, Natraj, and Van Reenen \(2014\)](#), [Frey and Osborne \(2017\)](#), [Autor and Salomons \(2018\)](#), and [Acemoglu and Restrepo \(2020\)](#).

⁶Closely related is the literature on assortative matching, which includes [Eeckhout and Kircher \(2011\)](#), [Lise, Meghir, and Robin \(2016\)](#), [Hagedorn, Law, and Manovskii \(2017\)](#), [Bagger and Lentz \(2019\)](#), and [Bonhomme, Lamadon, and Manresa \(2019\)](#).

to job tenure, worker experience, and establishment age, implying that sorting patterns are not constrained by informational frictions.

This chapter also contributes to the literature on industry and occupation wage gaps. Previous studies have looked at the between-industry and between-occupation components of rising wage dispersion from firm effects (Abowd et al. 2012; Akerman et al. 2013; Card, Heining, and Kline 2013; Barth, Bryson, Davis, and Freeman (2014); Torres et al. 2018), but have not directly explored their relationship to, or implications for, the theoretical literature on inequality. In a closely related paper, Haltiwanger and Spletzer (2020) examine matched data from the United States and show that industrial and occupational classifications explain the majority of rising wage variance over the years 1996-2015. Those authors find, as I do, that the principal drivers are changes to occupation mean wages and occupation-industry sorting. They do not decompose earnings into firm and worker components, however, and their results are silent on the ability of industry-average firm premia to explain the rise in U.S. wage sorting observed by Song et al. (2019).⁷ Given the similarity of the wage trends in the U.S. and Germany, it is likely that the results shown in this chapter will also hold for the United States.

Next I give a brief description of the theoretical mechanisms evaluated in this study. In section 2, I provide background on the AKM wage decomposition and German wage sorting, as well as results attesting to the robustness of wage sorting and its contribution to rising wage inequality. Section 3 contains the main results on between- versus within-market components of wage sorting. The observable correlates of between-market wage sorting are studied in section 4. Section 5 concludes this chapter.

1.1.1 Labor Sorting and Wage Sorting

Labor sorting is a natural prediction of two classes of models. The first class of models studies the technical demand for skilled and unskilled labor, which is generally supposed to vary across regions or industries. An example of this is Autor and Dorn (2013), who consider how a decline in demand for manufacturing labor affects skill premia. Although this literature

⁷In a recent working paper, Haltiwanger, Hyatt, and Spletzer (2022) do employ an AKM decomposition; their preliminary results indicate that roughly 1/3 of the rise in U.S. wage variance over 1996-2018 is related to industry firm premia and the sorting of workers across industries.

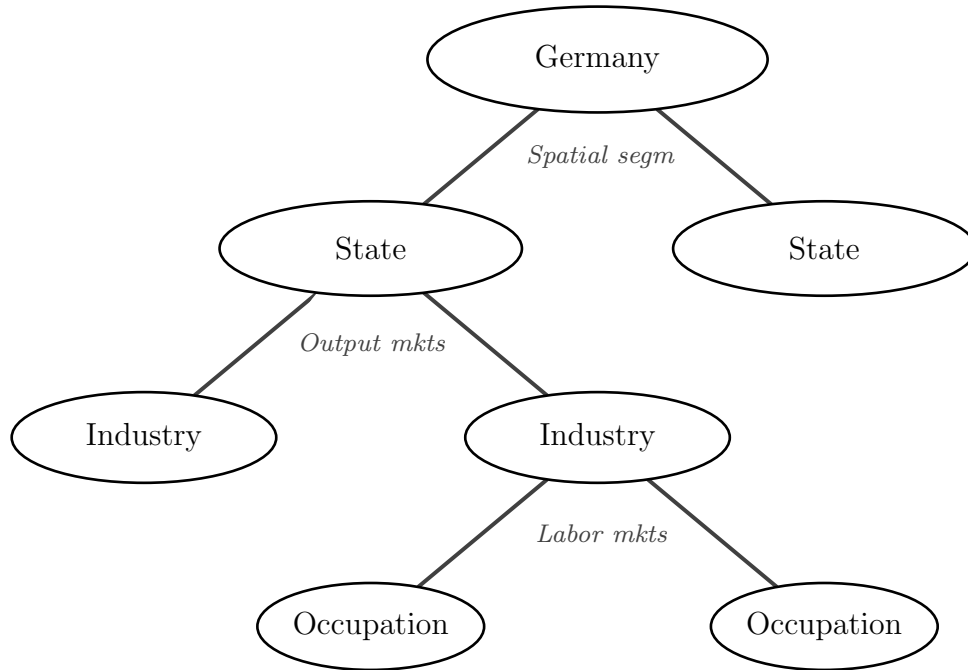


Figure 1.2: Labor Market Segmentation

usually assumes that labor markets are perfectly competitive, and that firms play no direct role in wage-setting, extension to the case of imperfect competition is straightforward. The second class of models studies worker-firm or worker-worker complementarities, which lead to sorting and - when markets are frictional - create rents that may pass through to wages. The canonical example is [Becker's \(1973\)](#) model of marriage markets, which was extended to the frictional case by [Shimer and Smith \(2000\)](#). In this framework, matching between agents is *assortative* if high-ability workers seek out high-productivity firms.

The key difference between these two environments is that match-based sorting has strong implications for the wage function, while technical sorting does not. If sorting is technical in nature and occurs *across* segmented labor markets - for example regions, industries, and occupations, as illustrated in 1.2 - then higher match rents in a particular market will induce entry by firms and/or workers, and either quantities or prices will adjust until the rents are eliminated. In other words, between-market wage differentials will reflect types and technologies rather than complementarities, though this does not rule out an equilibrium relationship between labor sorting patterns and wages. By contrast, in models of assortative matching firms produce homogeneous output and type distributions are fixed. Rents cannot

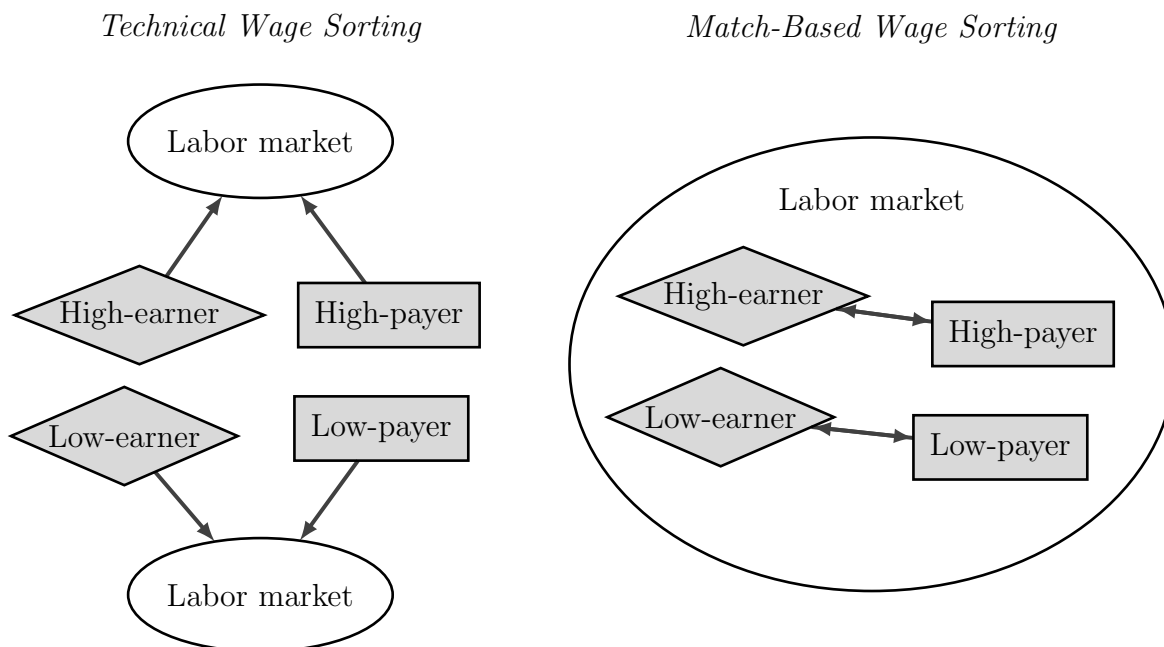


Figure 1.3: Technical and Match-Based Wage Sorting

be eliminated through price or quantity adjustments, and so complementarities show up directly in wages.

The distinction between technical and match-based sorting is important for understanding how environmental change (including policy) affects wages. The past studies motivating this chapter have shown that wage inequality, and wage trends, are substantially impacted by patterns of sorting in labor markets. But how does labor sorting depend on the environment? And how do sorting patterns impact wages? These questions rest critically on the underlying mechanism. If sorting is technical in nature then quantities will be responsive, and the association between firms and workers more variable. If sorting is match-based then quantities are less prone to respond, but wages more so. The key behavioral margins are different: entry and exit in the first case, and job search in the second.

A more practical difference between these two theories of wage sorting is the extent to which they can be empirically characterized. Because market segmentation is observed, differences between markets can be characterized, even if only in terms of quantities and prices. But it is also more likely that we can characterize wage sorting in terms of observable characteristics of agents and markets, and there is more hope in such a case of

evaluating specific causal mechanisms, and identifying *which* features of markets, and which dimensions of heterogeneity are important. If instead wage sorting is a story about match effects and unobservable differences in worker ability or firm productivity, then the empirical challenges are significant. Indeed the wage decomposition shown in figure 1.1 becomes difficult to interpret in this case, because such decompositions are mis-specified when match complementarities are present.⁸ Hence our ability to make empirical progress on the causal origins of wage sorting depends, to a large extent, on whether it is a between-market or within-market phenomenon.

1.2 West German Wage Sorting, 1993-2017

In this section I describe the dataset and the AKM wage decomposition used to identify worker and firm wage effects. I then introduce the trend that is the focus of this chapter: a rise in the covariance of person and firm effects over the 1993-2017 period. The remainder of the section focuses on the robustness of this trend. I present results by group, which attest to the pervasiveness of wage sorting in German labor markets; and I consider whether the “limited-mobility bias” known to affect AKM variance moments is cause for concern when comparing measures of wage sorting over time.

1.2.1 Data and Decomposition

LIAB linked employer-employee dataset. The primary dataset used in this analysis is the German linked employee-employer dataset (LIAB), provided by the Institute for Employment Research (IAB). Every year since 1993 the IAB has conducted a stratified survey of German business establishments, collecting information on operational activities, investment, and hiring.⁹ A matched employee-employer dataset is constructed by pulling administrative social security records for all individuals employed by the surveyed established on June 30 of the survey year. The social security records include the occupation and (top-coded) daily wage associated with the employment spell, along with basic information on demographics and

⁸See e.g. [Lopes de Melo \(2018\)](#).

⁹An establishment is defined as a physical workplace, however locations may be aggregated when they share the same corporate ownership, industry classification, and municipal code.

education.¹⁰ In any given year, the LIAB includes between 4 and 15 thousand establishments and 1.5 to 2.5 million workers, representing approximately 5% of the German workforce.

The approach taken in this chapter builds on that of [Card, Heining, and Kline \(2013\)](#) - henceforth CHK - who were the first to document the contribution of wage sorting to rising German wage inequality. Like those authors, I limit attention to full-time West German workers aged 20-60. In addition to allowing for greater comparability with CHK, this restriction ensures that wages are measured in a consistent manner. Part-time labor is problematic, both because it is recorded inconsistently over time as a result of changes to social security tax codes, and because I do not observe hours worked. Apprentices are paid both in wages and non-wage benefits (e.g. certification), making comparisons with the larger workforce problematic. For similar reasons it is sensible to exclude workers at the extremes of the age distribution. East German establishments are only observed beginning in 1996, and are omitted for consistency; however results for East Germany and for the combined East and West are provided in the appendix. Unlike Card et al., who separated males and females for computational reasons, I consider both sexes jointly.

For analyses involving industry and occupation, I rely on aggregate groupings that preserve - to the extent possible - between-group differences in the AKM wage effects discussed below. On the one hand this allows me to avoid small numbers of establishments when estimating group-specific results, and to better satisfy the confidentiality requirements imposed by the data provider. On the other, aggregation minimizes the impact of coding changes over time, most importantly a shift to a new occupational coding system in 2011.¹¹ The appendix contains details on the industry and occupation groups used, as well as (in many cases) results using less-aggregated codes. Industry aggregations are similar but not identical to NACE sections, which in practice provide too much differentiation of service sectors and not enough between goods-producing industries. Occupational codes are based on the KLDB 1988 system, which is not hierarchical. [Blossfeld \(1985\)](#) and [Schimpl-Neimanns \(2003\)](#) proposed a system of aggregation into 12 occupational groups, which I find to be too aggregated

¹⁰Wages are top-coded at the social security contribution thresholds. Tobit regressions are used to impute affected values following CHK; see appendix for details.

¹¹Time-consistent codes are provided by IAB, but in some cases these rely on imputation; therefore I propagate industry and occupation codes forwards or backwards (as applicable) when a job spell is observed on both sides of a coding change.

to preserve the AKM wage structure; for example there are notable differences, in terms of both wage effects, between jobs that involve processing goods (e.g. on an assembly line) and those that involve direct operation of machines. I find that 12 industry groups and 15 occupation groups are sufficient to capture the vast majority (nine-tenths) of the overall variation of AKM wage effects across 3-digit industries and occupations.

AKM wage variance decomposition. Provided with the LIAB are updated versions of the person and employer wage effects estimated by CHK, following the now-standard approach of [Abowd, Kramarz, and Margolis \(1999\)](#). Effects are obtained by implementing the panel wage regression

$$w_{i,t} = \pi_i + \phi_{j(i,t)} + x'_{i,t}\beta + \epsilon_{i,t} , \quad (1)$$

where $w(i, t)$ is the log daily wage of person i in year t , π is a time-invariant wage effect associated with person i , ϕ is a time-invariant effect associated with the establishment j that employs i at time t , and x is a vector containing year fixed-effects and a cubic polynomial in worker age, interacted with dummies for educational attainment. Estimation is performed on the population-level datasets from which the LIAB is extracted, in four partially-overlapping panels that each span 7-8 years. From equation (1), the wage variance can be decomposed as a sum of the variances and covariances of the estimated regression effects, which allows one to see how the underlying sources of wage dispersion have evolved over time. Panel decompositions for West Germany are shown in table 1.4. Although regression (1) may be estimated directly on the LIAB, in practice this is undesirable. Identification of the wage effects comes from job-movers - individuals who move between establishments *within the sample* - and because the LIAB represents but a small portion of the German workforce, the large majority of employer transitions result in entry or exit from the sample. This results in a severe loss of sample, especially among smaller employers, and it exacerbates the limited-mobility bias discussed later in this section.

It is important to note that there are several sources of error and bias that may affect the variance moments in table 1.1. First, wages are top-coded at the upper contribution

Table 1.1: AKM Variance Decomposition, 1993-2017

	1993-99	1998-04	2003-10	2010-17
$\text{Var}(w)$	0.1684	0.1997	0.2321	0.2316
$\text{Var}(\pi)$	0.1088	0.1238	0.1399	0.1415
$\text{Var}(\phi)$	0.0310	0.0381	0.0518	0.0399
$\text{Var}(x'\beta)$	0.0039	0.0054	0.0054	0.0131
$\text{Var}(\epsilon)$	0.0126	0.0149	0.0157	0.0181
$2 \times \text{Cov}(\pi, \phi)$	0.0160	0.0228	0.0250	0.0338
$2 \times \text{Cov}(\pi, x'\beta)$	0.0018	0.0000	0.0000	-0.0186
$2 \times \text{Cov}(\phi, x'\beta)$	0.0014	0.0016	0.0024	0.0008
Observations	10,645,769	9,185,412	9,511,130	7,080,688
Persons	3,351,593	3,301,936	3,097,049	2,347,598
Establishments	8,151	18,518	19,989	17,684

NOTE: Daily wage w denoted in log 1995 euros. Variance components sum to equal $\text{Var}(w)$. Time-varying effects $x'\beta$ are not provided, and are estimated *ex post* and therefore subject to error.

levels for social security, and the affected values are imputed *via* Tobit regressions. Second, time-varying effects are not provided and I estimate them *ex post*; moments based on these effects are provided for reference, and should be interpreted with caution. Third, as mentioned above the data is limited to a subset of the employed German workforce. While this allows for greater comparability over time and with the extant literature, it also means that results may differ from those that would be obtained for the full population of German workers. Finally, it is well-known that the variances $\text{Var}(\pi)$ and $\text{Var}(\phi)$ are biased upwards, while $\text{Cov}(\pi, \phi)$ is biased downwards. As this bias directly affects the moment under study - that is, $\text{Cov}(\pi, \phi)$ - I devote a separate section to it below.

1.2.2 West German Wage Sorting

Wage sorting, or the covariance of the person and firm effects in table 1.1, accounts for 9.5% of total wage variance over the years 1993-1999. By 2010-2017 this contribution increases to 14.6%. In terms of trends, the increase in dispersion from wage sorting is mechanically responsible for 29% of the rise in wage variance over this period. How should we interpret this result? The AKM person effect is the wage level associated with a given worker, controlling for their place of employment. This term will reflect the return to ability, human capital, and

“skill” more broadly, but it may also pick up higher wages associated with age and experience, which are imperfectly controlled for by the time-varying effects in regression (1); and since there is no match-specific term in the regression equation, and a typical person is observed at only a few employers over an 8-year period, π may reflect higher than average match effects. The establishment wage effect is a measure of the average wage change associated with workers moving to that establishment, and will capture higher wages due to firm-specific characteristics like bargaining agreements, rents and rent-sharing, and compensating differentials. It may also, however, capture firm-level variation in average match effects, local price levels, and other circumstances that do not reflect firm heterogeneity but nevertheless impact ϕ . A large and positive value of $Cov(\pi, \phi)$ indicates that higher-earning workers tend to work for higher-paying establishments. This could be the result of sorting on agents’ characteristics, but it may also be due to factors unrelated to agent heterogeneity that exert a similar effect both π and ϕ .




Figure 1.4: West German Wage Sorting, 1993-2017

NOTE: Solid (dashed) line indicates panel (annual) wage sorting value of $2 \times Cov(\pi, \phi)$. Panels are estimated over 1993-99, 1998-04, 2003-10, and 2010-17.

In figure 1.4 we may observe that $Cov(\pi, \phi)$ has risen both across panels and within them, suggesting that the upward trend is due solely to changes in the wage effects associated with specific workers and establishments.¹² The increase is roughly linear, though the covariance

¹²Care must be taken when drawing this conclusion, as there is entry and exit of both workers and establishments. It is therefore possible for nominal changes to show up within-panel. If, for example, there is an increase each year in the return to skill, this may manifest as higher person effects for skilled workers who only show up in the final years of that panel.

is markedly lower during the 2003-2010 panel than in the panels before and after. This deviation from trend appears to be due, at least in part, to statistical bias, which I discuss in the next section. Note that as the covariance of the wage effects has experienced a proportionally greater rise than either of the variances, the *correlation* of π and ϕ also increases over this period, from .14 to .22; this result is provided in the appendix. The upward trend in figure 1.4 is smaller than that shown by CHK, reflecting the inclusion of female workers in the same; $Cov(\pi, \phi)$ is initially higher among females but rises more quickly among males. There is in addition some wage sorting across genders - male workers tend to be both higher-earning and to work for higher-paying firms - but this pattern has weakened over time, further dampening the overall trend.

Table 1.2: Wage Sorting By Group

	1993-99	1998-04	2003-10	2010-17
Male	0.0064	0.0168	0.0224	0.0344
Female	0.0138	0.0152	0.0082	0.0216
Aged 20-30	0.0046	0.0104	0.0084	0.0224
Aged 31-40	0.0150	0.0218	0.0232	0.0342
Aged 41-50	0.0206	0.0252	0.0286	0.0362
Aged 51-60	0.0182	0.0266	0.0262	0.0326
Lower secondary ed.	0.0196	0.0210	0.0196	0.0314
Apprenticeship	0.0100	0.0144	0.0140	0.0236
Upper secondary ed.	0.0116	0.0194	0.0166	0.0194
University degree	0.0168	0.0198	0.0234	0.0138
Schleswig-Holstein	0.0136	0.0248	0.0160	0.0210
Hamburg	0.0172	0.0262	0.0408	0.0020
Lower Saxony	0.0128	0.0142	0.0138	0.0404
Bremen	0.0210	0.0262	0.0322	0.0412
North Rhine-Westphalia	0.0130	0.0222	0.0196	0.0302
Hesse	0.0164	0.0256	0.0226	0.0318
Rhineland-Palatinate	0.0112	0.0188	0.0176	0.0398
Wurttemberg-Baden	0.0156	0.0254	0.0250	0.0378
Bavaria	0.0152	0.0184	0.0288	0.0322
Saarland	0.0102	0.0060	0.0218	0.0272
Berlin	0.0260	0.0336	0.0408	0.0434

NOTE: Shown are the within-group values of $2 \times Cov(\pi, \phi)$, where π (ϕ) is the person (establishment) wage effect.

Group-level results are given in table 1.2, and they show that wage sorting is demograph-

ically robust. The covariance of the wage effects is in all cases positive, and almost uniformly large. Similar levels and trends are observed across most German states, and across different age groups. Nevertheless some differences are evidence. As mentioned above, the increase in wage variance is concentrated among male workers, and a similar pattern may be observed for educational attainment: over the full sample period, the wage effects covariance increases substantially for workers with lower levels of education, but not for those with a university degree or higher¹³ The results shown in the next section suggest an explanation for this: skilled labor is heavily concentrated in a small set of high-paying industries, whereas lower-skilled labor is more sectorally diverse, and hence there is greater variation in the firm effect. This in turn translates into a higher covariance. For the same reason, changes to the industry composition of employment - and in particular a shift from middle-paying to low-paying sectors - have had a greater effect on wage sorting among low-skilled workers. A similar dynamic can account for the difference between sexes, as female workers are overwhelmingly employed in service sectors that have been relatively immune from changes to the industry composition of employment.

1.2.3 Limited-Mobility Bias

A well-known shortcoming of the AKM wage decomposition is that, due to the incidental parameters problem, the variance components in table 1.1 are biased.¹⁴ As mentioned above, the wage effects are identified from job-movers, and the number of job-movers associated with a given employer may be quite small, even when observed over a period of 7-8 years. This is particularly the case for small establishments. For this reason the law of large numbers does not apply, and estimation of the AKM effects is unbiased but inconsistent. Because wages are additive, positive errors in the estimation of ϕ show up negatively in π , biasing downwards $Cov(\pi, \phi)$, whereas $Var(\pi)$ and $Var(\phi)$ are biased upwards. This “limited-mobility bias” is reduced when using population-level data, as both employers in a job transition will be contained within the sample, but it cannot be altogether eliminated. Therefore an important

¹³Note that educational composition is not stable over this period, as the share of workers with a university education has increased substantially over this period. In addition, the share of missing values increases over time; the data provider uses harmonization and imputation when possible to recover these values, but with imperfect success and with the addition of errors from the imputation process.

¹⁴See for example [Andrews et al. \(2008\)](#) and [Bonhomme et al. \(2020\)](#).

question is whether this bias varies over time, as any time variation would directly impact the wage sorting trend being studied.

Table 1.3: AKM Identification Statistics

	1993-99	1998-04	2003-10	2010-17
<i>Unidentified Sample (%)</i>				
All establishments	0.0219	0.0278	0.0296	0.0226
1-4 employees	0.1038	0.1239	0.1328	0.1181
10-24 employees	0.0052	0.0065	0.0131	0.0047
25-99 employees	0.0039	0.0046	0.0043	0.0029
100-499 employees	0.0035	0.0040	0.0049	0.0040
500+ employees	0.0029	0.0047	0.0046	0.0028
<i>Job Transitions (%)</i>				
Job entry	0.151	0.150	0.133	0.154
Job exit	0.172	0.168	0.153	0.166

NOTE: Job entry (exit) is the number of new (ending) job spells associated with full-time workers at continuing establishments, as a percent of total full-time employment.

The identification statistics shown in table 1.3 indicate important differences between panels in terms of the quality of the AKM effect estimation. A rough measure of the error associated with the wage effects is the proportion of the sample that cannot be identified, due to a lack of job-movers connecting individual establishments to the rest of the sample. The unidentified percentage is slightly above 2% for the 1993-1999 and 2010-2017 panels, but closer to 3% for the intervening panels. This could in principle be due to greater employment at small employers during the early 2000's, as these suffer from higher rates of non-identification; but a similar pattern is observed when we look only at establishments with less than 10 employees. Rather the explanation appears to be a decline in job transition rates during the 2000's, which shows up in both hires and separations.¹⁵ Fewer job movers means that wage effects are less precisely estimated, indicating that limited-mobility bias is likely to be more severe during the 1998-2004 and 2003-2010 panels.

Further evidence of time-varying bias is shown in figure 1.5, where wage sorting is plotted by establishment size. A decline in job transitions should have the greatest effect on estimation error - and hence on limited-mobility bias - at small employers, among which

¹⁵This decline in entry/exit rates is concurrent with the Hartz reforms (2003-2005), and lasts until the early 2010's.

the deviation from trend should be greatest. This is indeed what we see. The lower covariance during 2003-2010 is concentrated among very small establishments, and non-existent for those with 25 or more workers.¹⁶ While it is always possible that some other change occurred during this period that disproportionately affected smaller establishments, the pattern in figure 1.5 is consistent with greater statistical bias during the 2003-2010 period, and this would seem to be the simplest expl




Figure 1.5: Wage Sorting By Establishment Size

NOTE: Lines indicate value of $2 \times Cov(\pi, \phi)$. Employer size is calculated as the number of full-time workers.

To summarize, lower job transition rates during the first decade of the 2000's appear to have exacerbated limited mobility bias, and are likely to explain the deviation from trend observed in figure 1.4 for this period. Bias appears to be less of a concern over the full sample period; identification statistics are comparable for the 1993-1999 and 2010-2017 panels, which provides reassurance that the secular trend in $Cov(\pi, \phi)$ is real and not a statistical artifact. The results shown above are nevertheless concerning, as it is becoming more common for studies to perform panel implementations of the AKM variance decomposition in order to make comparisons over time, and limited-mobility bias is still perceived as an issue peculiar to *cross-sectional* comparisons of the variance moments. The AKM effects studied here are estimated under the best possible conditions - population-level data and relatively long panels - yet time-variation in the statistical bias associated with $Cov(\pi, \phi)$ is still large enough to

¹⁶The trend for establishments with 100-499 employees is similar to that shown for the 25-99 size class; for very large establishments wage sorting is negligible in all panels, reflecting the fact that such establishments are concentrated in a handful of similar sectors, and exhibit little variation in ϕ .

offset the empirical trend being studied. Nor is there any reason to expect that this is an anomalous result, as labor mobility is understood to be an endogenous outcome, depending for example on economic conditions and labor protection laws. While there is no direct way to measure limited-mobility bias, summary statistics like those in table 1.3 can provide an informal test of whether comparisons between two periods of time are likely to be valid, and would constitute good practice for future studies.

1.3 Is Wage Sorting Within or Between Markets?

In this section I use group decomposition methods to study the within-market and between-market components of wage sorting. First I describe the decomposition, and provide results for the three kinds of labor market differentiation discussed in the introduction: geography, industry, and occupation. These results indicate that wage sorting is entirely between industry-occupation groups, and inconsistent with match-level (within-market) sorting. I provide further evidence that this is the case by considering the relationship between labor market exposure - person experience, establishment age, and job tenure - and wage sorting, which in the case of match complementarities should be stronger as agents select into “better” matches. Finally, I examine how changes to wages and shares at the industry-occupation level have contributed to rising dispersion from wage sorting. By conducting a set of reduced-form counterfactuals, I disentangle the role of changes to what firms pay - for example due to shifts in market structure - and changes to what individuals earn and where they work, most likely reflecting developments on the demand side of labor markets.

1.3.1 Between-Market Wage Sorting

The group decompositions used in this section draw on the fact that any covariance can be decomposed into two components: the average of the covariance within each of a set of groups, and the covariance of the group means. More precisely, for any two variables x and y and a partition of the sample $G = \{1, \dots, N\}$, where ω_g is the employment share of group

g , we have

$$Cov(x, y) = \underbrace{\sum_{g \in G} \omega_g (x - \mathbb{E}_g[x]) (y - \mathbb{E}_g[y])}_{\text{within-group covariance}} + \underbrace{\sum_{g \in G} \omega_g (\mathbb{E}_g[x] - \mathbb{E}[x]) (\mathbb{E}_g[y] - \mathbb{E}[y])}_{\text{between-group covariance}}, \quad (2)$$

where \mathbb{E}_g and \mathbb{E} are the within-group and unconditional expectations. Supposing that G corresponds to occupational classifications, the within-group component of $Cov(\pi, \phi)$ would be positive if high-paying employers tend to hire better-skilled workers within a particular type of job; for example if more capable or experienced managers are more likely to benefit from a large value of ϕ . The between-group component would be positive if higher-paying (e.g. skilled) occupations comprise a large portion of the workforce at high-paying employers, for example if these employers require a disproportionately large number of managers.

In addition to offering a straightforward approach for separating labor sorting patterns into between- and within-market components, group decompositions are especially well-suited to the study of empirical wage effects. The reason is that *mean* wage effects are consistent even when the effects themselves are not. Provided that there is a sufficient number of firms within each group, limited-mobility bias will affect the within-group moments (variances and/or covariances) but not the between-group moments. This will still be approximately true even if groups are dominated by a handful of large employers, since as was shown in the previous section, limited-mobility bias is predominantly driven by estimation errors associated with the wage effects of small employers.

Table 1.4: Between-Group Wage Sorting (% Total)

	1993-99	1998-04	2003-10	2010-17
Industry (12)	0.845	0.823	1.002	0.829
Occupation (15)	0.945	0.841	0.981	0.734
State (11)	0.067	0.037	0.051	0.027
Ind. \times Occ.	1.155	1.066	1.285	1.035
Ind. \times State	0.893	0.852	1.032	0.827
Occ. \times State	0.981	0.867	1.020	0.758
Ind. \times Occ. \times State	1.182	1.091	1.305	1.037

NOTE: Value shown is the between-group component of $Cov(\pi, \phi)$ divided by the total covariance. Group size N is indicated in parentheses.

Results in table 1.4 show that while wage sorting is almost entirely within German states,¹⁷ the preponderance of West German wage sorting is between-industry and between-occupation. In each of the four panels, more than 4/5 of the wage effects covariance is between the twelve aggregated industry groups. A similar proportion is between the fifteen occupational groups, though over time an increasing proportion of wage sorting is within-occupation. The intersection of industry and occupation accounts for more than the entire positive value of $Cov(\pi, \phi)$, with the negative within-group component likely reflecting limited-mobility bias as discussed above. Hence industries with high-paying establishments tend also to employ high-earning workers, while occupations with large average person effects are relatively more concentrated within high-paying employers. Note that the ability of industry and occupation to account for the covariance of the wage effects is not due to their overall explanatory power: I show in the appendix that most of the variation in π and ϕ is within industry-occupation cells. Results are similar when using disaggregated industry and occupation codes, though unsurprisingly these are better able to explain the individual variances of wage effects.

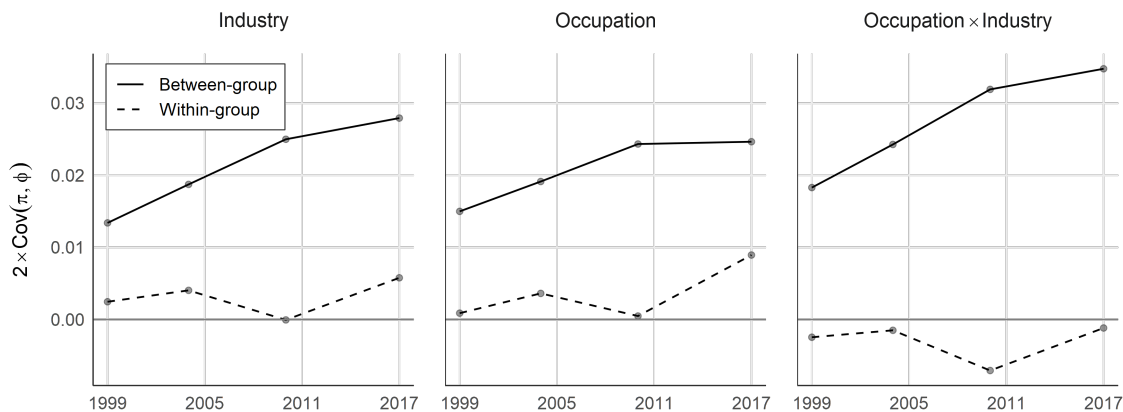


Figure 1.6: Between-Group Wage Sorting, Industry and Occupation

NOTE: Between- and within-group values of $2 \times Cov(\pi, \phi)$ for 15 KLDB 1988 occupational groups, 12 WZ 2008 industry groups, and 180 occupation \times industry groups.

The between- and within-group components for industry and occupation are plotted in

¹⁷Due to censoring, the version of the LIAB used in this study lacks municipal codes. An important question is whether wage sorting reflects differences between urban and rural wage markets. It is important to note, however, that German states are comparable in size to large U.S. counties, and vary substantially in degree of urbanization; hence we would expect to see urban-rural sorting manifest to some extent across states.

figure 1.6, from which it is evident that the deviation from trend in the early 2000's is entirely a within-industry, within-occupation phenomenon. This is further evidence that limited-mobility bias is the likely culprit, since as discussed above we would expect the bias to mostly affect within-group variance moments. The between-group trend is linear over the first half of the sample, but diminishes somewhat in the second half. The evolution of overall wage variance exhibits a similar pattern, suggesting that these trends are more tightly correlated than would be apparent from table 1.1. It is also evident that there is substantial overlap in the amount of covariance "explained" by industry and occupation, reflecting the fact that occupations are often industry-specific, while occupational composition varies markedly between industries.

By itself the group decomposition provides little intuition about the underlying structure of wages, and so in figure 1.7 I plot mean wage effects by industry and occupation. Two patterns are evident. First, the between-group correlations are very high, generally falling in the .6-.8 range. In other words, the average person effect is strongly predictive of the average employer effect, and vice versa. This is necessary to explain why industry and occupation are able to account for most of the covariance $Cov(\pi, \phi)$ despite explaining little of the overall variances of π and ϕ . Second, the group-level correlations would be even higher were it not for a distinct manufacturing wage premium: manual labor occupations associated with goods production tend to earn a higher ϕ than do comparable service sector jobs, while manufacturing sectors pay a substantial premium relative to service sectors with a similar mean person effect. This manufacturing-services wage gap tends to reduce wage sorting, and I find in the next section that it is important for understanding the growing contribution of wage sorting to overall inequality.

Concluding, wage sorting in West Germany is *entirely* a story about jobs: the industrial sector and the occupation in which an individual works. Different types of jobs are associated with different types of workers, as well as different types of employers, and jobs associated with high-paying employers tend to be filled by high-earning workers. Based on the limited information available for this analysis, geographic sorting does to appear to play an important role. Match complementarities are also unlikely to drive wage sorting. As discussed in the introduction, such complementarities will only have strong wage implications if types are

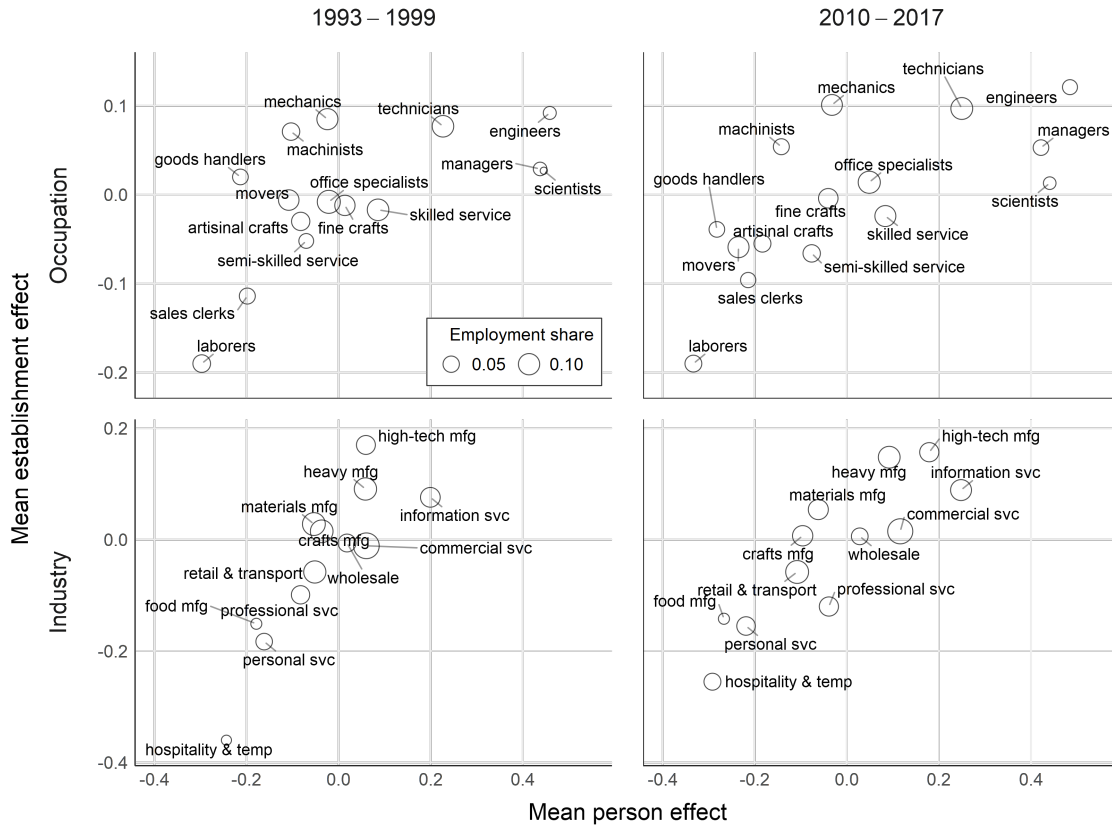


Figure 1.7: Wage Effects by Occupation and Industry

NOTE: Weighted average of AKM wage effects by aggregated KLDB 1988 occupation and WZ 2008 industry.

unobservable, preventing prices and quantities from adjusting so as to eliminate rents. But industry and occupation are observable, and moreover, they are precisely the margins along which price and quantity adjustments in labor markets are usually studied. In the next section, I provide additional evidence ruling out match-based sorting as a causal explanation for the wage patterns studied in this chapter.

1.3.2 Wage Sorting and Match Selection

In dynamic models where agent types are unobservable and sorting occurs due to match complementarities, which in turn generate higher wages, a general prediction is that over time there should be *match selection*. If for example worker skill and employer productivity are complements in match production, then through voluntary separations or on-the-job search, matches between skilled workers and low-productivity firms should dissolve at a faster rate.

In the simple framework of [Shimer and Smith \(2000\)](#), selection occurs immediately through match acceptance sets; one or both parties will reject the match if the match product is too low. But in general we would expect selection to be a gradual process, due to learning over time about match products, culling of poor matches during market downturns, dynamic behaviors such as on-the-job search, and so forth. For example in the framework of [Lentz \(2010\)](#) and [Bagger and Lentz \(2019\)](#), sorting becomes stronger as workers move up the job ladder. Predictions will vary with assumptions, but in general, if wage sorting is driven by match complementarities, we should expect it to be stronger among agents that have had more time or opportunities to find a better match.

To further assess whether the evidence supports a role for match-based sorting, I analyze the relationship between wage sorting and several measures of agents’ opportunity to select into better matches: job tenure, worker time in the labor force, and establishment age. Jobs with greater tenure are more likely to reflect “good” matches, as agents will presumably have had more opportunities to dissolve the match in favor of an alternative one. Likewise, newly-entered workers and establishments are likely to have little choice in the matches they form. I focus on the 2010-2017 panel in which the covariance of the wage effects is largest, and in [figure 1.8](#) I plot the relationship between wage sorting and the measures just described.




Figure 1.8: Wage Sorting and Match Selection, 2010-2017

NOTE: Time in workforce dated from first payment into social security.

For all three measures there is a negative relationship between tenure/experience and wage sorting, in contrast with the positive relationship that would be predicted by match complementarities. The covariance of the wage effects is strongest for workers at the be-

ginning of their careers, newly formed establishments, and freshly-formed job matches. The negative slopes in the figure can be largely accounted by differences in composition; new workers are more likely to work in the service sector, for example, where wage sorting is greater than it is in the aggregate. But these compositional differences occur along the same dimensions as wage sorting - industry and occupation - and they therefore offer little support to match-based sorting as a causal mechanism. Either industry and occupation represent heterogeneity unobservable to agents, and match selection over time is in the wrong direction; or industry and occupation are seen by agents, in which case wage sorting must reflect observable rather than idiosyncratic differences between agents. Hence the patterns in figure 1.8 cast further doubt on ability of match complementarities to explain West German wage sorting.

1.3.3 Decomposing the Between-Group Trend

In this section I consider the mechanical drivers of the increase in $Cov(\pi, \phi)$, making use of the fact that nearly the entire trend is between, and not within industry-occupation groups. This allows the change in covariance to be decomposed into three potential sources: group shares, group mean person effects, and group mean establishment effects. While these sources are not directly informative about the underlying causal mechanism(s), they do tell us about the nature of the equilibrium response. It is then possible to reason backwards about whether the changes we observe are consistent with the those predicted by a particular mechanism. If demand has increased for skilled labor, for example, then we would expect that to manifest through occupational shares and average person effect. If on the other hand increased rents have pushed up wages in the tech sector, we would expect that to materialize as an increase in the industry-average ϕ .

My approach is to conduct a set of reduced-form counterfactuals that isolate the group-level changes responsible for the rise in wage effects covariance. For this exercise I focus on 180 industry-occupation groups, which as shown previously are able to account for the entire covariance in each of the four panels. I pose the question: what would have been the counterfactual change in $Cov(\pi, \phi)$ if group shares and/or wages had remained constant? Moments are held constant in one of three ways: for each the 180 industry-occupation cells;

for the 12 industry groups, while allowing for within-industry, between-occupation variation over time; and for the 15 occupational groups, allowing for within-occupation but between-industry variation. The results of these experiments are given in table 1.5, first for group shares and wage effects separately, and then in combination.

Table 1.5: Counterfactual Between-Group Wage Sorting (% Trend)

	1993-99	1998-04	2003-10	2010-17
Total between-group	0.000	0.361	0.831	1.000
Constant group π	0.157	0.446	0.849	0.747
By industry only	0.187	0.464	0.843	0.759
By occupation only	0.048	0.380	0.813	0.886
Constant group ϕ	0.030	0.271	0.524	0.982
By industry only	0.006	0.277	0.608	1.054
By occupation only	0.133	0.325	0.530	0.777
Constant group ω	0.241	0.518	0.729	0.717
By industry only	0.259	0.506	0.753	0.735
By occupation only	0.030	0.380	0.783	0.952
Constant group (π, ω)	0.482	0.645	0.801	0.530
By (occupation, industry)	0.313	0.524	0.759	0.651
By (occupation, both)	0.331	0.560	0.747	0.645
By (both, industry)	0.470	0.620	0.813	0.542

NOTE: Value shown is the counterfactual between-group component of $Cov(\pi, \phi)$ divided by the total between-group trend for 1993-2017. Groups are 180 industry \times occupation pairs. Person effects (π), establishment effects (ϕ), and employment shares (ω) are held constant at their 1993-99 and 2010-17 values, as indicated, with the average covariance reported.

Comparing the results over the full sample period (columns 1 and 4), changes to the distribution of the person effect π and to group shares ω have each contributed substantially to the rise in covariance, while changes to the group establishment effect ϕ have had almost no impact. Holding fixed the group mean person effect reduces the overall trend by 40%. The main contribution appears to be from changes to industry rather than occupation means; π has increased in high-paying industries. Holding fixed the group shares reduces the overall trend by slightly more than half. This also appears to be the result of industry-level changes, indicating that the distribution of employment across industries has shifted over time so as to increase the extent of wage sorting. Occupational composition appears to have played little role, and to the extent that *conditional* distributions have changed, this has contributed

negligibly to the trend. That is, we do not appear see an increased concentration of high-earning occupations in high-paying industries. Finally, holding fixed both π and ω eliminates virtually all - 95% - of the trend in $Cov(\pi, \phi)$.

Results for the intermediate two periods are somewhat different, and indicate a larger role for changes to the establishment effect. This could in principle reflect either short-term changes to firm pay, or shifts in the composition of firms. A similar trend is also seen in the between-group variance of ϕ ,¹⁸ which exhibits a hump-shaped pattern. Changes to the group establishment effect are unlikely to be the result of bias, both for the reasons stated at the beginning of this section, and because the changes were not random but tended to enlarge the existing pay gaps between industries. Regardless, they were also short-lived and did not persist into the 2010's.

What was the nature of the changes to group shares and person effects? These changes are shown at the industry and occupation level in figure 1.9. Across industries, two developments have occurred. The average person effect has declined in industries with a low π relative to those with a high person effect, most likely reflecting the rising wage gap between skilled and unskilled occupations (as seen in row 2). Because person and establishment effects are so highly correlated across industries, this in turn means that the average person effect has declined in industries with a low ϕ , thereby raising the wage effects covariance. Meanwhile industry growth has been concentrated in high- π , high- ϕ and low- π , low- ϕ industries, with those in the middle - mainly industries related to commodities manufacturing and trade - have lost employment share. Industry polarization has tended to increase the between-industry variance of both wage effects, and because these are correlated, it also pushes upwards their covariance. Across occupations these same trends have been weaker or absent, indicating that the rise in $Cov(\pi, \phi)$ is fundamentally a story about widening inter-industry wage gaps.

It is perhaps surprising that between-group changes to the establishment wage effect have not played a greater role, given the observation of [Card, Heining, and Kline \(2013\)](#) that the increase in (population) $Var(\phi)$ is evidently related to the decline in collective bargaining coverage that occurred following reunification. In un-shown results (see appendix), I find that

¹⁸The between-group variances are shown in the appendix.

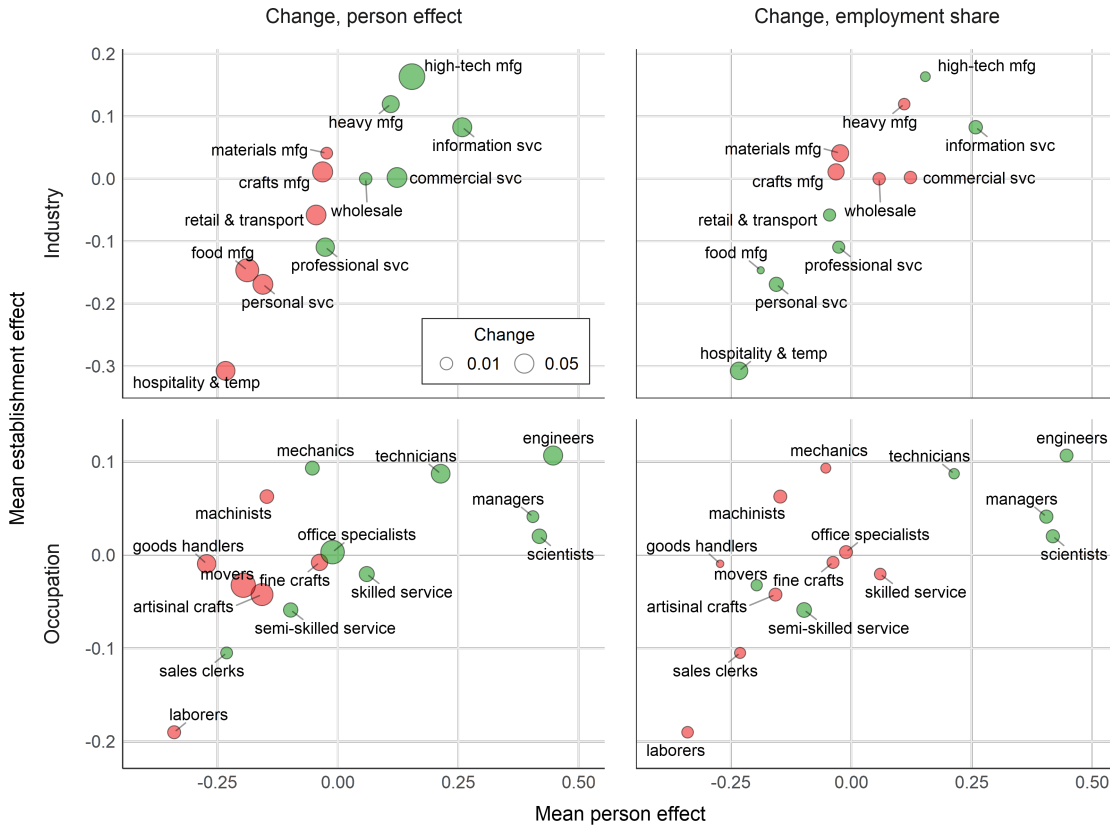


Figure 1.9: Evolution of Mean Person Effects and Shares, 1993-2017

NOTE: Axes indicate group mean wage effects, averaged across 1993-1999 and 2010-2017. Green (red) bubbles indicate an increase (decrease) in group mean person effect and employment shares, measured as the difference between 1993-1999 and 2010-2017 values.

the ϕ -gap between small and large employers has increased in many industries, notably those within manufacturing. Similarly the wage premium associated with collective bargaining coverage has increased, likely reflecting a decline in the extent to which industry bargaining agreements are observed by establishments not formally subject to them. Large, covered firms tend to employ workers with a higher average π - a pattern that holds within industries as well as between them - and so one might expect widening within-industry pay gaps to contribute to wage sorting. At the same time, however, the mean establishment effect has risen for the hospitality and temp agency industry group, which overwhelmingly employs lower-paid workers. These between-industry and within-industry (but between-occupation) changes have offset in the aggregate, resulting in a negligible overall contribution to the trend in $Cov(\pi, \phi)$.

1.3.4 Discussion

Summarizing this section, West German wage sorting is entirely between industry-occupation groups, and there is little evidence of assortative matching within markets. The principle changes over time are a widening of inter-industry π -differentials, and a shift in employment towards low- π , low- ϕ and high- π , high- ϕ industries. At an aggregate level, establishment pay has not been fully stable over the 1993-2017 period, but has contributed little to the overall trend in wage effects covariance. This suggests that the rise in dispersion from wage sorting is not driven by changes to rents, rent-sharing, or market structure, but is instead likely to reflect developments on the demand side of labor markets. This would be consistent with a vast literature relating increased wage inequality to structural change as a result of globalization, and rising skill premia as a result of technology-driven demand for skilled labor.

These results provide definition for the trend discovered by [Card et al. \(2013\)](#), and they indicate that wage sorting - a trend in unobservable wage effects - is closely related to inter-industry wage gaps, and fundamentally a story about observed dimensions of heterogeneity. In the next section I explore this heterogeneity, and provide descriptive results on the characteristics of workers and employers most predictive of wage sorting.

1.4 Descriptive Evidence: Skill, Scale, and Technology

What observable characteristics of workers and firms are associated - causally or otherwise - with wage sorting? The patterns in figure 1.7 are suggestive but not definitive. High-paying industries like vehicle manufacturing and information technology are associated with large, technology-intensive employers. High-earning occupations are those casually thought of as high-skill: doctors, engineers, and so forth. To what extent can these associations be formalized, and what do they suggest about the underlying mechanism?

In this section I quantify the relationship between observable characteristics and wage sorting. I leverage the rich employer survey data present in the LIAB, and supplement this with person-level survey data on occupational tasks. To constrain the exercise, I focus on variables related to the casual associations described above - scale, skill, and technology -

and to the mechanisms discussed in previous sections, namely match quality and market structure. In addition I limit the analysis to the 2003-2010 period, first because the survey data are in general not comparable over time, and second because I wish to abstract from changes in the relationship between industries, occupations, and observable characteristics. The 2003-2010 panel is especially attract as during these years the IAB survey has greater coverage of variables relating to investment and innovation, while the 2005-2006 panel of the BIBB survey is particularly well-suited to this analysis as it features a comprehensive set of task-related questions and is the last survey to include KLDB 1988 occupation codes.

I focus on the correlation of wage effects $Cor(\pi, \phi)$, and how this is changed when controlling for observable characteristics. I prefer correlation for this exercise because, as a normalized measure, it is less influenced by the dispersion of the wage effects, and will not be affected if an observable characteristic is predictive of only one of the two wage effects. The outcomes of interest are partial and semi-partial correlations. Letting X denote an observable characteristic, linear regressions of the wage effects on X yield residuals π_X and ϕ_X . The partial correlation is defined as $Cor(\pi_X, \phi_X)$, whereas there are two semi-partial correlations $Cor(\pi_X, \phi)$ and $Cor(\pi, \phi_X)$. These measures have different interpretations. If for example X is education, then the semipartial correlations will tell us (1) how person-level wage gaps unrelated to education are correlated with firm premia, and (2) how person wage effects are correlated with the component of firm premia unrelated to the educational composition of the firm. The partial correlation indicates the relationship between person wage gaps orthogonal to education and employer wage gaps unrelated to educational composition. Intuitively, while partial correlations provide information about the overall ability of X to account for wage sorting, semipartial correlations are informative about where this explanatory power comes from - i.e. whether it tells us about worker traits or employer characteristics associated with wage sorting.

1.4.1 Wage Sorting and Establishment Characteristics

I begin with an analysis of establishment characteristics, drawing on the employer survey portion of the LIAB. In general the response rate to survey questions is is high, and some variables such as employment are present in all cases as they can be calculated from adminis-

trative records. There are cases where non-responses are problematic, however. Value-added is missing in around a quarter of cases, for example, because a significant number of establishments choose not to report sales revenue, and some of those that do report revenue do not report cost of goods sold (and vice versa). There is no basis for assuming that non-responses are random, and so I address this issue by imputing values for missing cases. I regress each of the variables below on log employment, fixed effects for detailed industry groups, dummy variables for five employment size categories interacted with three industry groups, and year dummies. Logistic regression is used when the response variable is binary. To obtain industry-occupation means, I run a second-stage regression with industry-occupation fixed effects and year dummies. The predicted values for 2003 are used below.

I consider two sets of variables. The first set includes employer characteristics related to scale - a well-known covariate of higher pay - and activities related to technology adoption and innovation, which are widely thought to increase skill requirements at the firm level. Scale, technology, and innovation are not independent; there is, for example, a robust between firm size and information technology adoption.¹⁹ Hence this set of variables is of particular interest as regards wage sorting, in addition to being suggested by the patterns in figure 1.7. The second set of variables I examine is related to market structure. Although the results in the previous section are inconsistent with an important role for *changes* to market structure driving changes to wage sorting, it may nevertheless be that output market competition or wage agreements are explanatory in the cross-section. It may for example be that skill-intensive industries are by nature less competitive, or that bargaining agreements induce firms to employ a higher-skilled workforce; neither mechanism is novel, and both could potentially explain wage sorting patterns.

The relationships between establishment characteristics, wage effects, and wage sorting are given in table 1.6, and it is evident that at this aggregated level, wage variation is well-explained by observable traits. The first two columns in the table give the unconditional correlations between characteristics and AKM wage effects; the third and fourth columns give the semi-partial correlations of the wage effects, as described above; and the fifth and final column gives the key result, which is the partial correlation of π and ϕ . Measures

¹⁹See e.g. the meta-analysis by Lee and Xia (2006).

Table 1.6: Wage Effects and Establishment Characteristics, 2003-2010

Variable (X)	Correlations		Partial correlations		
	$\text{Cor}(X, \pi)$	$\text{Cor}(X, \phi)$	$\text{Cor}(\pi_X, \phi)$	$\text{Cor}(\pi, \phi_X)$	$\text{Cor}(\pi_X, \phi_X)$
No control			0.548	0.548	0.548
Log employment	0.404	0.751	0.267	0.370	0.404
Log value-added	0.461	0.844	0.179	0.296	0.334
Multi-estab. firm (%)	0.490	0.536	0.327	0.338	0.387
Log investment/worker	0.400	0.861	0.221	0.399	0.435
Log ICT inv./worker	0.647	0.809	0.032	0.041	0.054
Product dev. (%)	0.427	0.782	0.236	0.343	0.379
Process impr. (%)	0.413	0.805	0.237	0.363	0.399
Collective barg. (%)	0.070	0.336	0.525	0.556	0.558
Positive profits (%)	0.261	0.019	0.562	0.543	0.562
Competitive mkt. (%)	-0.110	0.265	0.580	0.598	0.602

NOTE: Correlations estimated over 180 industry-occupation cells. Missing establishment characteristics are imputed, and wage effects and characteristics are regressed on fixed effects for year and industry-occupation, with results using predicted 2003 values. Terms π_x and ϕ_x indicate residuals from a regression of wage effects on establishment characteristics. All results weighted by employment; see appendix for confidence intervals.

of establishment scale - employment and value-added - are highly predictive of both wage effects, with unconditional correlations around .4 (person effect) and .8 (firm effect). It is also clear that wage sorting is related to establishment scale, as when this is controlled for the wage effects correlation falls from .55 to .33-.40. Similarly, industry-occupation cells associated with multi-establishment firms (a dummy variable) tend to have high mean values of both π and ϕ . Every several years, establishments are asked whether they have engaged in product development or process improvements; these variables are also predictive of both wage effects, which is unsurprising as they are highly correlated with establishment size. On the other hand measures related to rents and rent-sharing - the percentages of establishments reporting positive profits, a competitive market environment, or binding collective bargaining agreements - are mostly uninformative, as they are strongly correlated with only one of the two wage effects.

The second result of interest in table 1.6 concerns investment per worker, which is a proxy for the capital-labor ratio.²⁰ Capital hold-up is a commonly studied mechanism for generat-

²⁰I do not observe capital stocks and, as the typical firm is only observed for a few years, estimation through the perpetual inventory method is not practical.

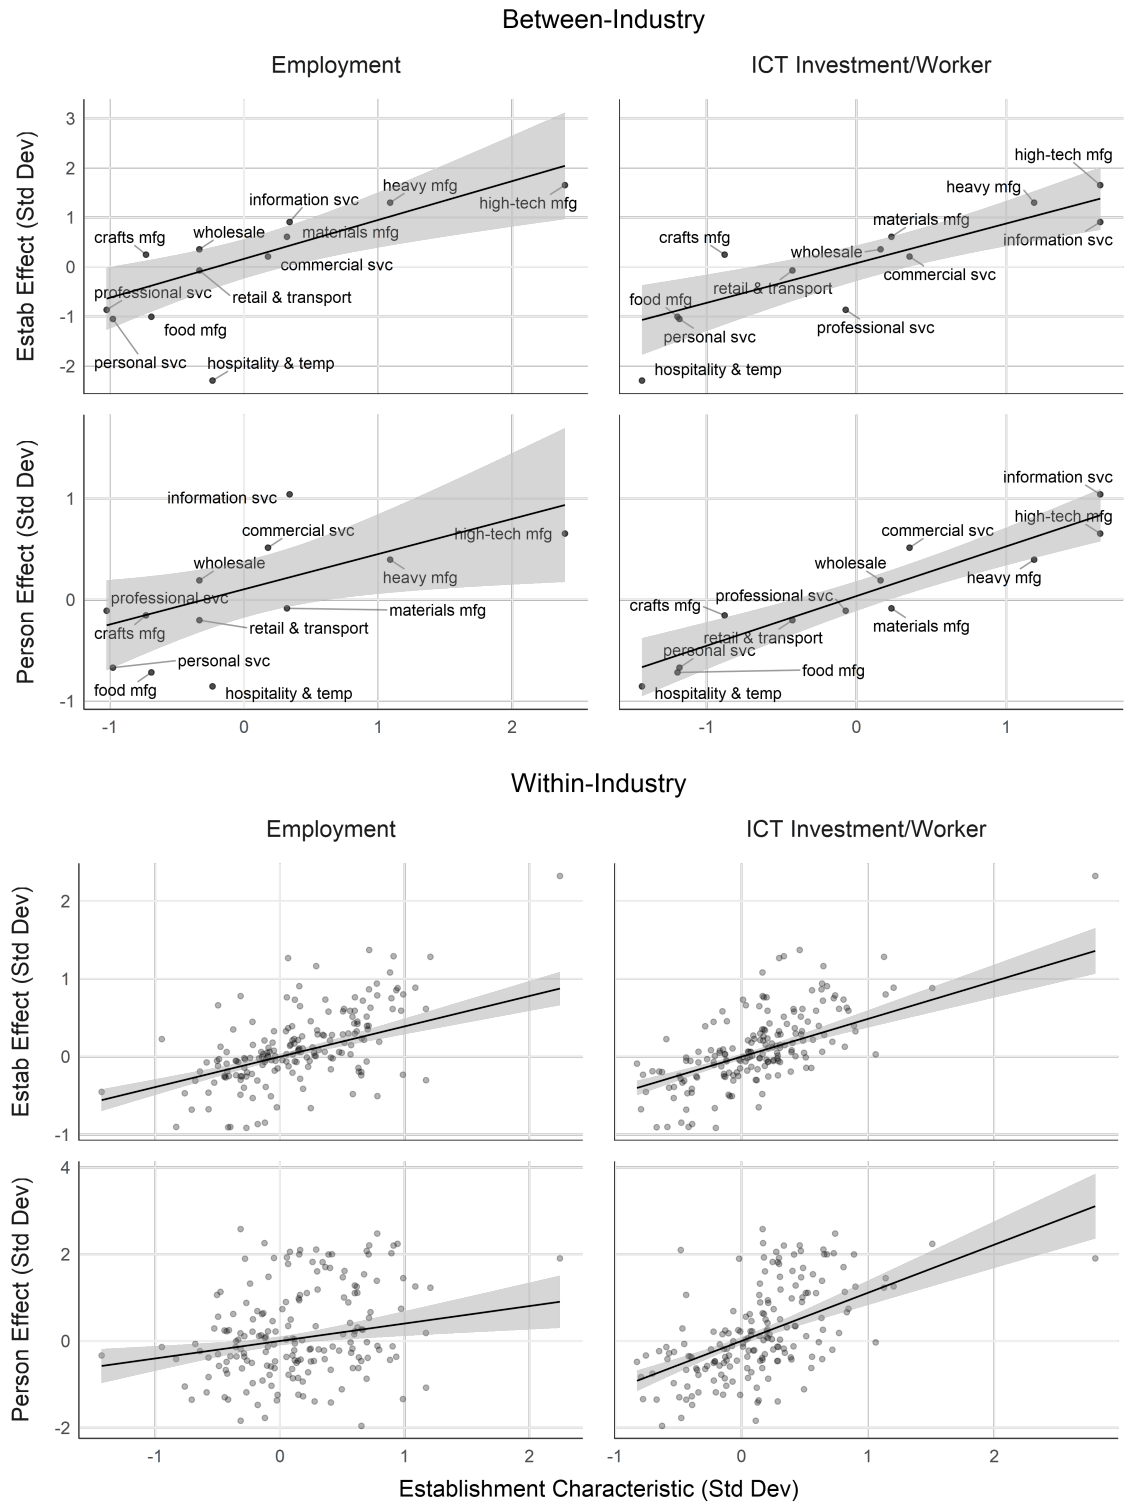


Figure 1.10: Establishment Scale, ICT Investment, and Wage Effects
 NOTE: First two rows: industry mean values. Last two rows: industry-occupation mean values, demeaned by industry. Shaded areas indicate 95% confidence interval of least-squares relationship.

ing pass-through of firm rents, and indeed investment/worker is of all of the variables the one most correlated with the firm wage effect. Investment in information and communications technology is nearly as predictive of ϕ , but is also highly correlated with the person effect π ; and when controlling for ICT investment, the wage effects correlation falls close to zero. This result is quite striking, but also intuitive. The link between information technology and worker skill is well-known and widely studied, and (as shown later in this section) measures of skill are strongly associated with the person effect π . At the same time ICT investment is highly correlated with measures of establishment scale, which can almost entirely account for between-group variation in ϕ .

Why is information technology so capable of predicting both wage effects? In figure 1.10 I compare ICT investment and establishment employment, by plotting the within- and between-industry relationship of each variable with π and ϕ . Between industries, employment is weakly correlated with the person effect. The reason is that size varies less between service sectors than between manufacturing sectors, yielding different linear relationships; hence size is more explanatory within either set of industries than it is overall. This is not true of investment, which exhibits a consistent relationship with π across industries. The difference between the variables is even more stark within-industry, where ICT investment remains strongly associated with person effect but employment does not.

These results cannot speak to causality, though they do suggest a set of plausible mechanisms. That ICT-intensive establishments are skill-intensive may reflect technological skill-bias, or simply greater demand for information technology among skill-intensive firms. That these establishments are also high-paying is consistent with monopsony, insofar as ICT investment per worker is correlated with size; it is consistent with higher wages due to capital hold-up; and it is consistent with other forms of rent-sharing, if capital represents fixed costs that translate into a larger match surplus. The superstar firm hypothesis, as stated by Autor, Dorn, Katz, Patterson, and Van Reenen (2020), predicts a relationship between firm scale, markups (and therefore rents), and information technology. Although the focus of that paper is on changes to market concentration, a similar story would be consistent with the results shown here.

1.4.2 Wage Sorting and Job Characteristics

What characteristics of workers and matches are related to wage sorting? To answer this I supplement data on educational attainment and experience in the LIAB, with task data from the 2005-2006 BIBB employment survey.²¹ The BIBB surveys, which have been performed every 5-6 years since 1979, draw on a random sample of the employed German labor force and ask respondents a range of questions concerning job task content, the use of technology and tools, and other aspects of the job environment and the individual's work history. I limit attention to questions concerning tasks performed on the job, which are answered on a frequency scale ("never", "sometimes", "always") from which I impute numerical values (0, .5, 1). For some industry-occupation cells I observe few or no workers, so I estimate task values through a fixed-effects regression on dummy variables for industry and occupation separately, and interactions between three industry and four occupation groups.²² In principle the 2005-06 survey might be supplemented with earlier and later surveys, however this is the last version of the BIBB survey containing KLDB 1988 occupational codes consistent with the LIAB, while the questions concerning job tasks are different in previous versions of the survey. I apply the same sample selection criteria as with the LIAB - full-time workers aged 20-60 - but it is important to note that there may nevertheless be differences between the two samples. In the LIAB, workers who are not observed are those who do not pay social security taxes, whereas for the BIBB it is those who choose not to answer the survey. Because of this and the differences in survey design (stratified versus randomly sampled), it is likely that the two samples are not perfectly comparable.

Below I focus on a representative set of BIBB task measures, with additional results shown in the appendix. The value of job tasks is that they describe the type of work being done, and for that reason are informative about technical aspects of the job. For this reason tasks are especially useful for the study of interactions between technology and labor markets. From the LIAB I include job tenure, as dynamic search models would predict that this is

²¹The BIBB employment surveys are the work of, and provided by the Federal Institute for Vocational Education and Training (BIBB), in partnership with the Federal Institute for Occupational Health and Safety.

²²Results obtained using cell means are little different from those shown in this section. Because the BIBB surveys are randomly sampled, industry-occupation groups with few respondents account for only a small portion of the German workforce.

associated with match quality and/or higher bargained wages; and I include several measure of human capital: the percentage of workers with a college degree, average years of education imputed as in [Card, Heining, and Kline \(2013\)](#),²³, and time in the labor force as a measure of experience.

Table 1.7: Wage Effects and Job Characteristics, 2003-2010

Variable (X)	Correlations		Partial correlations		
	$\text{Cor}(X, \pi)$	$\text{Cor}(X, \phi)$	$\text{Cor}(\pi_X, \phi)$	$\text{Cor}(\pi, \phi_X)$	$\text{Cor}(\pi_X, \phi_X)$
<i>LIAB Person Characteristics</i>					
No control			0.548	0.548	0.548
College degree (%)	0.823	0.253	0.598	0.351	0.618
Years of education	0.873	0.271	0.638	0.323	0.663
Job tenure	0.162	0.778	0.427	0.671	0.680
Years in labor force	0.193	0.625	0.435	0.547	0.558
<i>BIBB Task Characteristics</i>					
No control			0.548	0.548	0.548
Research, design	0.723	0.455	0.317	0.246	0.356
Use PCs	0.781	0.380	0.402	0.271	0.435
Manage others	0.613	-0.025	0.712	0.563	0.712
Buy, sell	-0.052	-0.504	0.522	0.604	0.604
Control machines	-0.241	0.191	0.612	0.605	0.623
Transport goods	-0.619	-0.266	0.488	0.397	0.506
Host, serve	-0.273	-0.632	0.390	0.484	0.503
Clean	-0.762	-0.407	0.367	0.260	0.402

NOTE: Correlations estimated over 180 industry-occupation cells. Worker characteristics are regressed on fixed effects for year and industry-occupation, with results using predicted 2003 values; regressions for job tasks exclude year dummies. Terms π_x and ϕ_x indicate residuals from a regression of wage effects on establishment characteristics. All results weighted by employment; see appendix for confidence intervals.

The unconditional and partial correlations are shown in table 1.7, from which it is evidence that person and task characteristics are strongly associated with between-group wage effects, as was the case with establishment covariates in the previous section. Experience and job tenure are highly associated with the employer effect ϕ , though only weakly with the person effect; workers in low-paying service sectors are generally younger, and job spells shorter. Tenure and experience do not yield a lower partial correlation, however, indicating

²³The educational categories and assigned years of education used by Card et al. are: missing (10.5 years), less than high school (11 years), vocational education (13 years), upper secondary (15 years), and university degree or higher (18 years).

that wage sorting is at least as strong when controlling for these variables. Education is associated with both wage effects, and accounts for most of the between-group variation in the person effect π .²⁴ Controlling for education yields a higher partial correlation, however, and the semi-partial correlations indicate that while wage sorting is associated with education-intensive employers, it occurs mostly within educational groups.

Task measures are more successful at predicting wage sorting, in particular those related to (1) information-intensive activities and (2) unskilled labor. R&D and PC use are positively and strongly related to both wage effects, while cleaning, serving, and transportation (e.g. packaging and shipping) are negatively related to both π and ϕ . Because wage sorting is present within the manufacturing and service sectors, as well as between them, the most explanatory tasks are those that are not inherently sector-specific; cleaning-related work, for example, is widely dispersed across industries. Many of these tasks are strongly associated with educational attainment, but unlike education, they are predictive at the person level as well as the firm level. The semi-partial correlations show that it is not just R&D-intensive establishments that are associated with wage sorting, but jobs specifically engaged in R&D-related tasks.

Comparing education and the task measure for R&D in figure 1.11, we can see several reasons for the poor explanatory power of the former. Skilled and semi-skilled service jobs are education-intensive, but not R&D intensive; these jobs generally involve application of learned knowledge to a stable set of domains. Mechanical and technical jobs are more likely to require interaction with, or creation of, new tools, technologies, or products. They are also more likely to occur in the context of a high-paying and - most likely - large establishment. Within-occupation, the difference between the two variables is that education is unrelated to ϕ whereas research-related tasks are correlated with both wage effects.

These results are consistent with those for establishment characteristics, and point to an information- or technology-based explanation for wage sorting. This is indicated by the positive role for R&D and PC use, but also by the negative role for unskilled, manual labor tasks, as these two task groups tend not to occur in the same job and have a strong negative

²⁴This result implies that the return to education is not a pure measure of the return to skill, and that the coefficient on education in a standard Mincer regression will reflect sorting of educated workers into high-paying firms.

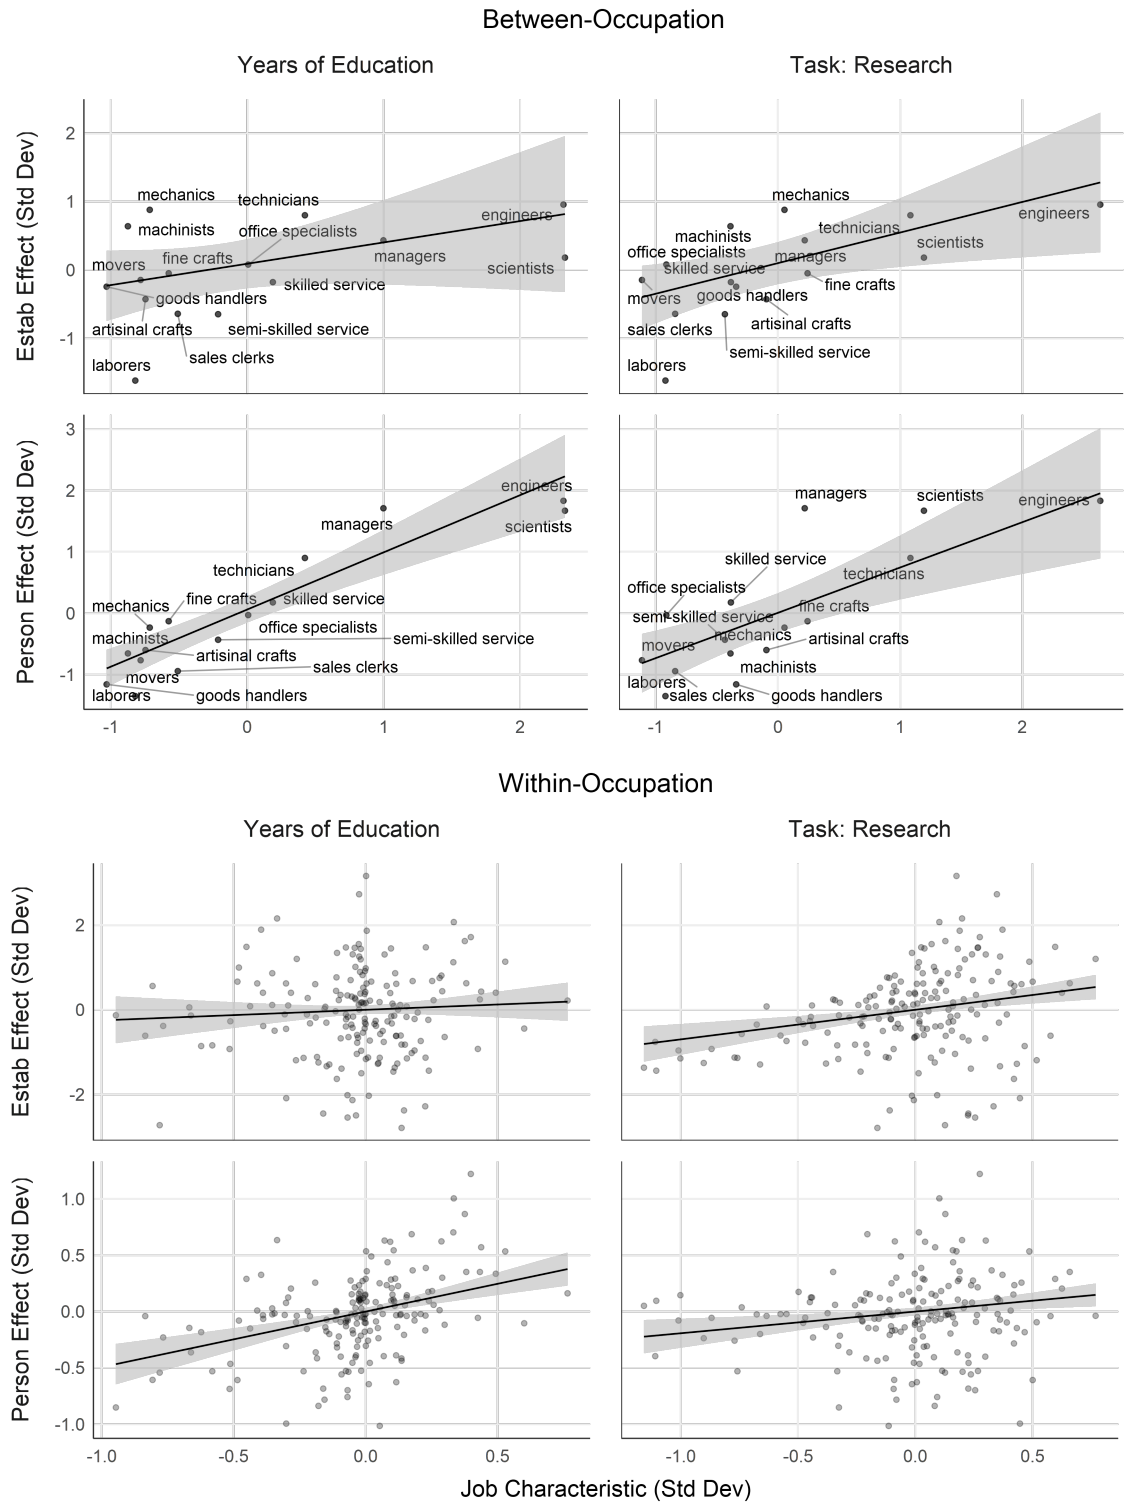


Figure 1.11: Worker Skill, PC Use, and Wage Effects

NOTE: First two rows: occupation mean values. Last two rows: industry-occupation mean values, demeaned by occupation. Shaded areas indicate 95% confidence interval of least-squares relationship.

correlation. The ability of these measures to account for within-occupation variation in π and ϕ is important. It suggests a consistent relationship between wage sorting and observable characteristics, in which case it is more likely that wage sorting reflects a fundamental relationship and not a chance outcome, and that a precise causal explanation might be found and studied.

An additional implication of these results is that, conditional on occupation, industry is informative about task content and worker type. This is interesting because job tasks are usually studied at the occupational level,²⁵ with intra-occupational variation usually seen as idiosyncratic. The results in this section suggest that occupational codes imperfectly capture the task structure in labor markets, and that industry codes can provide a useful supplement. At the same time it is common for papers to study the relationship between occupation-mean characteristics and wages, and in some cases to interpret these relationships in terms of prices. This practice may be problematic if different occupations are associated with different types of employers, and wages reflect not just market prices but differences in the average value of ϕ .

1.5 Conclusion

What do we know about wage sorting? It is, first of all, a robust feature of the German labor market, and not limited to particular geographic regions or demographic groups. Although I find that longitudinal comparisons of wage sorting can be adversely affected by statistical bias that varies with labor market conditions, this does not appear to explain the growing contribution of wage sorting to German wage inequality over 1993-2017. Secondly, we know that in the German case, wage sorting is entirely a between-industry, between-occupation phenomenon. While theoretical studies of sorting in labor markets have typically focused on match-level complementarities, these do not seem to explain the aggregate nature of wage sorting in Germany, which is more consistent with a mechanism based on heterogeneous factor demand across sectors. There is also little evidence that changes to market structure

²⁵See for example [Autor, Levy, Murnane \(2003\)](#), [Goos and Manning \(2007\)](#), [Acemoglu and Autor \(2011\)](#), and [Autor and Handel \(2013\)](#), and for studies specific to Germany [Spitz-Oener \(2006\)](#) and [Gathmann and Schonberg \(2010\)](#).

have contributed to the trend in wage sorting; industry wage gaps associated with firm pay are stable over the full sample period. Instead, the increase in dispersion from wage sorting is well-accounted for by developments associated with technical and structural change: a shift of low-skill labor from manufacturing to services, and rising wages in skill-intensive jobs that, in all periods, are more common at high-paying firms. We know, finally, that observable characteristics of agents are effective at predicting person and firm wage effects at the industry-occupation level where sorting occurs, and that the association between worker earnings and firm pay is substantially or altogether eliminated by controlling for information- and technology-related differences between jobs and workplaces. Taken together, these results indicate a technical explanation for wage sorting, as well as the increased contribution of wage sorting to overall inequality.

One implication of these results, explored in the next chapter, is a likely connection between wage sorting and skill-bias. If technology adoption is associated not only with skilled labor but also with high-paying employers - either causally, or indirectly through an association with firm scale - then there are several channels by which skill-biased technical change may affect wages. Conversely, because industry and occupation capture substantial differences in employer pay as well as worker-firm sorting, technology adoption at the industry-occupation level is unlikely to provide clean results on the presence or magnitude of skill-bias. This is potentially important for empirical studies exploiting such variation.

A second implication concerns the competitiveness of labor and goods markets: if wage sorting reflects fundamental (e.g. technological) differences between markets rather than lack of competition, then the trade-offs associated with firm-side wage policies become more adverse. This problem is highlighted by the finding that across labor markets, a higher firm-related component of pay is unrelated to profitability and positively correlated with the perceived degree of market competition. For a policy targeting firm rents or rent-sharing to reduce wage sorting, it will have to target high-paying, skill-intensive sectors; but if these sectors are high-paying not as a result of greater market imperfections but due, for example, to greater economies of scale, then there may be strong implications for efficiency. This further highlights the need for a comprehensive understanding of wage sorting and the factors that influence firm wage policies.

Chapter 2

Skill-Bias, Firm-Bias, and Wage Inequality

2.1 Introduction

Rising OECD wage inequality is widely attributed to changes in the relative demand for different types of labor. Technological change is thought to have complemented knowledge-intensive occupations, increasing the demand for skilled workers, while a structural shift away from manufacturing has reduced demand for goods-production jobs, and for the less-skilled parts of the labor force traditionally employed in industry.²⁶ The effect of a *skill-biased* demand shock in a competitive labor market is widely-studied, and well-known: an increase in the relative wage paid to skill workers, and a general rise in wage inequality. This explanation for observed wage trends is straightforward and intuitively compelling, but challenged by a large body of evidence indicating that labor markets are not perfectly competitive. Different employers are found to pay different wages to similar workers - a phenomenon widely attributed to frictional labor markets and differences in firm or match rents.²⁷ Employer wage differentials are found, moreover, to be related to occupational and industry wage gaps, an important example being the manufacturing wage premium observed in many developed countries.²⁸ The presence of large, firm-specific wage premia implies that demand shocks may be *firm-biased* as well as skill-biased. They may affect employer composition and hence the distribution of firm premia; and they may change the

²⁶See for example Krusell et al. (2000), Costinot and Vogel (2010), Acemoglu and Autor (2011), Autor and Dorn (2013), Autor, Dorn, and Hanson (2013), Dauth, Findeisen, and Suedekum (2014), Frey and Osborne (2017), Burstein and Vogel (2017), and Acemoglu and Restrepo (2018)

²⁷Card, Cardoso, Heining, and Kline (2018) survey this literature. A typical estimate is that 15% of observed wage variance is due to employer wage differentials.

²⁸See Abowd, Kramarz, Lengermann, McKinney, and Roux (2012), Akerman, Helpman, Itskhoki, Muendler, and Redding (2013), Card, Heining, and Kline (2013), and Torres, Portugal, Addison, and Guimaraes (2018).

incidence of firm premia across the wage distribution, and the extent to which labor market rents are captured by skilled workers.

This chapter is a theoretical and quantitative investigation of how skill-bias and firm-bias jointly determine the effect of demand shocks on the wage distribution. I develop a model with (1) explicit notions of industry and occupation demand, (2) labor supply formalized as an assignment problem, in which differentially skilled workers sort into occupations, (3) firms that earn positive flow rents due to up-front hiring costs, and (4) search frictions that result in rent-sharing between firms and workers. Skill premia and firm premia are equilibrium outcomes, but are also *inputs* to the hiring decisions of firms and the job search behavior of workers. Despite allowing for rich interactions between prices and quantities, the model is both theoretically and empirically tractable. There exist closed-form wage and policy functions that map directly into wage effects from an AKM regression,²⁹ allowing the key distributional parameters to be non-parametrically identified. I estimate the model using German matched employer-employee data that spans the period 1993-2017, and I conduct counterfactual experiments addressing two key questions of interest. First, what is the historical role of firm-bias in explaining German wage trends? And second, what are the implications of firm-bias for policies that affect the wages firms pay?

I find that industry and occupation demand account for two-thirds of the rise in West German wage variance, of which at least one-half is explained by firm-bias and interactions between firm-bias and skill-bias. To obtain this result I conduct counterfactual experiments in which firm-bias is “shut down” by equalizing the match rents earned by firms. Firm-bias is especially important for industry demand, which I estimate to explain one-sixth of the overall trend but which would have exerted a small, *negative* impact on wage inequality if firm premia were homogeneous. Disaggregating further, I find that the effect of firm-bias varies substantially between individual industries and occupations, and I show that conditioning on worker skill does not allow one to predict the direction of the effect on wage variance; an increase in demand for low-skill jobs can increase wage inequality if it occurs in low-paying sectors, while demand for skilled jobs can have the opposite effect. I show in addition that the variability introduced by firm-bias can help to account for regional trends

²⁹After [Abowd, Kramarz, and Margolis \(AKM, 1999\)](#).

in wage dispersion, and that differences in the distribution of firm premia across industries and occupations can substantially explain the smaller increase in wage variance observed in East Germany.

Finally, I consider the joint implications of skill-bias and firm-bias for policies that target firm premia. I conduct two experiments. The first, motivated by proposals to extend collective bargaining coverage to German temp agency workers, involves eliminating the wage gap between temp agencies and the firms that employ temp agency labor. In the second experiment I consider how wage inequality is affected by a compression of firm entry costs, rationalized as the outcome of an exogenous reduction in anti-competitive barriers to entry. I show that in equilibrium, demand and supply responses largely offset, but on net tend to dampen the effects of such policies on wage inequality. The reason is that firm-side wage policies are firm-biased by design, and when firm premia and skill premia are correlated, these policies will also be skill-biased.

The main contribution of this chapter is to the macroeconomic literature on wage inequality, and the contributing role of skill-biased demand shocks due to technological change and trade. Models of skill-bias commonly assume frictionless labor markets: for example [Krusell et al. \(2000\)](#), [Acemoglu and Autor \(2011\)](#), and [Acemoglu and Restrepo \(2018\)](#) on technological change, and [Costinot and Vogel \(2010\)](#), [Autor and Dorn \(2013\)](#), and [Burstein and Vogel \(2017\)](#) in relation to trade. Studies that allow for imperfectly competitive labor markets, such as [Helpman, Itskhoki, and Redding \(2010\)](#), abstract from relative labor demand and the price of skill in order to focus on residual wage inequality within narrow labor markets. In this chapter I consider interactions between the demand for skill and firm premia, which with few exceptions (discussed below) have not been studied in the literature. In doing so I combine and generalize the occupational framework of [Acemoglu and Restrepo \(2018\)](#) and [Autor and Dorn's \(2013\)](#) model of industry demand, and extend the skill-bias paradigm to the case of imperfectly competitive labor markets with search frictions. In this environment there may be important interactions between occupational choice and employer search, and wage premia may influence the job search behavior of workers and the hiring decisions of firms.

A second key contribution is the development of an empirically tractable equilibrium

model of the wage distribution. This is difficult to achieve because neither skill nor firm premia are observable, and equilibrium models with imperfectly competitive labor markets do not, in general, map into empirical wage effects. Type distributions must be parameterized, resulting in loss of information about the underlying wage and employment structures that constitute the key objects of study. I achieve tractability through two strong but reasonable assumptions. First, I assume that labor search is directed across industries and occupations, and not random as is commonly studied in models of worker-firm sorting. Under an efficiency restriction on the unemployment insurance payout,³⁰ directed search results in a closed-form, log-additive wage function. Second, I assume that workers choose their occupations, formalized as an assignment problem building on [Costinot and Vogel \(2010\)](#). Assignment yields skill premia that are monotonically increasing in type, allowing worker skill to be identified from empirical wage effects in a time-consistent manner. Although tractable, the model allows for rich interactions between prices and quantities. Workers may seek out high-paying firms and occupations associated with high-paying firms, though I allow for the possibility of offsetting non-pecuniary amenities.

This approach contrasts with past studies relying on structural models. Wage accounting frameworks such as [Feenstra and Hanson \(1999\)](#), [Goos, Manning, and Salomons \(2014\)](#), and [Lee and Wolpin \(2010\)](#) are highly stylized, and are either out-of-equilibrium or limited to a small set of agent types. The equilibrium model I develop allows for an arbitrary number of industries and a continuum of skill types and occupations. Closer to this study are papers by [Card, Cardoso, Heining, and Kline \(2018\)](#), who develop a reduced-form model in which demand is differentiated across skill types and firms face upward-sloping supply curves; and [Haanwinckel \(2021\)](#), who extends the Card et al. framework to an equilibrium setting. This approach yields a separable wage function when firms have homogeneous demand for skill, but becomes analytically and empirically intractable when skill demand is heterogeneous. The model developed here retains wage separability under arbitrary patterns of labor demand, and is therefore uniquely suited to the quantitative study of interactions between skill demand and firm premia. Finally, while a number of papers have used occupational

³⁰Directed search may be inefficient when (1) the environment is dynamic and (2) submarkets produce differentiated outputs. A dynamic setting will tend to create congestion in particular submarkets, which leads to inefficiency when submarket output is not perfectly substitutable.

assignment as a framework for modeling skill premia, this is one of the first to structurally estimate such a model. The only other study to do so - [Ales, Kurnaz, and Sleet \(2015\)](#) - imposes parametric assumptions on match production, whereas the structural estimation method pursued in this chapter can replicate fine, local features of the wage and employment distributions.

Finally, this chapter contributes to our knowledge of the empirical and theoretical sources of wage dispersion, and bridges the gap between two empirical literatures. On the one hand are studies that attribute wage trends to skill-biased demand, such as [Acemoglu and Autor \(2011\)](#), [Autor and Dorn \(2013\)](#), [Hummels et al. \(2014\)](#), [Autor, Dorn, and Hanson \(2013\)](#), [Ebenstein et al. \(2014\)](#), [Frey and Osborne \(2017\)](#), and [Acemoglu and Restrepo \(2020\)](#). These efforts focus on industry- and occupation-level outcomes, and the role of market-clearing prices in explaining an overall widening of skill premia. Separately, a number of papers employ the wage decomposition proposed by [Abowd, Kramarz, and Margolis \(AKM, 1999\)](#) to study, in a reduced form setting, the mechanical sources of rising wage dispersion. In a closely related study, [Card, Heining, and Kline \(2013\)](#) estimate that one-fourth of the post-1980's rise in West German wage variance is due to greater dispersion of firm AKM wage effects, and one-third to increased covariance of the person and firm effects.³¹ I show that the different sources of wage dispersion studied in these papers are not independent, but have interacted over time. The results I present suggest that the aggregate effect of a skill-biased demand shock is sensitive to the distribution of firm premia across industries and occupations, whereas the presence of equilibrium interactions between supply, demand, and wage premia indicates that reduced-form models of the wage distribution will not be able to disentangle wage dispersion related to market-clearing prices, from that associated with firm premia.

The outline of this chapter is as follows. After a brief motivating discussion, I develop the equilibrium model in section 2. Structural estimation is described in depth in section 3, as this is a key contribution of the chapter. Quantitative experiments and results are presented in section 4.

³¹See also [Song et al. \(2019\)](#), [Bagger, Sorensen, and Vejlin \(2013\)](#), and [Hakanson, Lindqvist, and Vlachos \(2021\)](#) for similar results in other countries, and [Torres et al. \(2018\)](#) and [Alvarez, Benguria, Engbom, and Moser \(2018\)](#) for cases in which firm wage effects are associated with a decline in wage dispersion.

2.1.1 Motivation

The motivation for this chapter comes from an empirical literature on the sources of wage dispersion, that builds on the regression approach of [Abowd, Kramarz, and Margolis \(AKM, 1999\)](#) to study the contribution of firm heterogeneity to wage dispersion. The AKM framework assumes that wage is the product of a person-specific and an employer-specific term, each capturing both observable and unobservable agent heterogeneity. These effects are estimated from matched employer-employee data, and identified from wage changes associated with transitions between employers. Variants of this basic approach have been applied in a variety of contexts and countries. The general finding is that firm premia constitute a small but important source of wage dispersion, directly accounting for between 10% and 20% of wage variance ([Card, Cardoso, Heining, and Kline, 2018](#)). As discussed in the introduction, a number of studies find that the contribution of firm premia has grown over time, most notably due to a strengthening pattern of association between high-earning workers and high-paying firms.

In contrast, the literature on skill-bias has focused almost exclusively on worker heterogeneity in explaining wage trends, with strong emphasis on changes to labor demand at the industry and occupation level.³² This approach is justified if firm premia are entirely a within-industry, within-occupation phenomenon, but past research suggests this is not the case. Inter-industry wage gaps have been found to reflect not only differences in worker ability, but in firm pay as well.³³ The role of firm heterogeneity in driving occupational wage gaps is not as well-studied, but in their analysis of German wage dispersion [Card, Heining, and Kline \(2013\)](#) find large inter-occupation differences in firm premia, while [Torres, Portugal, Addison, and Guimaraes \(2018\)](#) estimate an AKM regression with occupation fixed effects on Portuguese data, and they find these effects to strongly covary with firm premia. When industry and occupation capture wage-related heterogeneity on both sides of the labor

³²See for example [Goos and Manning \(2007\)](#), [Acemoglu and Autor \(2011\)](#), [Autor and Dorn \(2013\)](#), [Autor, Dorn, and Hanson \(2013\)](#), [Goos, Manning, and Salomons \(2014\)](#), [Michaels, Natraj, and Van Reenen \(2014\)](#), [Frey and Osborne \(2017\)](#), [Autor and Salomons \(2018\)](#), and [Acemoglu and Restrepo \(2020\)](#).

³³See [Gibbons and Katz \(1992\)](#) for an early study, and [Abowd, Kramarz, Lengeremann, McKinney, and Roux \(2012\)](#) for more recent work. The important of firm premia varies by country; [Goux and Maurin \(1999\)](#) found them to be unimportant in the case of France, but for Germany [Card, Heining, and Kline \(2013\)](#) found that the variance of the between-industry component of firm premia is large and growing over time.

market, changes to labor demand are likely to affect not just skill premia, but also the firm premia that workers of a given type receive.

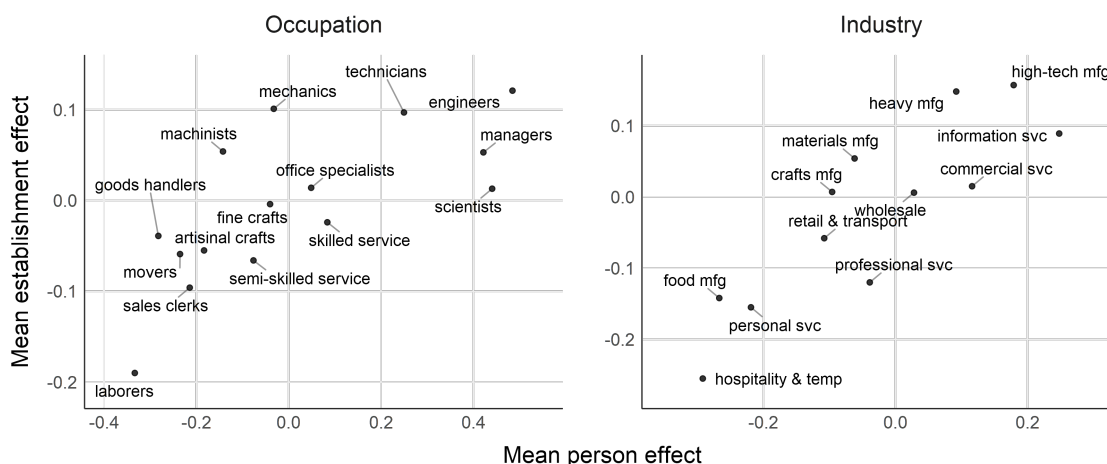


Figure 2.1: Wage Effects by Occupation and Industry, 2010-2017

SOURCE: German linked employer-employee dataset (LIAB). NOTE: Weighted average of AKM wage effects by aggregated KLDB 1988 occupation and WZ 2008 industry. Wage effects estimated following specification of [Card, Heining, and Kline \(2013\)](#).

That this consideration is relevant for West German wages can be seen in figure 2.1, which plots occupation- and industry-mean wage effects over the years 2010-2017. There are substantial differences in firm premia across segments of the labor market. Considered jointly, industry and occupation are as good as explaining firm premia as they are skill premia - they account for one-third of either variance - indicating that firm heterogeneity is no less important at the industry and occupation level than it is in the aggregate. Moreover, it is evident from the figure that these two sources of wage dispersion - person and employer - are not independent across occupations and industries, but are strongly correlated. Skilled occupations tend to be concentrated in high-paying firms, while high-paying industries employ a high-earning (and presumably highly-skilled) workforce.

The implication of the patterns observed in figure 2.1 is that changes to labor demand can affect wage dispersion in one (or all) of three ways.³⁴ First, through the well-known mechanism of skill-bias, labor demand can impact the relative wages earned by skilled work-

³⁴This assumes that firm premia are the result of environmental characteristics unrelated to labor demand, for example differences in employer scale together with monopsonistic wage-setting. This is a reasonable assumption given uncertainty as to the underlying mechanism generating firm pay gaps.

ers. Second, labor demand can disproportionately or disparately affect high-paying and low-paying firms, thereby changing the distribution of labor across firms and hence the distribution of firm premia across job matches. I refer to this mechanism as firm-bias. Finally, if there is sorting of skilled labor into high-paying firms, skill-bias and firm-bias may interact. Changes to skill premia may be concentrated in high-paying firms, thereby having a greater affect on wage variance than they otherwise would have; and shifts in the composition of employers can affect the incidence of firm premia across the wage distribution, and the extent to which high-skill workers also benefit from a high-paying employer. The third channel in particular depends on the equilibrium responses of wages and quantities, and is the primary motivation for the model developed below.

2.2 A Model of Firm-Bias and Skill-Bias

In this section I develop a model that formalizes, in a tractable way, the linkages between labor demand, skill premia, and firm premia. Labor demand is differentiated by industry and occupation, generalizing an approach taken previously by [Acemoglu and Autor \(2011\)](#) and [Autor and Dorn \(2013\)](#). Firm premia arise due to the presence of search frictions, which create a rent-sharing motive and a wage gap between high-rent and low-rent firms. Skill premia are modeled as the result of an occupational assignment problem following [Costinot and Vogel \(2010\)](#), adapted here to the case where firms are heterogeneous and labor markets are frictional. A key prediction of the model is that under a set of restrictions that may be interpreted as efficiency conditions, the wage function is separable in worker and firm type, providing a rationale and a microfoundation for the common empirical assumption of log-additivity. Skill premia and firm premia are not, however, independent; the model predicts a set of potentially important interactions between supply, demand, and wage premia, indicating that an equilibrium framework is necessary for disentangling the roles of skill- and firm-biased changes to demand. One implication of this, explored further in the quantitative portion of the chapter, is that wage policies targeting firm premia will in general have implications for skill premia, suggesting that the distributional effects of such policies are not straightforward. I close the section with a set of numerical comparative statics that illustrate the equilibrium properties of the model.

2.2.1 Environment

The economy is set in continuous time. There exists a continuum of infinitely-lived workers, heterogeneous in a scalar skill variable $s \in [\underline{s}, \bar{s}]$ whose distribution is given by the continuously differentiable and strictly increasing function $\nu(s)$. Firms are endogenous in measure and exist in two types. *Employers* are single-vacancy firms that hire workers to produce an (i, j) -specific labor output, where $i \in \{1, \dots, I\}$ is industry and $j \in [0, 1]$ denotes occupation.³⁵ *Industry aggregators* (or simply *aggregators*) combine labor outputs within an industry to produce an i -specific final good, which is then sold to, and consumed by, a representative household. All agents are risk-neutral and discount the future at rate ρ . As I will focus on the steady-state equilibrium, time subscripts are suppressed in this section.

Employers. An employer enters the market by paying a vacancy flow cost $C(i, j)$ and posting a wage offer w , with entry assumed to be otherwise free. Vacancies are filled at a rate $q(\theta(i, j))$ with θ denoting market tightness. An (i, j) vacancy filled by an s -worker generates match output $m(j, s)$, which may be sold to i -aggregators at a unit price of $p(i, j)$. To preclude the screening of workers by firms, I assume that the wage offer w is denominated in units of match output, so that the amount paid to a hired worker is $wm(j, s)$.³⁶ Existing matches dissolve at an exogenous rate δ that is the same across types.³⁷ I assume that match output is increasing in skill, and that skilled workers are relatively better at high- j (i.e. *skill-intensive*) occupations:³⁸

³⁵The assumption of a discrete number of industries is empirically motivated, and may be relaxed without affecting the results shown below.

³⁶With search frictions, the worker-optimal and firm-optimal occupational assignments are generally different. This assumption imposes the worker-optimal assignment, which I favor for its intuitive appeal. From a quantitative standpoint there is no loss of generality, as the two approaches yield different parameter estimates but identical model relationships.

³⁷My rationale for this assumption is that I am not attempting to replicate empirical worker flows. While it will be important that the model captures the response of labor supply to posted wages, this requires only a single margin of worker choice (job search), and it is therefore convenient to assume that δ is exogenous and to account for variation in separation rates when estimating the empirical elasticity of labor supply. A similar point may be made regarding the absence of on-the-job (OTJ) search, which has implications primarily for worker flows and not wages given that wages in this environment are not the outcome of bargaining.

³⁸Comparative advantage is required for uniqueness of the optimal assignment, while the ranking of j is simply a normalization. That skill possesses an absolute advantage is intuitive in this context, and a similar assumption is made by [Teulings \(1995\)](#). The optimal assignment is simplified by having m independent of i , and there is little loss of generality given that skill is one-dimensional.

Assumption 1 (*Match productivity*): $m(j, s)$ is a continuously differentiable function, and for any j and s we have $\frac{\partial}{\partial s} m(j, s) > 0$ and $\frac{\partial^2}{\partial j \partial s} \log m(j, s) > 0$.

The problem facing an employer is therefore to choose the wage that maximizes the expected value of a vacancy posting, taking as given equilibrium prices and market tightness:

$$\begin{aligned} \max_w S^V(i, j, w) \\ \text{s.t. } \rho S^V(i, j, w) &= -C(i, j) + q(\theta(i, j)) [\mathbb{E}_s S^J(i, j, s, w) - S^V(i, j, w)] \\ \rho S^J(i, j, s, w) &= m(j, s)(p(i, j) - w) - \delta S^J(i, j, s, w) . \end{aligned} \quad (3)$$

Here S^V is the net present value of a posted vacancy and S^J the NPV of a filled vacancy. Worker skill is an expectation as the firm does not choose s and it is *ex ante* possible that multiple worker types could apply to submarket (i, j) , though in equilibrium this will not be the case.

Workers. At any point in time workers are either employed or unemployed. Employed workers earn the wage $w \times m$ and realize an industry-specific, non-pecuniary amenity $A(i)$, for a total flow value from employment of $u(wm, A)$. Unemployed workers become employed by directing their search across occupations (j) and industries (i). While searching for a job they receive an insurance payment $B(u, \theta)$ that is allowed to depend on market tightness and the flow utility of employment, subject to the constraint that any match must be accepted (i.e. there is a perfectly enforced work-search requirement).³⁹ Jobs are found at a rate $f(\theta(i, j))$. I focus on the symmetric equilibrium in which workers choose a mixed strategy $\phi(i, j, s)$ over submarkets that solves the problem:

$$\begin{aligned} \max_{\phi(i, j)} \sum_i \phi(i, j) S^U(i, j, s) \\ \text{s.t. } \rho S^U(i, j, s) &= B(u(i, s, j), \theta(i, j)) + f(\theta(i, j)) [S^W(i, j, s) - S^U(i, j, s)] \end{aligned}$$

³⁹The motivation for $B(u, \theta)$ is discussed in more detail later in this section; for now I simply give the intuition that because employment flow payoffs, job-finding rates, and B will jointly determine workers' job-seeking behavior, a policymaker will optimally take u and θ into account when setting unemployment benefits.

$$\begin{aligned} \rho S^W(i, j, s) &= u(w(i, j)m(j, s), A(i)) + \delta [\max_{i', j'; s} S^U(i', j', s) - S^W(i, j, s)] \\ \sum_i \int \phi(i, j) dj &= 1, \end{aligned} \quad (4)$$

where the terms S^U and S^W are the net present values of job search and a matched job, respectively.⁴⁰

Aggregators and household. Match output is sold to industry aggregators as an intermediate good $y(i, j)$. Aggregators then produce an industry-specific final good, using a constant returns CES technology.⁴¹ All industries face the same elasticity of substitution $\sigma > 0$ across j -goods, but occupational shares $\alpha(i, j)$ are allowed to be industry-specific. Final goods are sold to the representative household at a price $P(i)$. Aggregator profits will be maximized by the bundle of intermediate goods that solves the problem

$$\max_{y(j)} \left(\int \alpha(i, j) y(j)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}} - \int p(i, j) y(j) dj. \quad (5)$$

Here $\int \alpha(i, j) dj = 1$, and I assume that α is continuous in j . The household is assumed to have CES preferences with elasticity $\tau > 0$, and will maximize utility by solving

$$\max_{Y(i)} \left(\sum_i \beta(i) Y(i)^{\frac{\tau-1}{\tau}} \right)^{\frac{\tau}{\tau-1}} - \sum_i P(i) Y(i). \quad (6)$$

I abstract from the household's budget constraint as it affects neither relative prices nor quantities, but note in passing that the model may be easily closed by repatriating all wages and profits to the household. For similar reasons I omit the government's budget constraint, with unemployment payouts B assumed to take the form of an endowment.

⁴⁰Although the supply of skill is assumed to be exogenous, it may easily be made endogenous to reflect e.g. the influence of wages on labor market participation or human capital investment. The main challenge in doing so is empirical: because wages are generally used to identify worker skill - a latent variable - any causal relationship between wages and skill is thereby lost.

⁴¹For simplicity I omit capital from production, though the entry costs $C(i, j)$ may be thought of as reflecting capital or any other fixed costs incurred prior to job creation, in addition to competitive barriers to entry.

2.2.2 Equilibrium

Markets. Goods markets are perfectly competitive. Free entry of employers implies that the value of a vacancy is non-positive:

$$\rho S^V(i, j) \leq 0, \quad (7)$$

with equality in any submarket for which entry is positive. Subject to perfect competition, aggregators earn no profits.

Labor markets are frictional in the sense that, at any point in time, some vacancies go unfilled and some workers fail to find an employer. Given N unemployed workers and V vacant employers, matches will form at a rate $M(N, V)$. For tractability reasons I assume this function to be Cobb Douglas:

Assumption 2 (*Matching function*): $M(N, V) = \zeta V^\eta N^{1-\eta}$.

With Cobb Douglas matching the vacancy-filling and job-finding rates will be $q(\theta) = \zeta \theta^{\eta-1}$ and $f(\theta) = \zeta \theta^\eta$, respectively.

The price of the final good may be normalized, and a convenient normalization is $P = \delta/Y$ where Y is household utility (or equivalently, “aggregate output”). I maintain this assumption in what follows.

Equilibrium. In equilibrium, market tightness must be consistent with the optimal job search behavior of workers and employer vacancy policies. Defining $N(s)$ to be the unemployed mass of s -workers and $V(i, j)$ the total number of (i, j) -vacancies in steady-state, we require that

$$\theta(i, j) = \frac{V(i, j)}{\int \phi^*(i, j, s) N(s) ds}. \quad (8)$$

It must also be that the markets for intermediate and final goods each clear:

$$y^*(i, j) = \frac{q(\theta(i, j)) V(i, j)}{\delta} \times \frac{\int N(s) \phi^*(i, j, s) m(j, s) ds}{\int N(s) \phi^*(i, j, s) ds} \quad (9)$$

$$Y^*(i) = \left(\int \alpha(i, j) y^*(i, j)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}, \quad (10)$$

where the two terms on the right-hand side of (9) correspond to total employment multiplied by average labor productivity in submarket (i, j) . Finally, in steady-state we must have parity between flows into and out of unemployment:

$$q(\theta(i, j))V(i, j) = \delta(\nu(s) - N(s)). \quad (11)$$

A steady-state equilibrium in this economy is employer wage posting decisions $w^*(i, j)$ that solve problem (21), search behavior $\phi^*(i, j, s)$ solving (19), and output quantities $y^*(i, j)$ and $Y^*(i)$ that solve problems (5)-(6), together with market tightness $\theta(i, j)$, unemployment and vacancy distributions $N(s)$ and $V(i, j)$, and prices $p(i, j)$ and $P(i)$ such that conditions (7) and (8)-(11) hold.

2.2.3 The Restricted Model

Further characterization of the model requires that assumptions be placed on the worker flow payoffs u and B . These are important because they determine the relative values that workers place on the different forms of “payment” associated with a given submarket: wages, amenities, and job-finding rates. In this section I impose conditions on u and B such that worker preferences are homogeneous. Homogeneity is an attractive property in this case for several reasons. First, it allows the equilibrium to be described in closed-form. Second, it is a natural assumption in this environment, insofar as it is a necessary condition for the allocative efficiency usually associated with directed search. And third, homogeneous preferences are a sufficient condition for the model to generate a separable wage function - which provides both a direct path to empirical identification of the model primitives, and a microfoundation for the common empirical assumption of a log-additive wage function.⁴²

⁴²The AKM wage regression assumes log-additivity, but many authors have observed that separability is not consistent with random search models - see for example the criticism by [Lopes de Melo \(2018\)](#). The monopsony framework of [Card, Cardoso, Heining, and Kline \(2018\)](#) does generate a separable wage function, but only when firms have identical technical demand for skill; separability is lost when there is sorting as in [Haanwinckel \(2021\)](#). The key advantage of wage separability is that it allows agent types and model primitives to be identified separately, which in turn allows for a non-parametric estimation approach as employed in this chapter.

Restrictions. The conditions to be imposed are that (1) flow utility u is a homogeneous function of wages, and (2) unemployment payoffs B are set in such a way that the value of labor search (i.e. S^U) is a homogeneous function of u . I implement these conditions through the following assumptions:

Assumption 3 (*Flow payoffs*): *The flow payoffs of employment and unemployment take the functional forms*

$$u(\omega, A) = \omega^\psi A^{1-\psi} \quad (12)$$

$$B(u, \theta) = \frac{(\rho + \delta + f(\theta))u^{\kappa-1} - 1}{\rho + \delta} f(\theta)u . \quad (13)$$

where $\psi \in (0, 1]$ and $\kappa \geq 0$.

The restriction on u lacks an intuitive interpretation, insofar as amenities are exogenous and play a structural rather than a theoretical role in this chapter.⁴³ Condition (13), on the other hand, says that workers with a high job-finding rate are incentivized to place a greater value on additional increases in market tightness, since due to the dynamic nature of the environment they would otherwise be less sensitive to changes in the job-finding rate.

Interpretation of B. The restriction on B has an efficiency interpretation: it is a necessary condition for an efficient allocation of labor across industries and occupations.⁴⁴ Perhaps surprisingly, directed search in this environment is not generally efficient; inefficiency arises due to the fact that with differentiated output and labor markets, labor demand cannot be perfectly substituted across submarkets. When preferences are non-homogeneous, the effective “price” of skill will vary across industries, which distorts labor shares and reduces output due to the quasi-concavity of the aggregator’s production function. Condition (13)

⁴³It is straightforward also to include occupational amenities, though characterization of the optimal assignment becomes problematic unless these are continuously differentiable.

⁴⁴In the appendix I provide an informal proof for the simplified case in which occupational assignment is exogenous and amenities are homogeneous. Efficiency is not key to the results in this chapter, and is discussed here for purpose of providing intuition for condition (13).

does not eliminate this congestion externality, but constrains it to be independent of skill. Workers may still “over-apply” to industries offering a high wage ($\kappa > 1$), or to low-paying firms offering a higher job-finding rate ($\kappa < 1$), but all workers will behave similarly regardless of s . When $\kappa = 1$ the externality is removed altogether, and this can be interpreted as the special case in which the efficiency properties of directed search are preserved under differentiated labor and product markets.

Occupational assignment. An immediate implication of restrictions (12) and (13) is that the assignment of workers to occupations will be independent of industry. Following Costinot and Vogel (2010) in defining a correspondence $\lambda : i, s \rightarrow j$ that gives the set of jobs j chosen by s -workers in industry i , and provided that $\alpha(i, j)$ and $C(i, j)$ are continuous, we can show the following:

Proposition 1. *Under assumptions 1-3, $\lambda(s)$ is a strictly increasing, differentiable function independent of i , that satisfies the system of equations*

$$\frac{dU(s)}{ds} = \kappa\psi \frac{m_s(\lambda(s), s)}{m(\lambda(s), s)} U(s) \quad (14)$$

$$\frac{d\lambda(s)}{ds} = \frac{m(\lambda(s), s)N(s)U(s)}{\sum_i \left[\left(\frac{\kappa\psi(1-\eta)}{\eta+\kappa\psi(1-\eta)} p(i, \lambda(s)) m(\lambda(s), s) \right)^\psi A(i)^{(1-\psi)} \right]^\kappa \left(\frac{\alpha(i, \lambda(s))}{p(i, \lambda(s))} \right)^\sigma \frac{\beta^\tau}{P(i)^{\tau-\sigma}}} \quad (15)$$

where $\lambda(\underline{s}) = 0$ and $\lambda(\bar{s}) = 1$, $P(i)$ is the industry price index, $p(i, j)$ is the price of output in submarket (i, j) , and $U(s) \equiv \rho \max_{i,j} S^U(i, j, s)$ is the reservation value of type s .

Equation (14) states that at any point, the wage function and the match production function increase at the same rate in s , because only when this is true will workers not realize a larger payoff by moving to a higher or a lower k . Equation (15) describes the sorting of workers across occupations. Higher-skilled workers are assigned to higher- j jobs, at which they have a comparative advantage; and this assignment is constant across industries. Note that assignment depends not just on labor demand (through y), but also on wages (through p) and industry amenities.

2.2.4 Demand, Supply, and Wages

The model developed thus far makes two important predictions: that the wage function is separable, and that wage premia have behavioral implications. In this section I briefly discuss both predictions. I then illustrate, by way of a numerical exercise, how the skill-bias and firm-bias of a demand shock jointly determine the effect on wages.

Wage separability. As a consequence of conditions (12) and (13), the equilibrium wage function can be shown to be a multiplicative function of a *firm premium* that depends on i and j , and a worker or *skill premium* that depends on s :

$$w^*(i, \lambda(s)) = \underbrace{\left(\frac{\kappa\psi(1-\eta)(\rho+\delta)C(i, \lambda(s))}{\eta\zeta^{\frac{1}{\eta}}A(i)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}}} \right)^{\frac{\eta}{\eta+\kappa\psi(1-\eta)}}}_{\text{firm premium}} \underbrace{U(s)^{\frac{1-\eta}{\eta+\kappa\psi(1-\eta)}}}_{\text{skill premium}} \quad (16)$$

$$\equiv FP(i, \lambda(s))SP(s) .$$

The firm premium is a function of entry costs and amenities; because these increase or reduce the cost of hiring a worker, they affect firm entry, and hence output prices and wages as well. The skill premium is a function of the worker's outside option U , which is determined in equilibrium and reflects the demand for different types of jobs as well as the substitutability of skill types across these jobs.

Behavioral response to wage premia. The supply of labor to a given submarket will depend on the matching function - for which there is no closed-form solution - and on the search probability ϕ :

$$\phi^*(i, j, s) = \frac{V(i, \lambda(s)) \left(FP(i, \lambda(s))^\psi A(i)^{1-\psi} \right)^{\frac{\kappa}{\eta}}}{\sum_k V(k, \lambda(s)) \left(FP(k, \lambda(s))^\psi A(k)^{1-\psi} \right)^{\frac{\kappa}{\eta}}} . \quad (17)$$

Equation (17) states that workers are more likely to apply to a submarket when there are more vacancies, and when firms in this submarket are higher-paying or have relatively larger

amenities. Less obvious is that, through the matching function (15), workers will also tend to seek out high-paying *occupations*. Hence labor demand V and the distribution of firm premia will jointly determine occupational assignment and therefore distribution of skill premia SP . A key role here is played by the term κ/η , which gives the λ -conditional elasticity of labor supply with respect to the flow payoff $FP^\psi A^{1-\psi}$. When κ is large, workers will respond more strongly to firm-specific payoffs, and when $\kappa = 0$ they will not respond at all. Similarly, a low value of η means that the job-finding rate is unresponsive to vacancies, and as a result workers will respond relatively less to V and relatively more to FP and A .

The natural measure of labor demand in this environment is the mass of vacancies in submarket (i, j) , which is given by the equation

$$\begin{aligned}
V(i, \lambda(s)) &= \alpha(i, \lambda(s))^\sigma \beta(i)^\tau \frac{m(\lambda(s), s)^{\sigma-1}}{A(i)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}} FP(i, \lambda(s))^{\frac{\eta\tau + \kappa\psi(1-\eta)}{\eta}} SP(s)^{\sigma-1}} \\
&\times \left(\int \alpha(i, \lambda(k))^\sigma \left[\frac{m(\lambda(k), k)}{SP(k)} \right]^{\sigma-1} dk \right)^{\frac{\tau-\sigma}{\sigma-1}} \bar{V}, \tag{18}
\end{aligned}$$

where \bar{V} is a function of aggregate parameters. Aside from the technical demand parameters α and β , vacancies will also depend on the industry amenity and the firm premium, which both tend to increase the vacancy-filling rate and hence to reduce output prices given V . Two substitution effects are also possible: across occupations within an industry whenever $\sigma \neq 1$, and across industries when $\tau \neq \sigma$. Both substitution effects depend inversely on the term SP/m , or the unit labor cost associated with s -workers. This ratio is the outcome of equation (14), and will only be constant in s if $\kappa = 1/\psi$. In other words changes to skill premia will generally exert an equilibrium effect on demand, though the direction and magnitude of that effect will depend critically on the demand elasticities σ and τ .

Comparative statics. To better illustrate the model's predictions regarding wages, I consider the effects of a skill-biased demand shock in the context of a simple, two-industry version of the model. I assume that share functions and amenities are initially the same in both industries; that type distributions are uniform, and match production is an exponential

function of s and j ; and that $\kappa = \sigma = \tau = 1$. Parameter values are roughly calibrated to match wage variance moments for West Germany over the period 1993-1999. The demand shock is modeled as a change to occupational demand α_{low} and α_{high} such that the *aggregate* vacancy shares of skilled jobs increases by a fixed amount: that is, for any j' , the mass of vacancies where $j > j'$ is now greater than the mass of vacancies for which $j < j'$ by a proportion held constant across experiments. Whether this change in demand is also *firm-biased* depends upon whether it is driven by changes to α_{low} or α_{high} , with different cases represented in different columns of figure 2.2. The first row of plots gives results as a function of the worker's wage percentile, to show the effects of the shock on the wage distribution; and the second row as a function of the worker's skill percentile, to better illustrate the implications for wages earned by a particular type.

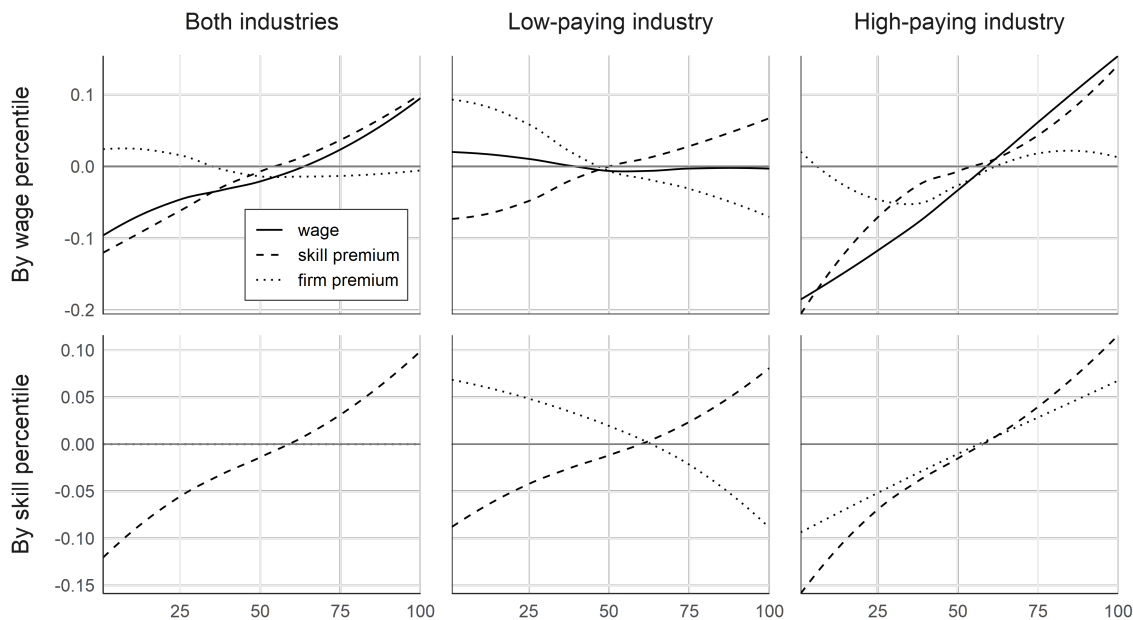


Figure 2.2: Simulated Effects of a Skill-Biased Demand Shock

When the increase in demand for skilled jobs is identical across industries (first column), there is no change to the average firm premium earned by different worker types. Because the optimal occupational assignment will now have each skill type assigned to a higher- j job than before, the return to skill and hence $Var(SP)$ increases, as shown by the upward-sloping dashed line. Critically, however, wage variance $Var(w)$ increases by a smaller amount than $Var(SP)$, reflecting the fact that wage and skill are imperfectly correlated; for this

reason there is incomplete “pass-through” from the distribution of skill premia to the wage distribution. The extent to which pass-through is incomplete will depend on the initial distribution of jobs across industries, and will become greater when high- j occupations are initially more common in the high-paying sector.

The second and third columns show wage outcomes when the skill-biased shock is concentrated in the low-paying and the high-paying industry, respectively. If rising skill demand is driven by the low-paying industry, then skill premia and firm premia become negatively correlated and there are offsetting effects on wages. On the one hand, $Var(SP)$ increases as in the case with no firm-bias. On the other hand, $Cov(SP, FP)$ now declines. In this particular example the sorting effect dominates, and wage variance decreases overall. The third column shows the effects of skill-bias when these two channels are complementary, and skilled workers are sorted into the high-paying industry. We now have $Var(w)$ increasing by a larger amount than $Var(SP)$, with firm-bias amplifying the effects of the demand shock on wages.

By assumption only the supply side responds behaviorally to the change in demand, and this response may be seen by examining the effect on wages by skill percentile (row 2). A given increase in skilled vacancies will have a larger effect on worker search behavior, and hence on occupational assignment and the distribution of skill premia, when this increase occurs in the high-paying industry. This is because, as shown in equation (17), vacancies and firm premia play a similar role in determining workers’ job choice. An increase in the average firm premium associated with a particular occupation acts like increase in vacancies. Hence the “effective” skill-bias of the demand shock will depend upon the industry that it predominantly affects.

2.2.5 Discussion

Summarizing the results shown above, the model suggests four ways in which a change to the environment (such as a demand shock) might impact the wage distribution:

1. **Skill-Bias:** by affecting the occupational assignment of workers.
2. **Firm-Bias:** by affecting relative demand or hiring costs facing employers.

3. **Statistical Interactions:** changes to skill premia that are correlated with firm premia, or vice versa.
4. **Equilibrium Interactions:** behavioral supply and/or demand responses that are themselves skill- or firm-biased.

The first channel is the standard one considered in the literature. The second is non-standard but straightforward: any change that impacts the distribution of labor across firms may also impact the distribution of firm premia across skill types. The third and fourth channels are not straightforward, and arise when occupations are unevenly distributed across industries, or when employer rents are different for skilled and unskilled occupations (which may be true even if occupational demand is homogeneous across industries). In this case, a firm- or skill-biased change to the environment that affects wage premia may also impact the statistical association between skill and firm premia. It may also elicit behavioral responses by workers, firms, or both, for the reasons explained above; and these behavioral responses may themselves be firm- or skill-biased.

While the prediction of a separable wage function is somewhat unique among models of this class, and will be useful in the next section, the predictions regarding agent behavior are of deeper significance. They suggest that in the presence of labor market sorting, skill-bias and firm-bias are not separable in theory or in practice. As the numerical exercises in this section indicate, it is possible for a skill-biased shock to reduce the wage gap between low-earners and high-earners - an outcome not only quantitatively but qualitatively inconsistent with canonical models of skill-bias. At the same time, a strong role for behavioral responses to wage premia is also inconsistent with common empirical experiments, in which wages are decomposed into firm- and worker-specific effects, and the wage distribution is then manipulated along one axis (e.g. firm effects) but held fixed along the other, in order to ask a counterfactual “what if.” Such experiments cannot reliably measure the impact of firm premia on the wage distribution; and they are likely to suffer from greater mis-specification when used to draw inferences about the wage impacts of environmental changes that are inherently firm-biased. These issues are explored in more depth in the quantitative portion of the chapter.

2.3 Structural Estimation

This section describes the structural estimation procedure for the model developed above. I begin with a description of the German matched data used for this exercise. I then describe the non-parametric identification of the latent type distributions, made possible by the log separability of the model-predicted wage function. Next I discuss estimation of the match primitives, followed by the distributional parameters characterizing labor demand, entry costs, and amenities. Lastly I turn to the aggregate parameters and the elasticity terms σ , τ , and κ . The model moments used in estimation are provided in the appendix, along with a characterization of model performance and goodness-of-fit.

2.3.1 Data

German linked employee-employer dataset (LIAB). The principal dataset used in estimation contains administrative and survey data on German establishments and their employees, and covers the years 1993-2017. This data is provided by the Institute for Employment Research (IAB), who each year conduct a large, stratified survey of German private sector establishments.⁴⁵ The survey is then matched with social security records for all employees of record as of the survey date. Key variables in the establishment portion of the dataset include detailed industry codes, current vacancy postings, and cross-sectional weights to correct for stratification.⁴⁶ Establishments are surveyed in waves lasting several years, making it possible to construct measures of establishment-level job flows. Employee data include information on daily earnings, the occupational code assigned by the employer, and basic demographic characteristics including sex and educational attainment. Workers that do not pay social security taxes (e.g. those in marginal employment) are not observed, and as part-time employees are recorded inconsistently over the sample period, I exclude them as well. Following past studies I restrict the sample to individuals aged 20-60, leaving a dataset that consists annually of between .75 and 2.2 million individuals and between 4,000 and 9,000 thousand establishments.

⁴⁵An establishment is defined as a physical workplace, however locations may be aggregated when they share the same corporate ownership, industry classification, and municipal code.

⁴⁶See Chapter 1 for a description of how industry and occupation codes are aggregated.

Provided with the LIAB are updated versions of the person and employer wage effects estimated by [Card, Heining, and Kline \(2013\)](#), following the now-standard approach of [Abowd, Kramarz, and Margolis \(1999\)](#).⁴⁷ These effects are obtained from the panel wage regression

$$w_{i,t} = \pi_i + \phi_{j(i,t)} + x'_{i,t}\beta + \epsilon_{i,t} ,$$

where $w(i, t)$ is the log daily wage of person i in year t , π is a time-invariant wage effect associated with person i , ϕ is a time-invariant effect associated with i 's employer j , and x is a vector containing year fixed-effects and a cubic polynomial in worker age, interacted with dummies for educational attainment. Estimation is performed on the population-level datasets from which the LIAB is extracted, in four partially-overlapping panels spanning 7-8 years. While it is possible to estimate AKM wage effects from the LIAB, I prefer the data provider's estimates for two reasons. First, because wage effects are identified from job transitions, they will be more precisely estimated in datasets that contain the full universe of establishments. Second, because better identification means that wage effects can be estimated for a larger number of establishments, this allows me to obtain finer measurements of the empirical wage distribution without violating confidentiality restrictions.

For consistency with the model, I average both the person and the firm wage effect at the industry-occupation level using detailed (3-digit) codes. As there is substantial and apparently idiosyncratic variation in the extreme tails of the person effect distribution, even after averaging, I filter observations lying in the first or last percentiles. See the appendix to Chapter 1 for additional details on data cleaning.

Aggregate data. All nominal variables are converted to real values using German CPI data from the OECD.⁴⁸ As the LIAB contains data only on employed individuals, I obtain annual employment, unemployment, and vacancy data from the Federal Employment Agency. Numbers are reported separately for West and East Germany. Measures of unemployment and vacancies are not fully consistent over time, however these data inform only two aspects of the model - the matching efficiency ζ and aggregate market tightness (through the overall

⁴⁷See [Bellman \(2020\)](#) for a complete description of this procedure.

⁴⁸Monthly CPI of all items, averaged annually with a base year of 1995.

level of entry costs C) - neither of which is influential with respect to the quantitative results.

2.3.2 Type Distributions

For the 2010-2017 panel I normalize $\lambda(s)$ and $\nu(s)$, which serves to fix s and j in all panels.⁴⁹ In earlier panels, these distributions are estimated as follows. Using equation (14) and the estimated match productivity function m (discussed next), the empirical matching function is backed out from percentiles of AKM person effects. Once the empirical analogue of λ is estimated, other model primitives can then be estimated over j .




Figure 2.3: Estimated Type Distributions

The skill distribution is allowed to vary across panels, as the 1993-2017 period saw migration from East Germany, a general increase in German educational attainment, and a secular rise in the participation rate and fall in the unemployment rate following the Hartz reforms of 2003-2005.⁵⁰ To account as best as possible for these changes, I use AKM person effects, together with the model prediction of a strictly increasing skill premium $SP(s)$, to assign workers to wage percentiles calculated over each panel. I then estimate the annual distributions of labor across the (panel) wage percentiles in order to capture within-panel changes to the distribution of skill. I leverage the years of overlap between panels - 1998-1999, 2003-2004, and 2010 - to map between panels. In this way I am able to estimate the skill distribution ν in earlier panels. I find that each successive panel sees an increase in

⁴⁹These normalizations are WLOG since the match productivity function m is as yet unrestricted. Some care is required when fixing $\lambda(s)$, however, as a poor choice may result in non-monotonic matching functions for earlier panels. In practice the monotonicity condition arising from occupational assignment must be checked and verified, and the normalization for 2010-2017 adjusted if necessary.

⁵⁰See for example [Krueger and Pischke \(1995\)](#) and [Krause and Uhlig \(2012\)](#).

the supply of high-skill labor, consistent with increased workforce education; and following the early 2000's there is an upwards trend in the supply of very low-skill labor, as would be expected from greater participation of marginal workers.

2.3.3 Match Productivity

A functional form is required for the match production function. A natural approach is to assume that m is an exponential function.⁵¹ I generalize this assumption as follows.

Assumption 4 (Functional form of m): $m(j, s) = e^{G(s)\gamma_{sj} + G(j)}$, with G and F time-invariant.

Fixing an initial guess for unemployment distribution $N(s)$, I obtain the empirical analogue of the skill premium function $SP(s)$ from the distribution of AKM person effects in the LIAB. I then fit a quadratic spline and differentiate in order to arrive at $SP'(s)$, which may be substituted into equation (14) to non-parametrically obtain $G'(s)$ for the 2010-2017 period under the normalization $\gamma_{sj}^{2010-17} = 1$. The term $G(s)$ is recovered afterward by integrating.⁵² The occupational productivity multiplier $H(j)$ is also estimated over the 2010-2017 panel, under the restriction $\frac{SP(s)}{m(\lambda(s), s)} = 1$.⁵³ Once G is estimated it is then possible to back out $\lambda(s)$ for panels prior to 2010-2017 as described above. The term γ_{sj} is identified in each panel from the boundary condition $\lambda(\bar{s}) = 1$, and increases over the sample period from .72 in 1993-1999 to the normalized value of 1 in 2010-2017.

Note that all estimates are conditional on $N(s)$, because empirical wages are calculated as percentiles of the *employed* distribution and must be mapped into s . As the unemployment distribution is not initially known but must be solved for, this necessitates an iterative estimation procedure, which I describe in greater detail below.

⁵¹Examples include Teulings (1995) and Ales, Kurnaz, and Sleet (2015).

⁵²To prevent division by 0 when solving equation (14) for G' , it is necessary to add an intercept $G(s)\gamma_s$. I assume that $\gamma_s = .01$.

⁵³The intent behind this assumption is to equalize unit labor costs, the *levels* of which are unconstrained, prior to conducting forward-looking experiments that shift the occupational composition of industries. Note, however, that unit labor costs will deviate from 1 in the earlier panels, and may deviate from 1 in the policy experiments if the occupational assignment is affected.




Figure 2.4: Estimated Match Production Parameters

2.3.4 Demand, Amenities, and Costs

The demand functions α and β are estimated in each panel from industry (wage) cost shares and the intra-industry distributions of AKM person effects, averaged by 3-digit industry and occupation as described previously.⁵⁴ Entry costs are estimated from the inter- and intra-industry distributions of AKM establishment effects, with the level of entry costs calibrated so that market tightness is equal to its empirical counterpart over 1993-2017. I estimate this to be .146 using annual unemployment and vacancy statistics. Industry amenities are identified from industry-mean AKM effects and vacancy-filling rates. The equilibrium equations used in estimation are provided in the appendix.

Estimation is an iterative process for two reasons. First, the unemployment distribution is *ex ante* unknown, but appears in the equilibrium equations used to identify the technical parameters and those governing match production. Second, the coarseness of the empirical distributions, together with the interpolation steps used in estimating the distributional parameters, adversely affects model fit. I therefore begin with an initial guess for unemployment, after which estimation is performed based on this guess, and the model is solved conditional on the parameter estimates.⁵⁵ The guess for $N(s)$ is then updated, the model parameters are re-estimated, α is adjusted to close the gap between the empirical and model-generated matching functions, and the process is repeated until convergence. On top of this an outer loop is performed, in which industry demand and entry costs are adjusted

⁵⁴All intra-industry distributions are measured over 20 quantiles. A finer resolution is not possible given data confidentiality requirements that each statistic represent at least twenty unique establishments.

⁵⁵In particular a numerical solution must be obtained to the system of differential equations (??)-(15).

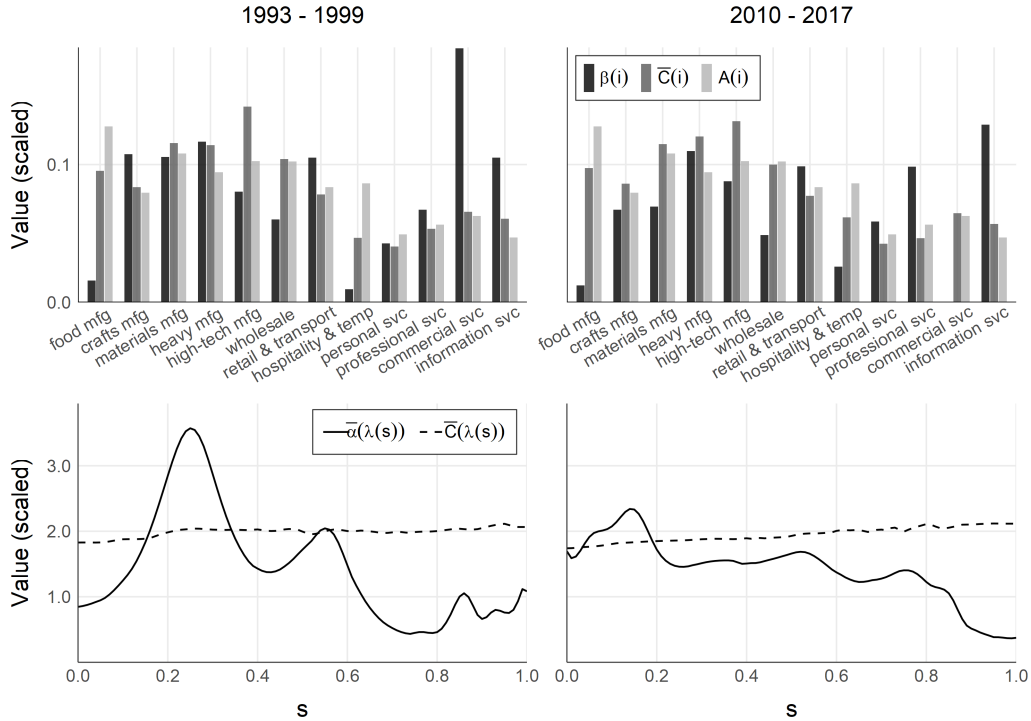


Figure 2.5: Estimated Technical Parameters

NOTE: terms with bars indicate averages across either i or j .

to minimize the distance between empirical and predicted (1) market tightness, (2) industry employment shares, and (3) industry mean firm premia.

2.3.5 Aggregate Parameters

The flow utility exponent ψ is not separately identified from the amenity term A , and I therefore impose the normalization $\psi = \eta/(1 - \eta)$, which has the convenient implication that the elasticity of job applicants with respect to FP will be exactly equal to κ . The discount rate ρ , which serves only as a scaling variable, is assigned a value corresponding to a discount rate of 0.96. Estimation of the matching elasticity η is possible using LIAB vacancies and hires together with aggregate unemployment flows, but in practice this results in an implausibly high value.⁵⁶ I therefore set $\eta = .35$ as estimated for Germany by [Kohlbrecher et al. \(2016\)](#). Match efficiency ζ is identified from aggregate unemployment, vacancies, and hires. The separation rate, which also plays a nominal role, is set equal to the proportion of

⁵⁶In their survey, [Petrongolo and Pissarides \(2001\)](#) suggest a plausible range of .3 to .5. I obtain a value of approximately .6, likely due to attenuation bias caused by noise in the LIAB-based hiring measures.

workers observed in year t but not year $t + 1$, among establishments observed in both years.

Table 2.1: Aggregate Parameters

Parameter	Definition	Value	Source
ψ	Amenity exponent	.538	Equal to $\frac{\eta}{1-\eta}$
ρ	Discount rate	.042	Discount factor of .96
δ	Separation rate	.165	$\frac{\text{Annual hires}}{\text{Total employment}}$
η	Match elasticity	.35	Kohlbrecher et al. (2016)
ζ	Match efficiency	2.39	$\frac{\text{Annual hires}}{\text{Predicted hires}}$

It remains to discuss the key elasticities in the model: of occupational demand (σ), industry demand (τ), and labor supply (κ). The demand elasticities σ and τ are fundamentally unidentified. In part this is because unit labor costs $SP(s)/m(\lambda(s), s)$ are unobserved and cannot be measured,⁵⁷ but also because α and β are allowed to vary by panel, and it is not possible to disentangle the role of changes to technical demand from that of substitution effects (e.g. in response to changing skill premia). For this reason, I assume that labor demand is unit elastic for all experiments spanning the 1993-2017 sample period. When eventually I turn to forward-looking policy experiments, I relax this assumption and consider the implications of elastic and inelastic demand, as in these cases firm vacancy postings will respond to changes in the “price” of skill. For all quantitative experiments I impose the following simplification:

Assumption 5 (Homogeneous demand elasticity): $\sigma = \tau$ and $\sigma \in \{.5, 1, 2\}$.

A common value allows for a smaller set of results to be considered, which avoids the need for a 3-dimensional grid of results given that κ is also set-valued. Unit elasticities provide a natural baseline, while the values of .5 and 2 are arbitrary and chosen for convenience. Note that because $G(j)$ is chosen so that SP/m is constant for the 2010-2017 panel, counterfactual experiments involving occupational demand α will only result in substitution effects if there are changes to FP or to SP/m . In particular, any changes to the distribution of skill types across industries will be neutral with respect to industry shares.

⁵⁷Wage data does not allow one to disentangle the components of wages - here p and m - whereas even if worker-level productivity were observed, it could not be meaningfully compared across occupations.

The elasticity of supply κ is also problematic for identification. Although I am able to estimate vacancy-filling rates from the LIAB, there are likely to be unobserved amenities that affect the job search behavior of workers, as assumed in the model. In regressions of vacancy-filling rates on skill premia and firm premia, I estimate an elasticity of labor supply with respect to firm premia of between .96 (East Germany) and 2.04 (West Germany). This is consistent with past studies, which suggest a plausible range of 1 – 4 when non-homogeneous separation rates are also taken into account.⁵⁸ I therefore consider values lying in the set $\kappa \in \{.5, 1, 2, 4\}$, though to conserve space I omit results for $\kappa = 4$ in some of the results that follow.

2.4 Quantitative Results

In this section I present the results of simulated counterfactual experiments, using the model developed and estimated as described thus far. I first consider historical experiments intended to address the question: given changes to labor demand over the 1993-2017 period, how did the presence of firm-bias affect the resulting wage trend? The second set of results concerns the predictiveness of skill-bias, and the extent to which firm-bias confounds the relationship between skill-biased demand shocks and wage inequality. Third, I consider the extent to which firm-bias can account for regional variation in wage trends, in this case comparing East and West Germany. Lastly I present several experiments that quantify the response of wages to wage policies, taking into account the equilibrium effects on worker job search and employer labor demand.

2.4.1 Historical Firm-Bias

To what extent does the contribution of labor demand to the historical trend in wage inequality reflect firm-bias, as opposed to skill-bias alone? To answer this question I consider how wage variance would have evolved under the counterfactual scenario in which all firms pay the same wage conditional on skill. This is done by manipulating the entry costs faced by firms so that firm premia are constant. As entry costs also exert a direct effect on vacancy

⁵⁸See for example [Bo, Finan, and Rossi \(2013\)](#), [Sokolova and Sorensen \(2021\)](#), and [Bassier, Dube, Naidu \(2020\)](#).

posting, I “compensate” changes to $C(i, j)$ by manipulating α and β so as to preserve the distribution of vacancies in the 1993-99 panel. From this starting point, changes over time to relative industry and occupation demand are preserved. Note that this experiment does not hold constant labor supply, and under the counterfactual scenario there will be a shift of labor away from submarkets that were formerly high-paying, and towards those that were previously low-paying.

In figure 2.6 I show the model-predicted wage trend, together with the counterfactual trend when firm premia are homogeneous. Results are given for different values of the supply elasticity κ and different sets of parameters. In the first row all time-varying parameters - α (which includes $\gamma_{i,j}$), β , C , and ν - are set at their contemporaneous values. Row 2 contains results when only the demand parameters are changed; shown is the average outcome from two experiments, in which other parameters are fixed at their 1993-99 and 2010-17 values. In the last two rows, industry and occupation demand are considered separately. Tabulated results are given in the appendix.

The predicted trend in wage variance is slightly less than two-thirds of the empirical increase, with the remainder due to widening residual wage inequality that is not captured by industry and occupation. When all parameters are allowed to vary, eliminating firm-bias reduces the predicted trend by between one-half and three-fifths. The demand parameters α and β account for most but not all of the model-predicted trend, the rest being primarily associated with changes to the skill distribution ν .⁵⁹ Hence the role of firm-bias is somewhat greater when we consider only the demand parameters: the predicted trend falls by roughly two-thirds. The interpretation of this result is that, *by itself*, skill-bias explains less than half of the effect of industry and occupation labor demand on the wage variance. The difference between the counterfactual and predicted trends reflects both firm-bias - changes to the distribution of FP - and interactions between skill-bias and firm-bias, due to a higher covariance of skill premia and firm premia. The differences between the columns are substantial, indicating that the response of labor supply to firm premia is an important determinant of the overall impact of firm-bias. When κ is large, the elimination of differences in FP results

⁵⁹Over time, labor shifts towards the extremes of the skill distribution, which naturally tends to increase wage inequality; see the discussion in the previous section.

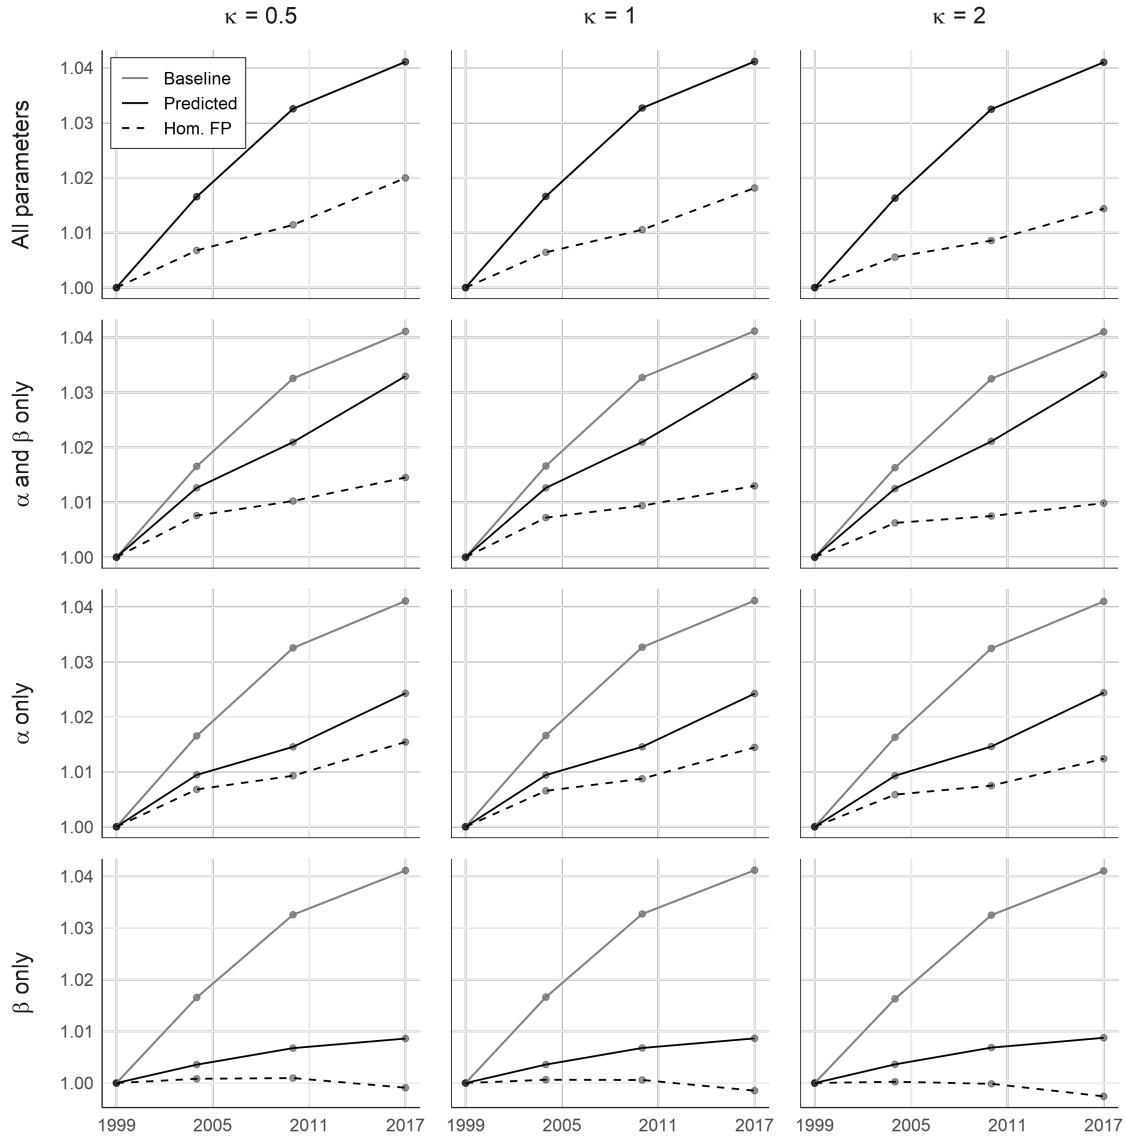


Figure 2.6: Predicted and Counterfactual Wage Trend, 1993-2017

in more workers shifting their job search towards lower-skill sectors. This further reduces the effect of a skill-biased shock to demand, because the impact of a proportional change in skill demand depends on the labor share - and not the vacancy share - of skilled jobs.

When occupation and industry demand are considered separately, it is evident that they are affected to different degrees by the presence of heterogeneous firm premia. The increase in wage variance associated with occupational demand is between one-third and one-half smaller under the counterfactual scenario. The effect of industry demand is more strongly intermediated by firm-bias; whereas β accounts for roughly one-quarter of the model-predicted trend

(an amount equal to one-sixth of the empirical trend), this contribution becomes *negative* when firm premia are homogeneous. The intuition for this result is that, through the lens of the model, the shift in employment from manufacturing to service sectors has not been skill-biased. Rather it has contributed to wage inequality through firm-bias, and through the interaction of firm-bias with pre-existing and largely stable inter-industry differences in worker skill.

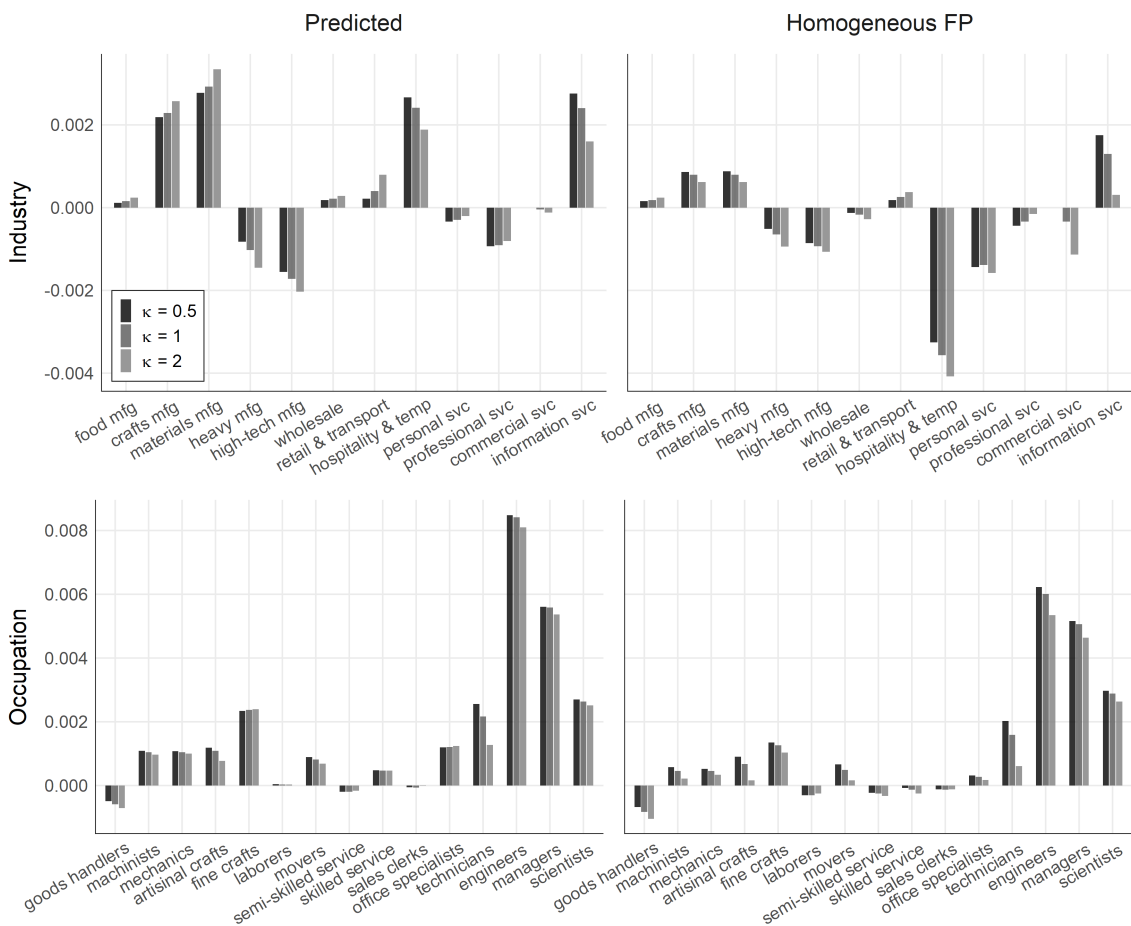


Figure 2.7: Contribution to Change in Wage Variance, 1993-2017

The model allows for further disaggregation of industries and occupations, for which I plot the individual contributions in figure 2.7.⁶⁰ Shown in the first column are the predicted effects on wage variance of changes to α and β between the 1993-99 and 2010-17 panels.

⁶⁰I define occupations as panel-specific distributions over i and j . In practice, confidentiality restrictions prevent me from obtaining the joint distribution for each occupation, which I instead impute from the unconditional distributions of occupations over industry and over worker type.

Counterfactual effects when firm premia are homogeneous are given in the second column. Comparing the two, we can see that firm-bias has tended to amplify the (positive) effects of industry-level demand on wage variance. This is especially true for crafts and material manufacturing, in which employment declines over this period; and similarly for the hospitality/temp, personal services, and information sectors, which have gained share. In the absence of heterogeneous firm premia, the increased use of temp labor would have exerted a large, positive effect on wage variance, versus an even larger negative effect when accounting for the low firm premia associated with this industry group. Turning to occupations, we see a similar overall pattern: firm-bias tended to amplify the distributional effect of skill-biased shocks, and reduce the effect of shocks favoring lower-skilled workers. The differences between the predicted and counterfactual scenarios are smaller, however, due to the fact that occupations are dispersed across industries and are characterized by smaller differences in firm premia. This is also reflected in the insensitivity of results to the job search elasticity κ .

The historical results have several implications. First, the impact of firm-bias is large, and would seem to be of first-order importance for understanding observed trends. Whether this conclusion extends beyond West Germany is unclear, though [Song et al. \(2019\)](#) document a mechanical contribution of firm premia to U.S. wage trends similar to that observed in Germany by [Card, Heining, and Kline \(2013\)](#). On the other hand, if the importance of firm-bias varies across regions then this is also of interest, because it would imply that firm-bias is important for understanding regional differences in wage trends - an idea I revisit when comparing East and West Germany later in this section. Second, the evident differences in firm-bias across industries and occupations suggests that, depending on the part of the labor market in which a skill-biased shock takes place, there may be substantially different effects on wages. In this respect, skill-bias may be poorly predictive of wage outcomes, a possibility I explore next.

2.4.2 Firm-Bias and the Effects of a Skill-Biased Shock

A number of papers study the effects of a skill-biased demand shock on wages, under the assumption of competitive labor markets in which firm heterogeneity is absent or irrelevant.

Is this approach justified? In this section I consider the implications of firm-bias for the predictive ability of simple models of skill-bias. The industry- and occupation-level results in the previous section indicate that even among similarly-skilled occupations and industries, there is substantial variation in firm-bias. Below I formalize this intuition, by considering an experiment in which the occupational distribution is divided into ten skill deciles, and the aggregate effect of a one-percentage point increase in vacancy share is simulated for each of 120 industry-decile cells. Results are shown in figure 2.8, with the first panel giving the effect on wage variance when firm premia are homogeneous (i.e. no firm-bias). The remaining three panels correspond to different values of the supply elasticity κ .

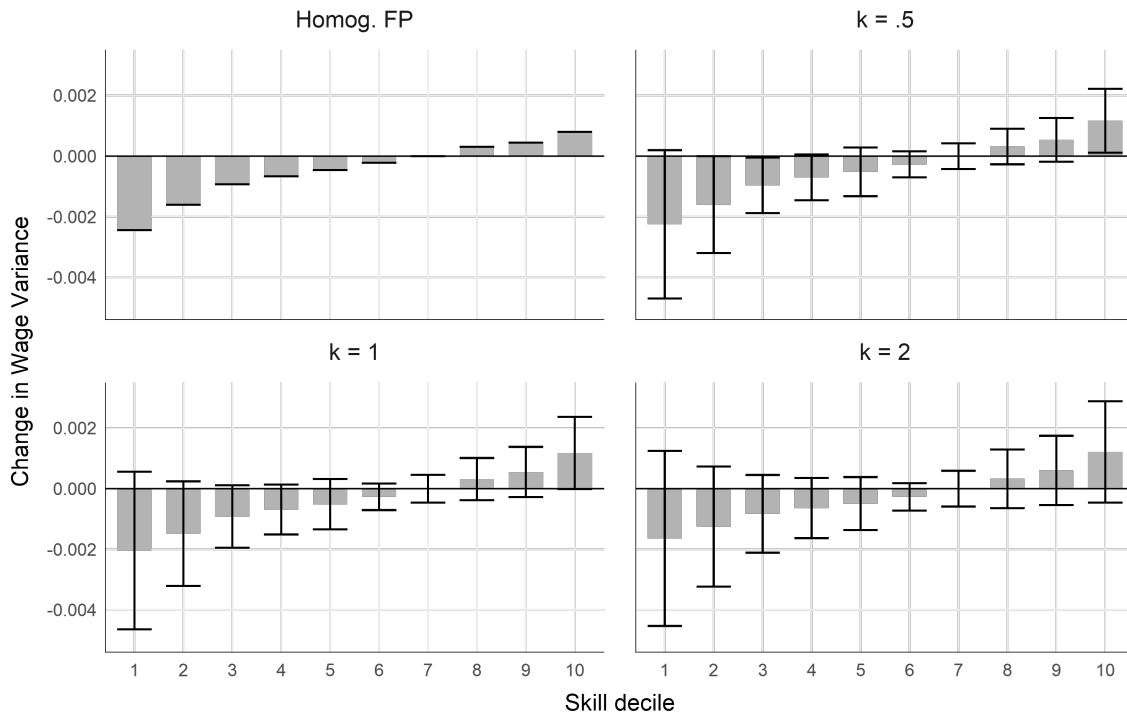


Figure 2.8: Effect of a Skill-Biased Demand Shock

NOTE: Effect on wage variance of a 1% increase in vacancy share, by skill decile and aggregate industry group. Columns (bars) indicate the mean (standard deviation) of the change in $Var(w)$. In the case of homogeneous firm premia the value of κ is set to 1.

Firm-bias can be seen to have two broad effects: it increases the variance of outcomes associated with demand shocks pertaining to any given level of skill, and for high levels of κ it weakens the average relationship between the skill-bias of a shock (moving from left to right on the graph) and the change in wage variance. As the supply elasticity becomes larger, the

predictive power of skill-bias grows weaker and the variance of outcomes correspondingly greater. Even in the case where labor supply is unresponsive ($\kappa = .5$), the influence of firm-bias is substantial; it is possible for an increase in demand in the lowest skill decile to increase wage variance, and an increase in the higher skill deciles to increase it.

These results have two substantive implications. First, firm-bias may act as a confounder in studies relating wage outcomes to environmental changes related to skill-bias. This is potentially important insofar as past researchers have argued that observable skill-biased shocks, such as aggregate adoption of information technology, are only weakly correlated with contemporaneous changes to wage inequality.⁶¹ The patterns in figure 2.8 suggest that, if firm-bias is as quantitatively important in other countries as in Germany, then the relationship between skill-biased demand shocks and wages is likely to vary considerably, and to be apparent only over large samples or long periods of time. The second implication concerns regional wage variation. Past studies have observed that firm premia vary geographically, potentially reflecting differences in market institutions or the organizational structure and characteristics of firms. Industry- and occupation-level shocks are also likely to differ across countries, even when resulting from a common technological trend or change to the trade environment. Hence firm-bias may be able to account for regional variation in wage trends, which is useful not only for our retrospective understanding of wage inequality but also for quantifying the importance of local institutions like collective bargaining. This insight motivates the next section.

2.4.3 Firm-Bias and Regional Wage Trends

In this section I turn to East Germany, where wage variance has increased at a markedly lower pace over the 1996-2017 period.⁶² Can firm-bias account for this difference? Estimating the model for East Germany, I find that the model-predicted trend in wage variance is only three-fourths as large in East Germany as it is in the West, which nevertheless understates the aggregate divergence between these regions: the variance of empirical wages has risen by less than half as much in the East, much of this due to a decrease in residual wage dispersion.

⁶¹See e.g. Card and Lemieux (2001), Card and DiNardo (2002), and Lemieux (2006).

⁶²East German establishments are only included in the LIAB beginning in 1996.

Although firm-bias cannot account for the different residual trends, it can indeed explain the portion predicted by the model. As shown in figure 2.10, the counterfactual trends with homogeneous firm premia are nearly identical between East and West, in particular for larger values of κ .

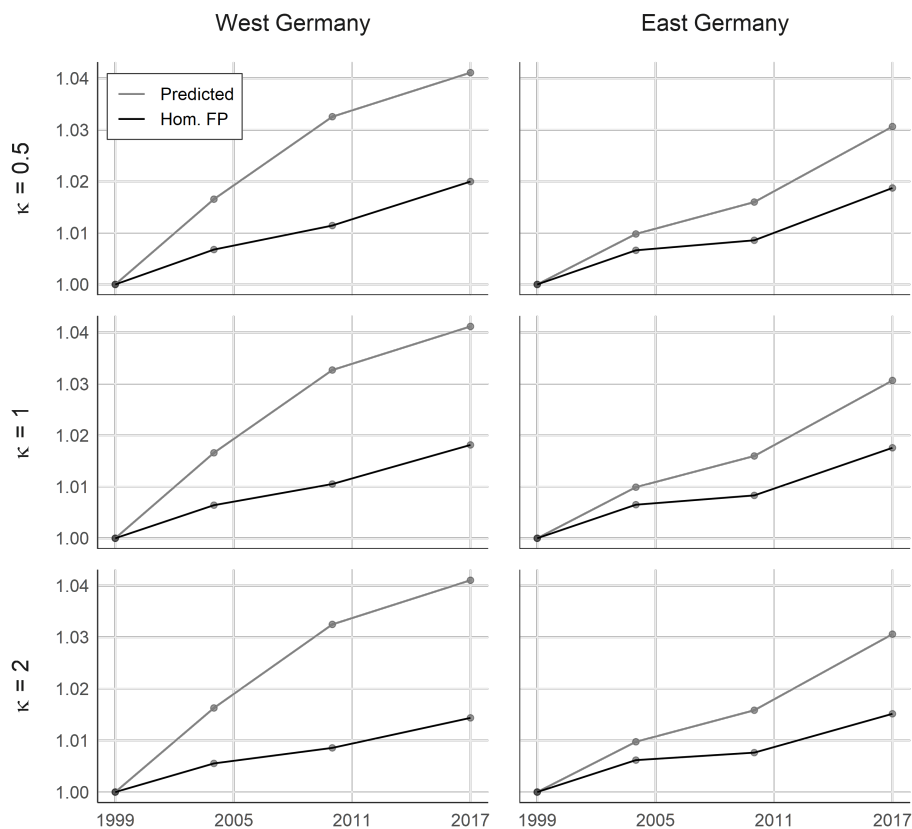


Figure 2.9: Predicted and Counterfactual Wage Variance, East Germany

In some respects the result in figure 2.9 is serendipitous, as changes to labor demand and the skill distribution are sharply different in East Germany. While new entrants to the West German labor force have tended to be either very low-skill or very high-skill, polarizing the estimated skill distribution, the opposite pattern has taken place in the East: there has been what appears to be an exodus of workers at the extremes of the wage distribution. Meanwhile the shift in employment towards service industries is less pronounced, and declining labor demand in the crafts sector has been offset to a larger extent by gains in other manufacturing industries. Hence there is no reason to expect that, in the absence of firm-bias, the two regions would show a similar increase in wage variance. This nevertheless turns out to be

the case, and the greater rise in West German wage dispersion entirely disappears when firm premia are made constant.

The lesser incidence of firm-bias in East Germany is evidently the result of differences in the distribution of firm premia across occupations and industries, and of changes to this distribution in the years following reunification. To show this I conduct a set of simulated experiments in which East German firms face different entry costs, and therefore pay different wages. Demand parameters are compensated so as to hold fixed the 1993-99 distribution of vacancies, as was done previously for the West German results. First I consider counterfactual scenarios in which firm premia are held constant at their 1993-99 and 2010-17 values. I then consider the model-predicted trend when East German firms pay contemporaneous West German wages. Results are plotted in figure 2.10.




Figure 2.10: East German Wage Trend, Alternative Entry Costs

We observe that, if East German firms paid West German wages, or if the wage structure in East Germany were the same in the 1990's as it is currently, the rise in wage variance would have been similar between the two regions. From this result it is no surprise that, by the end of the sample period, the distribution of firm premia is similar between the two regions. While in the years after reunification there was little difference in pay between manufacturing and service industries in East Germany, by the 2010's there is a wage gap similar in magnitude to that observed among West German firms. The changes to firm premia over time have tended to further dampen wage inequality in the East, as manufacturing workers in that

region are comparatively lower-earning, and consequently an increase in industrial wages disproportionately benefits those in the lower half of the wage distribution.

In conclusion, firm-bias can in part account for the slower rise in wage variance in East Germany, owing to differences in the distribution of firm premia that have largely dissipated by the end of the sample period. This result raises two questions of interest. First, what explains the different wage structure observed in East Germany during the 1990's? And second, what accounts for the convergence with West Germany over time? The most likely answer to both questions is institutional in nature. It was recognized after reunification that there were large differences in labor representation and bargaining between the two regions, and a concerted effort was made to export West German trade associations, work councils, and bargaining agreements to East Germany.⁶³ It is also likely that access to capital and exposure to external markets had an impact on the organizational structure (and notably the size) of East German establishments. Such explanations nevertheless suggest that there is an important idiosyncratic component to the distribution of firm premia within an economy, in which case firm-bias should be expected to vary across countries as well as over time.

2.4.4 Firm Premia and Wage Policies

As a final exercise I consider the quantitative effects of wage policies that target firm premia. In recent years, with increased empirical evidence on the contribution of firm premia to wage inequality and to wage trends, it has become more common to proposals for (1) policies targeting competition in output markets, so as to reduce anti-competitive rents, and (2) policies supporting worker bargaining power, to allow for greater pass-through of rents to workers.⁶⁴ On the one hand, such policies might be expected to adversely impact labor demand, and if they target labor markets associated with low-skill labor then any demand response is likely to be skill-biased. On the other hand if workers response to an increase in firm pay by increasing their likelihood of submitting a job application, then supply and demand responses may offset. The ability to quantify these equilibrium responses is a unique

⁶³See for example [Burda and Funke \(1993\)](#) and [Snower and Merkl \(2006\)](#) on extension of West German bargaining agreements to East Germany.

⁶⁴See for example [Tomaskovic-Devey et al. \(2020\)](#) and the [OECD \(2021\)](#) report “The Role of Firms in Wage Inequality: Policy Lessons from a Large Scale Cross-Country Study.”

advantage of the structurally estimated model developed in this chapter.

I consider two experiments, one directly related to German policy discussion and one that is somewhat more speculative. In the first experiment I consider the impact of a policy that enacts “equal pay” for temp agency workers. That temp workers receive lower pay than their full-time coworkers is well-known,⁶⁵ and in Germany this wage gap is usually attributed to differences in collective bargaining coverage. Temp worker pay has become a more important policy issue as the share of temp labor has tripled since the 1990’s, almost entirely due to outsourcing of manufacturing labor.⁶⁶ Current laws mandate that temp workers receive equal pay after 9 months, and these laws have been strengthened over time, but over the sample period the firm premia associated with temp agencies are lower than those in manufacturing by approximately 20-30%. I consider a policy in which rent-sharing ratios are set exogenously so as to close this wage gap - a policy that, while stylized, can be thought of as the extension of collective bargaining agreements to temp workers. To perform this experiment I use LIAB establishment data on temp employment to identify the industry-specific demand for temp labor, in which case the effects of an equal pay policy will reflect the sectoral distribution of establishments in which temp agency employees work.

In the second experiment, I consider how wage inequality is affected by compression of the distribution of firm entry costs, implemented as a concave, monotonic, and rank-preserving transformation of $C(i, j)$.⁶⁷ This experiment is intended to capture the effects of pro-competitive policies that reduce the costs of entry, without taking a stand on which submarkets are most characterized by a lack of competition. As argued by [Autor, Dorn, Katz, Patterson, and Van Reenen \(2020\)](#), firm-level rents may reflect economies of scale that are greater when output markets are competitive. Nevertheless the purpose of this experiment is not evaluate a particular policy but to quantify the equilibrium effects of such a policy on the wage distribution, under the assumption that the policy has its intended effect.

Results for the first experiment are shown in figure 2.11. As the response of labor demand

⁶⁵See for example [Garz \(2013\)](#) and [Jahn and Pozzoli \(2013\)](#).

⁶⁶See appendix for breakdowns of temp employment by industry.

⁶⁷Specifically I set entry costs equal to $\hat{C} = C^v$ where $v \in (0, 1)$; in the experiment shown below I set $v = .95$. Entry costs are then re-scaled so as to maintain the overall level of market tightness.




Figure 2.11: Simulated Effect of Equal Pay for Temp Workers

is a key effect of interest, and as I am considering only the 2010-2017 panel for this exercise and not making comparisons across time, I consider alternative values for the demand elasticities ρ and σ , and for each value three results are shown: (1) the change in wage variance when there is no supply or demand response, which is the value that one would obtain from a “reduced-form” empirical experiment; (2) the change in wage variance once temp agencies reduce labor demand in response to the policy, which has effectively raised their labor costs; and (3) the final effect on wage variance when labor supply is also allowed to respond, and job-seekers increase the likelihood at which they apply to temp agencies.

From the figure it is evident that the equilibrium responses offset, though not entirely; and on net they tend to dampen the the aggregate decline in wage inequality, with this tendency becoming greater as the supply and demand elasticities become large. The magnitude of the offset is modest, ranging from a negligible amount to approximately 20%. When ρ is large, demand for temp workers falls by a greater amount, and by extension demand for low-skill workers falls as well. This reduces employment at temp agencies that are now high-paying, while raising the variance of skill premia. The labor supply response in turn pushes temp agency employment close to its pre-policy level, but some of the increase in the variance of skill premia persists.

The effects of a compression of entry costs are shown in figure 2.12. In this case the demand and supply responses work to opposite effect: a reduction of entry costs in high-




Figure 2.12: Simulated Effect of Entry Cost Compression

FP markets encourages more vacancy postings, which reduces output prices and therefore further reduces firm premia. At the same time this increases the demand for skilled labor as a result of the positive correlation between *SP* and *FP*, and hence the demand response is skill-biased. Labor supply responds by shifting job search towards lower-paying submarkets, which offsets some of the skill-biased change to demand, and pushes firm premia back towards their original level as it now requires more vacancies to hire a given number of workers, effectively raising entry costs.

In both experiments, most of the remaining effect after both supply and demand responses are factored in is the result of greater dispersion of skill premia. The labor distribution is largely unaffected, since the two responses offset. Hence the combined equilibrium response does not reflect skill-bias or firm-bias individually, but the interaction of skill-bias and firm-bias. Each policy is firm-biased by nature, insofar as they disproportionately impact particular submarkets and types of firms; and it is only because these submarkets are also associated with different levels of worker skill that the equilibrium response dampens the overall effect of the policies.

2.5 Conclusion

In this chapter I have presented theoretical and quantitative evidence on the firm-bias of labor demand, its contribution to wage trends in West Germany, and its implications for the study of wage inequality and for policies intended to address inequality. The model I develop provides a rich but empirically tractable framework for studying the relationships between labor demand, supply, skill premia, and firm premia. I estimate that by itself, skill-bias accounts for less than half of the impact of changes to industry and occupation labor demand on German wage inequality, and explains none the rise in wage variance due to changing industry employment. I show that because jobs of all skill levels are stratified across industries, there is considerable variability in the quantitative effect of a change in job demand *conditional on skill*, implying poor predictive power of models that abstract from firm premia. I also show that the effectiveness of labor market policies that target firm premia depends to an important extent on the equilibrium responses of labor demand and labor supply, and on the extent to which skill and firm premia are positively associated across industries and occupations.

Broadly, these results point to an important role played by firms, which intermediate how changes to macroeconomic environment impact labor markets. The traditional approach to studying how environmental changes affect wages is to abstract from firm heterogeneity, but this approach is problematic for several reasons. First, if the structural linkages between workers and firms vary across time and place, then so too will the impact of environmental change. I find that the slower rise of wage inequality in East Germany owes much to differences in the distribution of firm premia across industries and occupations. One benefit of paying greater heed to heterogeneity on the demand side of labor markets is therefore that it may allow us to better understand regional wage trends,⁶⁸ and better quantify the role of local labor market institutions and policies. The equilibrium model developed in this chapter provides a means for studying comprehensively the effects of changes to labor demand, and may be readily applied to other countries for which matched employer-employee data are available.

⁶⁸These differences are often attributed to location institutions; see for example [Dustmann et al. \(2009\)](#) and [Antonczyk et al. \(2010\)](#).

A second key implication of the findings presented here is for policy. Wage inequality is an issue that generates widespread concern, but effective policies require a sound understanding of the underlying mechanisms generating wage gaps. Ignoring the role played by firms and firm premia may lead one to over-estimate the importance of skill, and the quantitative impact of policies that address rising skill premia such as re-training programs. On the other hand, policies targeting the wages set by firms have two issues to tackle. First, the trade-offs associated with these policies will be less adverse when firm premia represent anti-competitive rents in output markets; but firm premia are not themselves proof of anti-competitive rents. For example the environment studied in this chapter features competitive output markets. Second, because skill premia and firm premia are strongly associated across labor submarkets, policies that impact the distribution of firm premia are also likely to affect the distribution of skill premia by generating skill-biased demand responses. I show that the extent to which this is true will depend on the labor supply and demand elasticities, suggesting that these are important objects of interest when considering such policies.

Chapter 3

Task Automation and Labor Polarization

3.1 Introduction

In both the popular and the academic literature, automation is associated with the wholesale replacement of human workers with machines. It is common to read predictions of robots “taking jobs”.⁶⁹ This notion of one-to-one substitution has been formalized in theoretical models of labor-substituting technology and in empirical studies that estimate occupation-level automation risk.⁷⁰ Historical examples of *job*-level automation are nevertheless difficult to find,⁷¹ and case studies of technology adoption generally paint a different picture, in which new technologies substitute for labor at particular tasks within jobs, changing but not eliminating the role of human labor.⁷² Despite this, little is known about the macroeconomic implications of *task*-level automation when tasks are distinct from jobs.

In this chapter, I study how labor markets are impacted by task automation when jobs consist of multiple tasks. I begin with an empirical analysis that draws on four decades of German survey data to show two motivating results. First, virtually all workers report performing a variety of tasks on the job, including both routine and non-routine work. Second, the most important technological shift over this time period - computerization - is associated with an intra-occupational shift in the proportion of time spent on non-routine

⁶⁹For example, “[The Robots Are Coming For Phil in Accounting](#)”, The New York Times, March 6, 2021.

⁷⁰Prominent research includes [Acemoglu and Autor \(2011\)](#), [Frey and Osborne \(2017\)](#), and [Acemoglu and Restrepo \(2018\)](#).

⁷¹See [Autor \(2015\)](#) and [Bessen \(2016\)](#).

⁷²Academic case studies include [Levy and Murnane \(1996\)](#), [Fernandez \(2001\)](#), and [Autor, Levy, and Murnane \(2002, 2003\)](#). For an example from popular media see “[How the World’s Biggest Companies Are Fine-Tuning the Robot Revolution](#)”, The Wall Street Journal, May 14, 2018. Closely related are empirical studies on technological change and within-occupation task and skill requirements, such as [Cappelli \(1993\)](#), [Spitz-Oener \(2006\)](#), [Bartel, Ichniowski, and Shaw \(2007\)](#), [Firpo, Fortin, and Lemieux \(2011\)](#), and [Atalay, Phongthientham, Sotelo, and Tannenbaum \(2020\)](#).

tasks, indicating a within-job shift in task content consistent with the case study literature. To study the implications of task automation in equilibrium, I develop an occupational assignment-based model of labor markets in which jobs involve multiple tasks, and workers solve a time allocation problem. Automation is formalized as the replacement of labor with capital at a single, low-skill task. In this environment, automation has several effects: a standard *labor-substitution* effect in which less labor is needed to produce a given amount of output, a *skill-complementing* effect arising from the reallocation of workers' time towards more skilled task content, and a *demand* effect that depends on the response of unit labor costs, together with the elasticity of substitution across occupations.

The model delivers several main predictions. In the short-run, when the automating technology is expensive, it will be adopted only in skilled occupations. The reason is that skilled labor has a comparative disadvantage at the unskilled task. Automation in the short-run tends to polarize the wage and employment distributions: job losses are concentrated around the “marginal” occupation where firms are indifferent between automating or not, while occupations further from this threshold benefit from a larger decline in unit costs. Employment polarization in turn drives polarization of wages. As capital prices continue to fall, employment in affected occupations recovers provided that occupations are not perfect complements, since declining capital costs tend to reduce unit labor costs and increase relative demand. In the long-run, after adoption is universal and as the cost of the automating technology approaches zero, employment in low-skill occupations rises and wage inequality falls. Hence low-skill unemployment and wage inequality exhibit a hump-shaped pattern over time, and the disruptions caused by automation are less severe in the long-run than in the short-run.

In the final part of the chapter I test the model's short-run predictions, and structurally estimate a continuous version of the model in order to obtain long-run forecasts. Consistent with model predictions, computerization in West Germany exhibited a top-down pattern, with significant PC adoption in low-skill occupations only occurring in the 2000's. Increased PC use is associated with contemporaneous declines in occupational employment, while adoption in *previous* surveys is predictive of occupational growth. The estimated structural model predicts that job losses in middle-skill occupations have peaked, as has the 90/50

wage ratio, with future computerization expected to drive wage inequality in the lower half of the wage distribution. In the very long-run, employment in high-skill occupations declines relative to current levels, and the 90/10 wage ratio increases slightly. As a last exercise I use the quantitative model to explore implications for two commonly-studied elasticities. First, I show that the capital-labor elasticity of substitution is a complex function of rental rates and worker skill, and is neither constant nor exogenous as is commonly assumed. Second, higher values of the labor-labor elasticity of substitution are associated with greater short-run (but not long-run) changes to occupational employment. This elasticity is therefore important for understanding the persistence of job loss from automation, and the efficacy of supply-side policies such as retraining programs.

The main contribution of this paper is to the macroeconomic literature on the effects of automation. Substitution of capital and labor is traditionally modeled in simple fashion as depending on the parameters of an aggregate production function, a tradition carried over to the case of automation by [Zeira \(1998\)](#). Subsequent work has also taken this approach, including [Acemoglu and Autor \(2011\)](#), [Peretto and Seater \(2013\)](#), [Acemoglu and Restrepo \(2018\)](#), [Aghion, Jones, and Jones \(2019\)](#), and [Hemous and Olsen \(2020\)](#). Theoretical models of capital-labor substitution yield a simple, negative relationship between technology adoption and employment, that has been used to generate forecasts for future disemployment as in [Frey and Osborne \(2017\)](#). Several strands of criticism exist. A number of authors have observed that history offers few examples of wholesale automation of jobs, which are rarely reducible to a single, automatable task.⁷³ In addition, there have been significant within-occupation changes to task content that are difficult to explain as a result of job-level automation.⁷⁴ I reconcile these strands of literature with the idea of labor-substituting capital by relaxing the assumption of one-to-one substitution, and by providing both empirical evidence and theoretical results on the automation of tasks at a sub-job level.

This chapter also contributes to the literature on how occupational structure intermediates the effects of macroeconomic change. The framework developed here represents an

⁷³For example [Autor \(2015\)](#), [Bessen \(2016\)](#), [Arntz, Gregory, Zierahn \(2017\)](#), and [Dengler and Matthes \(2018\)](#)

⁷⁴See [Levy and Murnane \(1996\)](#), [Fernandez \(2001\)](#), [Autor, Levy, and Murnane \(2002, 2003\)](#), [Spitz-Oener \(2006\)](#), [Bartel, Ichniowski, and Shaw \(2007\)](#), [Firpo, Fortin, and Lemieux \(2011\)](#), and [Atalay, Phongthientham, Sotelo, and Tannenbaum \(2020\)](#)

intermediate point between models in which occupations consist of task bundles requiring task-specific skills, such as [Gathmann and Schonberg \(2010\)](#), [Yamaguchi \(2012\)](#), and [Autor and Handel \(2013\)](#), and models that abstract from occupational task structure and assign to each occupation an exogenous value (e.g. an index) describing the relationship between productivity and skill (or a skill composite), such as [Costinot and Vogel \(2010\)](#) and [Acemoglu and Autor \(2011\)](#). I am able to maintain the tractability in general equilibrium of the second class of models, while incorporating the occupational task structure studied by the first group. By explicitly modeling the allocation of workers' time across tasks, I am also able to incorporate the time-reallocation effect often mentioned in anecdotal and case study accounts of automation, and to estimate task production shares from empirical task frequencies. This approach is more general than that taken in the task bundling literature, where occupational output is usually constrained to be a linear function of task output.

The structure of this chapter is as follows. In section 2 I show the main descriptive results: that virtually all workers report performing both routine and non-routine tasks, and computerization is associated with greater time spent on non-routine tasks. In section 3 I develop a model of task-based automation with multi-task occupations. I begin with a simple environment in which all jobs are homogeneous within an occupation, which in turn yields a set of precise analytical predictions. I then allow for intra-occupational heterogeneity, which weakens the analytical results but allows the model to be taken more directly to the data. Results on the qualitative model predictions are shown in section 5, and quantitative estimation and predictions are discussed in section 6.

3.2 Descriptive Analysis

In this section I present motivating evidence on within-job task content. After describing the data and the task measures used in the analysis, I show two main facts. First, the vast majority of jobs involve both routine and non-routine tasks, indicating that “wholesale” automation of jobs is *ex ante* unlikely. Second, computerization of occupations is associated with a shift towards more non-routine task content, suggesting that *ex post*, technological automation has had an impact on the within-job distribution of task content. These facts are inconsistent with stylized models of automation that assume an equivalence between

tasks and jobs, and they constitute evidence that technological change has interacted in a substantive way with the distribution of tasks within occupations, as well as between them.

3.2.1 Occupational Routineness in the BIBB Surveys

This analysis draws on the seven BIBB Employment Surveys, collected by the Federal Institute for Vocational Education and Training (BIBB) over the period 1979-2018 in partnership with the Institute for Employment Research (IAB, 1979-1999) and the Federal Institute for Occupational Health and Safety (2006-2018). Each survey draws on a random sample of the employed German labor force, and asks respondents a range of questions concerning job task content, the use of technology and tools, and other aspects of the job environment and the individual’s work history. Also contained is information on monthly wage, which I convert to an hourly number using reported weekly hours.⁷⁵ The first two surveys do not contain data on East German workers and consequently, for comparability over time, I limit the analysis in this paper to West German workers between the ages of 18 and 65.

Table 3.1: Summary Statistics for BIBB Employment Surveys, 1979-2018

	1979	1986	1992	1999	2006	2012	2018
Observations	28,595	26,090	23,940	27,371	15,905	16,330	16,699
with task data	27,709	25,859	23,830	27,229	15,893	16,257	16,617
with wage data*	26,474	22,631	21,198	20,840	13,607	13,230	13,952
PC use (%)	5.5	17.3	34.1	54.7	83.0	85.3	89.9
weighted	4.6	15.6	32.9	52.1	78.6	82.0	84.0
3-digit occupations	314	314	308	297/350	289/347	297	338

The detailed information contained in the BIBB surveys, together with the long time-frame over which they have been collected, makes them especially well-suited to research on the impact of technological change. Past examples of such research include [Spitz-Oener \(2006\)](#) and [Bachmann, Cim, and Green \(2019\)](#). There are nevertheless several challenges that must be addressed when using the BIBB surveys to study questions relating technology to labor markets. First, survey questions are often inconsistent across panels. For this project I rely on two sets of survey questions: those concerning PC use on the job, and

⁷⁵Wage data are highly aggregated for surveys prior to 2005/06, and consequently of limited usefulness in this study.

those on task performance. Questions about PC use are broadly consistent over time, with one substantial change in format and wording occurring between 1992 and 1999. Questions regarding task content - for example, whether and how often workers “control machines”, “advise others”, and so forth - vary substantially across most survey years. For example, the 1979 survey considered more than 80 such tasks, while the 1999 survey contained only 13. Past research has relied on aggregation methods in order to obtain a smaller set of task categories that can be compared across time, but such methods are essentially *ad hoc* and [Rohrbach-Shmidt and Tiemann \(2013\)](#) show that different approaches can have a substantial impact on measures of occupational routineness.

A second issue is that job tasks do not map directly into conceptual frameworks like skillfulness or routineness, and researcher interpretation is therefore required when relating empirical results to hypotheses. This is especially problematic so when combined with the aggregation issues discussed above. Methods like factor analysis that provide a principled approach to dimensionality reduction are, by nature, unlikely to result in easily-interpreted task groups. Approaches based on intuition - combining tasks that “seem” similar - are perhaps more meaningful, but provide no guidance on precisely how tasks should be aggregated.

Table 3.2: Routine Risk Content in the BIBB Surveys

	Repeat tasks	Follow instructions	Adapt to new tasks	Improve procedures	Solve problems	Make decisions
Survey years						
1979-1992	.491	.654	.647	.455		
1999-2018	.504	.678	.694	.619	.799	.633

Education						
None	.530	.779	.585	.517	.583	.441
Vocational	.533	.767	.690	.619	.711	.587
University	.383	.546	.821	.752	.838	.766

Wage pct.						
1-25	.500	.780	.615	.554	.644	.490
26-50	.542	.749	.694	.624	.717	.586
51-75	.504	.703	.758	.680	.764	.655
76-100	.421	.586	.803	.732	.812	.750

NOTE. Survey answers re-coded as frequencies and means calculated using survey weights. Results by education and wage percentile are for the 2006 survey. See appendix for details.

For both of these reasons I depart from past literature and limit attention to a set of

survey questions about *task characteristics*. In all seven panels, and with minimal changes in wording and response categories, respondents are asked how often they find themselves (1) repeating the same work process, (2) following detailed instructions, (3) adapting to new tasks, and (4) improving existing procedures or trying something new. These questions would seem to bear directly on the amount of routine task content present in the job, which is typically described in the literature as some combination of task repetition and the ability to codify task performance into a set of steps that a machine or computer might execute. I also consider questions from the 2006-2018 panels on whether respondents must (5) react to problems and solve them or (6) make difficult decisions on their own. These last three panels also feature a consistent set of questions regarding job tasks, for which I provide reference results in the appendix. Responses consist of between three and five verbal frequencies - e.g. “often” or “never” - which I interpret numerically and assign in even intervals to $[0, 1]$.

Table 3.2 shows conditional mean values after splitting the sample by year and, for the 2006 year survey, by educational attainment and wage quartile. Less education and a lower wage is associated as expected with greater task repetition and instruction, and less adaptive and cognitive content. Jobs do not appear less routine in later years, although there is a shift towards greater cognitive content that is likely the result of increased employment in professional and technical occupations. Comparisons with task content (see appendix) indicate that repetition and detailed instructions are associated with manual labor tasks, while non-routine characteristics are associated with tasks relating to analysis, instruction, advising, and sales and marketing. Notably, production-related tasks are associated with both routine and non-routine characteristics, suggesting a potentially complicated relationship between tasks and task characteristics like routineness.

3.2.2 Task Variety Within Jobs

Are all tasks within a job equally susceptible to automation? In figure 3.1 I show, for each wage percentile, the distribution of survey responses for the 2006 survey.⁷⁶ Across all wage percentiles, the majority of respondents report performing both routine and non-routine

⁷⁶This is the first year containing all of the measures discussed in the previous section, and also the first year in which wage data are not aggregated into coarse bins

tasks. Although higher-earning respondents perform non-routine content more frequently, and routine content less so, much of the variation in responses is within rather than between wage percentiles. In terms of task repetition, the 25th and 75th wage percentiles look similar; and even in the lowest quartile, three-quarters of respondents report that their job involves decision-making, problem-solving, and other task content that previous literature asserts is difficult to automate.

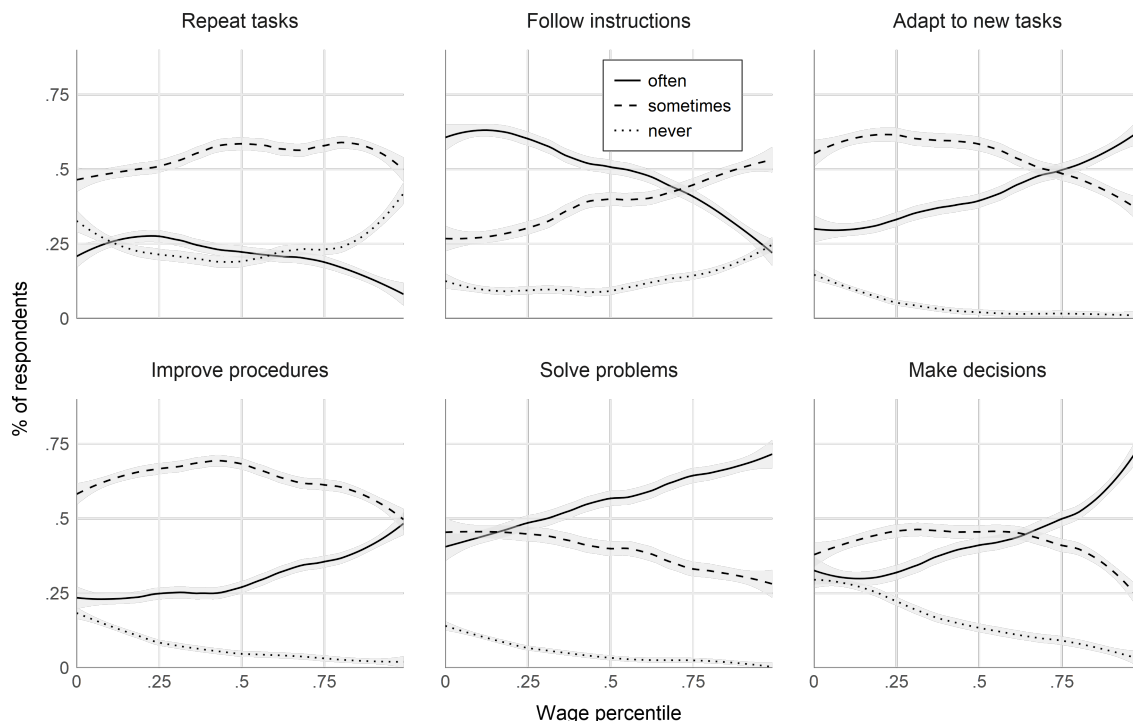


Figure 3.1: Routine and Non-Routine Activity by Wage Percentile, 2006

NOTE. LOESS interpolation of mean response by wage percentile. Gray regions indicate 95% confidence intervals.

Similar results can be shown for the *type* of tasks performed on the job, summarized in figure 3.2. Of the sixteen task categories present in the 2006 survey, four-fifths of respondents report performing at least 5. Of six broad task groups represented, four-fifths of respondents perform tasks in at least 3 of these groups. Similarly, when asked how often they must perform many different tasks at the same time, virtually all individuals report that this is at least sometimes true, and two-thirds that it is often the case. This and the previous figure offer a stark rebuttal to the assumption that jobs and tasks are equivalent. There exists a variety of tasks not just within a given occupation but within individual jobs, and in most

cases this variety includes both routine and non-routine task content.

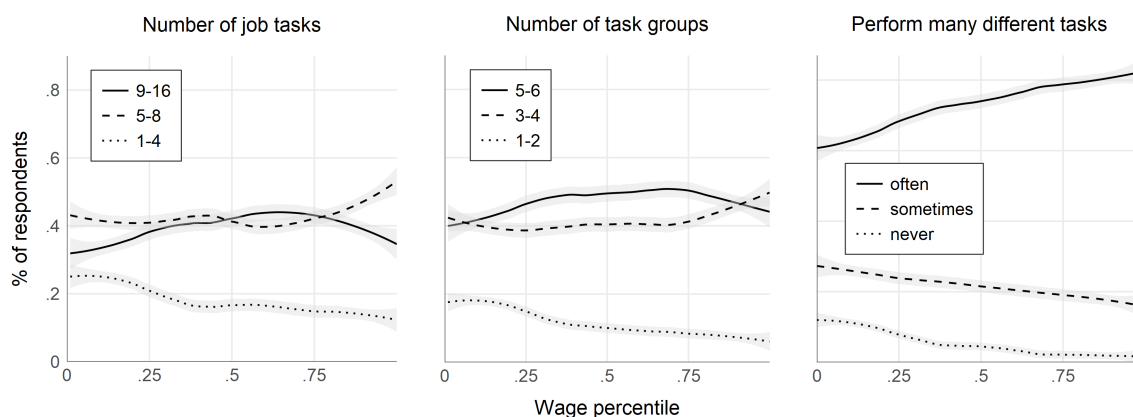


Figure 3.2: Task Variety by Wage Percentile, 2006

NOTE. LOESS interpolation of mean response by wage percentile, with gray regions indicating 95% confidence intervals. See appendix for description of tasks and task groups.

One notable aspect of the survey responses shown in figure 3.1 is the lack of any discernible discernible hump-shape in measures of routineness. Although Germany, like the United States, has experienced labor market polarization,⁷⁷ and employment has generally declined in “middle-skill” jobs, there is little evidence in the BIBB surveys that jobs in the middle of the wage distribution involve greater routine task content. Measures of routineness peak in the first quartile, while measures of non-routine task content are generally increasing in wage. If routineness were the *only* predictor of employment effects from technology adoption, then it is difficult to avoid the conclusion that disemployment should be greatest at the bottom of the wage distribution. More broadly, these results would seem to be inconsistent with theories that pose a simple relationship between task content and automation.

3.2.3 Computerization and Task Content

If automation is task-specific and jobs vary in their task content, then we would expect automation to affect the within-job distribution of tasks. A simple test of this hypothesis is to see whether technology adoption is associated with changes in the overall routineness of job tasks. Below I focus on the use of personal computers in the workplace. A large body of literature argues that computerization reduces time spent on routine tasks, and in this case

⁷⁷See e.g. [Dustmann, Ludsteck, and Schonberg \(2007\)](#) and [Bachmann, Cim, and Green \(2019\)](#).

we should observe less frequent routine task content as the use of computers increases. I begin with a cross-sectional analysis. In table 3.3 I show the marginal effect of computerization on task characteristics, estimated by performing fractional logit regressions of task content on computer use. All regressions control for 3-digit occupation, 1-digit industry, and year. Because the wording and format of the survey question on computer use changes between 1992 and 1999, I split the sample into two subsamples covering the 1979-1992 and 1999-2018 periods. Conditional on the type of work being done, computerization is associated with relatively less routine task content, and relatively more non-routine content. This result holds within educational groups and wage quantiles, as well as across panels.

Table 3.3: Task Characteristics and PC Use, 1979-2018

Sample	<i>Dependent variable</i>					
	Repeat tasks	Follow instructions	Adapt to new tasks	Improve procedures	Solve problems	Make decisions
<i>Survey years 1979 - 1992</i>						
Full sample	-0.054 (.005)	-0.046 (.005)	.103 (.004)	.094 (.004)		
Education						
None	-0.043 (.021)	-0.043 (.016)	.147 (.019)	.107 (.017)		
Vocational	-0.038 (.006)	-0.033 (.006)	.098 (.005)	.093 (.005)		
University	-0.039 (.009)	-0.039 (.010)	.051 (.007)	.062 (.008)		
Wage pct.						
1-25	-.024 (.012)	-.016 (.011)	.099 (.011)	.072 (.010)		
26-50	-0.046 (.011)	-.011 (.010)	.096 (.010)	.071 (.009)		
51-75	-0.034 (.010)	-0.036 (.009)	.080 (.009)	.072 (.009)		
76-100	-0.044 (.009)	-0.056 (.010)	.064 (.007)	.084 (.008)		
<i>Survey years 1999 - 2018</i>						
Full sample	-0.013 (.006)	-0.045 (.006)	.102 (.004)	.116 (.005)	.098 (.005)	.145 (.007)
Education						
None	.009 (.017)	-.016 (.016)	.119 (.014)	.130 (.014)	.128 (.018)	.162 (.020)
Vocational	.003 (.007)	-0.039 (.006)	.098 (.005)	.102 (.006)	.098 (.006)	.133 (.008)
University	-.020 (.013)	-0.041 (.015)	.090 (.008)	.103 (.010)	.082 (.009)	.147 (.018)
Wage pct.						
1-25	.004 (.012)	-0.032 (.010)	.107 (.009)	.102 (.009)	.089 (.011)	.134 (.013)
26-50	-.005 (.012)	-.016 (.011)	.079 (.009)	.083 (.010)	.079 (.011)	.082 (.014)
51-75	-.025 (.013)	-0.058 (.013)	.077 (.010)	.088 (.010)	.086 (.013)	.095 (.018)
76-100	-.030 (.019)	-0.050 (.021)	.100 (.012)	.121 (.014)	.068 (.017)	.165 (.025)

NOTE. Marginal effects and robust standard errors from fractional logit regressions of task characteristics on PC use, aggregated by 3-digit occupation. All regressions include dummies for year, 3-digit occupation, 1-digit industry. Bold results indicate 95% significance. See appendix for details.

To better the answer the question of how *changes* in computer use and task content are related, I aggregate survey responses by 3-digit occupation and perform a series of difference-in-difference regressions, shown in figure 3.4. Regressions are performed for consecutive panels in order to minimize confounding from time-variation in other variables that may affect task content; results for 1992-1999 are shown, but indicated in italics due to the change in question format. The results are comparable to those shown previously: when significant, the coefficients on routine task content are negative and those on non-routine content are positive. Marginal effects are large, and indicate that a 10% increase in occupational computer use is associated with a 1-2% change in task frequencies. Note that even if the presumed direction of causality is correct - i.e. from computerization to task frequencies - there are two ways in which this effect might occur. Task content may change within individual jobs,⁷⁸ or employment may substitute away from jobs in the same occupation but with less routine task content. Either story would nevertheless yield a similar implication: automation is unlikely to replace occupational labor wholesale, and employment outcomes would depend both on the degree to which labor is automated, and on any changes to occupational productivity.

Table 3.4: Occupation Mean Task Characteristics and PC Use (Two-Way FE)

Sample	<i>Dependent variable</i>					
	Repeat tasks	Follow instructions	Adapt to new tasks	Improve procedures	Solve problems	Make decisions
KLDB 1988						
1979-1986	-.049 (.034)	-.077 (.035)	<i>-.022 (.052)</i>	.073 (.036)		
1986-1992	-.220 (.024)	-.069 (.033)	.072 (.027)	.004 (.027)		
<i>1992-1999</i>	<i>.046 (.030)</i>	<i>.028 (.028)</i>	<i>-.025 (.025)</i>	.069 (.026)		
1999-2006	-.118 (.031)	-.106 (.030)	.111 (.019)	.164 (.023)		

KLDB 1992						
1999-2006	-.077 (.028)	-.086 (.025)	.114 (.018)	.144 (.020)		
2006-2012	-.018 (.036)	-.037 (.033)	.053 (.025)	.078 (.026)	.095 (.023)	.176 (.033)
2012-2018	-.068 (.040)	-.051 (.042)	.053 (.030)	.152 (.030)	.064 (.029)	.161 (.040)

NOTE. Marginal effects and robust standard errors from fractional logit regressions of task characteristics on PC use, aggregated by 3-digit occupation. All regressions include occupation and year dummies. Bold results indicate 95% significance. Italic results indicate a change in question wording between the two samples. See appendix for details.

These findings, like those in the previous section, are difficult to reconcile with standard

⁷⁸This would be consistent with case studies such as [Autor, Levy, and Murnane \(2002\)](#).

notions of automation as the wholesale substitution of capital for labor. Past authors have noted that it is difficult to find examples of jobs (i.e. occupations) eliminated by automation.⁷⁹ The results shown in this section indicate that it is difficult to find examples of jobs that *could* be entirely automated. The technological change most studied in the literature - widespread adoption of the PC - is associated with changes in occupational task content, indicating a reallocation of labor across tasks within occupations, and not just between occupations. These facts form the basis for the model developed in the next section.

3.3 A Model of Partial Automation

In this section I develop a model of automation where jobs consist of multiple tasks. As is common in the literature, automation is formalized as the substitution of capital for labor at a particular task. In contrast to one job-one task models such as [Acemoglu and Autor \(2011\)](#), the effects of automation are not limited to a reduction in the demand for labor. Automation will also affect workers' time allocation across tasks, and hence occupational returns to skill. And because automation will tend to reduce unit costs within a given occupation, the overall effect on employment is ambiguous, and will depend on the relative magnitudes of countervailing substitution effects.

3.3.1 Environment

The environment is static and consists of workers and firms. Workers are heterogeneous in a continuous skill variable $s \in [0, 1]$ with distribution $F(s)$. Intermediate good producers are heterogeneous in a continuous variable $j \in [0, 1]$ indicating the type of labor output they produce. Here, j is interpreted as referring to an “occupation”, with intermediate producers acting as aggregators of occupation-specific labor. The output of j -producers is in turn aggregated by final good producers that require no labor input.

Occupational output. Workers in a j -firm produce labor output by allocating a unit of time between a low-skill and a high-skill task, denoted l and h . Task output depends on

⁷⁹Examples include [Bessen \(2016\)](#) and [Autor \(2015\)](#).

worker skill, and is given by the functions $\gamma_l(s)$ and $\gamma_h(s)$. I assume that there are no agency problems: workers choose their time allocation so as to maximize firm output. The unskilled task may be partially automated, in which case a portion κ or less of task output may be performed by capital at a per-unit rental cost of r . Tasks are assumed to be perfect complements, with worker output given by the Leontief output function

$$y(j, s, K) = \min \left\{ \frac{t\gamma_h(s)}{\alpha(j)}, \frac{(1-t)\gamma_l(s) + K_l}{1 - \alpha(j)} \right\} .$$

In this equation t is the proportion of the worker's time allocated to the skilled task, and K_l the (per-worker) capital allocated to the unskilled task. The automation feasibility constraint implies that if K total capital is provided to the worker, then

$$K_l = \min \left\{ K, \frac{\kappa}{1 - \kappa} [1 - t]\gamma_l(s) \right\} .$$

The task share function $\alpha(j) > 0$ is assumed to be continuously differentiable and to lie strictly between 0 and 1 for all j , implying that there are no jobs consisting of a single task. Wages are denoted $w(s)$ and intermediate output prices $p(j)$. Per-worker profits of intermediate good producers can be written as

$$\pi_i(j, s, K) = p(j)y(j, s, K) - w(s) - rK .$$

I assume there is free entry of producers, and for now I abstract from producer scale.

Aggregate technology. Final good producers aggregate j -output into a consumption good using the technology

$$Y = \left(\int \beta(j)Y(j)^{\frac{\rho-1}{\rho}} dj \right)^{\frac{\rho}{\rho-1}} .$$

I assume that production of the final good requires only j -inputs, and normalize the price of the final good to 1. Given these assumptions, the profit of the representative final good

producer will be

$$\pi_f = Y - \int p(j)Y(j)dj .$$

Final good markets are assumed to be perfectly competitive.

Agents' problems. I now formalize the problems of workers and producers. Beginning with workers, the time spent on tasks is chosen to maximize output, given assigned capital K_l :

$$\max_{t \in [0,1]} \min \left\{ \frac{t\gamma_h(s)}{\alpha(j)}, \frac{(1-t)\gamma_l(s) + K_l}{1 - \alpha(j)} \right\} , \quad (19)$$

where K_l is subject to the feasibility constraint above. Intermediate good producers choose worker skill s and per-worker capital K in order to maximize profits:

$$\begin{aligned} \max_{K \geq 0, s \in [\underline{s}, \bar{s}]} p(j)y^*(j, s, K) - w(s) - rK \\ \text{s.t. } y^*(j, s, K) \text{ solves (19)} . \end{aligned} \quad (20)$$

Final good producers then choose the mix of j -goods that maximizes profits:

$$\max_{Y(j)} \left(\int \beta(j)Y(j)^{\frac{\rho-1}{\rho}} dj \right)^{\frac{\rho}{\rho-1}} - \int p(j)Y(j)dj . \quad (21)$$

Note that because capital costs are fixed when workers choose their time allocation, no inefficiency is introduced by separating the worker's time allocation problem from the intermediate producer's choice of capital.

3.3.2 Equilibrium

The worker's time allocation problem (19) yields two solutions, depending on whether the technological constrain $(1 - \kappa)K \leq \kappa[1 - t]\gamma_l(s)$ is binding:

$$t^*(j, s, K) = \begin{cases} \frac{\frac{\alpha(j)}{\gamma_h(s)} \left[1 + \frac{K}{\gamma_l(s)}\right]}{\frac{\alpha(j)}{\gamma_h(s)} + \frac{1-\alpha(j)}{\gamma_l(s)}} & K \leq \frac{\kappa(1-\alpha(j))}{\frac{\alpha(j)}{\gamma_h(s)} + \frac{1-\alpha(j)}{\gamma_l(s)}} \\ \frac{\frac{\alpha(j)}{\gamma_h(s)}}{\frac{\alpha(j)}{\gamma_h(s)} + (1-\kappa)\frac{1-\alpha(j)}{\gamma_l(s)}} & K > \frac{\kappa(1-\alpha(j))}{\frac{\alpha(j)}{\gamma_h(s)} + \frac{1-\alpha(j)}{\gamma_l(s)}} \end{cases} .$$

Both costs and output are linear in capital, and so profit maximization by producers (20) will entail a corner solution: producers will choose either $K = 0$ or $(1 - \kappa)K = \kappa[1 - t]\gamma_l(s)$, depending on whether it is cheaper to hire additional labor or to automate the low-skill task. Formally,

$$K^*(j, s) = \begin{cases} 0 & w(s) \leq r\gamma_l(s) \\ \frac{\kappa(1-\alpha(j))}{\frac{\alpha(j)}{\gamma_h(s)} + (1-\kappa)\frac{1-\alpha(j)}{\gamma_l(s)}} & w(s) > r\gamma_l(s) \end{cases} . \quad (22)$$

On the other hand the first-order condition for worker skill implies that in equilibrium we must have

$$w'(s) = \begin{cases} \frac{d}{ds} \left(\frac{\alpha(j)}{\gamma_h(s)} + \frac{1-\alpha(j)}{\gamma_l(s)} \right)^{-1} & w(s) \leq r\gamma_l(s) \\ \frac{d}{ds} \left(\frac{\alpha(j)}{\gamma_h(s)} + (1-\kappa)\frac{1-\alpha(j)}{\gamma_l(s)} \right)^{-1} & w(s) > r\gamma_l(s) \end{cases} .$$

The optimal choices of s and K will depend on the wage function. This is a potential problem, as it may not be possible to characterize the optimal assignment (and hence the wage function) without knowing producers' automation decisions. Nevertheless I show below that in this simple environment, K^* can be characterized *ex ante* to a degree sufficient for the optimal assignment and wage functions to be characterized and solved numerically without difficulty.

The final good producer's problem yields the first-order condition

$$Y(j) = \left(\frac{\beta(j)}{p(j)} \right)^{\rho} Y .$$

Free entry, on the other hand, has the implication that $p(j) = \frac{w(j)+rK(j)}{y(j)}$. Hence we can write total j -employment as

$$L(j, s) = \frac{y^*(j, s, K)^{\rho-1}}{\left(w(s) + rK^*(j, s)\right)^\rho} \beta(j)^\rho Y .$$

Fixing skill and wages, automation will increase labor demand whenever

$$\left(\frac{\frac{\alpha(j)}{\gamma_h(s)} + \frac{1-\alpha(j)}{\gamma_l(s)}}{\frac{\alpha(j)}{\gamma_h(s)} + (1-\kappa)\frac{1-\alpha(j)}{\gamma_l(s)}} \right)^{1-\rho} > \left(\frac{w(s) \left(\frac{\alpha(j)}{\gamma_h(s)} + (1-\kappa)\frac{1-\alpha(j)}{\gamma_l(s)} \right) + r(1-\alpha(j))\kappa}{w(s) \left(\frac{\alpha(j)}{\gamma_h(s)} + (1-\kappa)\frac{1-\alpha(j)}{\gamma_l(s)} \right)} \right)^\rho . \quad (23)$$

The right-hand side of (23) will always be greater than one, and hence for employment to increase it must be that automation decreases unit output costs, and that ρ is sufficiently large. Unit output costs will decrease by a greater amount when $\alpha(j)$ is large but automation has a large effect on labor productivity, e.g. when workers spend a large amount of time on the unskilled task because they have a comparative disadvantage at that task. Note, however, that $w(s)$ will be endogenous to r in equilibrium and so a critical factor in determining the employment effects of automation is the response of the wage function. On the other hand the capacity for automation κ will inform the overall magnitude of the employment effect. This is a departure from past, well-known models in that automation does not necessarily reduce employment, and the set of workers affected is an endogenous outcome depending on relative costs.

Optimal assignment. I now turn to the optimal assignment of workers to jobs, which will only be well-defined when different types of workers enjoy a comparative advantage at different types of jobs. I therefore impose the condition:

Assumption 6: $\frac{d}{ds} \log \gamma_h(s) > \frac{d}{ds} \log \gamma_l(s) > 0$ and $\alpha'(j) > 0$.

Task output is increasing in skill, and this increase is proportionally greater for skill-intensive tasks. That $\alpha(j)$ is increasing implies that jobs are ordered in terms of their return to skill, from which we can predict that in equilibrium, higher j will be associated with higher s . In

order for the optimal assignment to be a smooth function, I further impose the conditions that $\beta(j)$ and $F(s)$ are continuously differentiable, and that $F(s)$ is a strictly increasing function. With these restrictions, the model can be shown to possess two properties that allow automation decisions (i.e. K^*) to be characterized separately from the wage function, allowing for a precise characterization of the equilibrium. Denoting \bar{K} to be the optimal level of capital conditional on automation, the model satisfies the following conditions:

1. **rank-preserving:** for any two occupations j and j' , if $\frac{d}{ds} \log y^*(j', s, 0) > \frac{d}{ds} \log y^*(j, s, 0)$ then $\frac{d}{ds} \log \left[p(j)y^*(j', s, \bar{K}) - r\bar{K} \right] > \frac{d}{ds} \log \left[p(j)y^*(j, s, \bar{K}) - r\bar{K} \right]$
2. **bias-consistent:** for all j and all s we have $\frac{d}{ds} \log y^*(j, s, 0) > \frac{d}{ds} \log \gamma_l(s)$

Under the first property, the ordering of occupations from least to most skill-intensive is unchanged by automation, which greatly facilitates characterization of the optimal assignment. The second property implies that automation either everywhere increases the occupational return to skill, or everywhere decreases it, which simplifies the equilibrium pattern of automation across producers and allows it to be precisely characterized independent of the optimal assignment. These properties are satisfied in the present case due to the assumption that there are only two tasks; in an environment with three or more, they would require additional restrictions on worker productivity and task shares.

Now let $\lambda : s \rightarrow j$ indicate the set of jobs at which at least one s -worker is employed. Given assumption 1, λ will be a strictly increasing, piece-wise differentiable function. There will exist a unique s^* separating automated and non-automated labor types, where s^* may take an interior value or, in the cases where no or all employers automate, may be equal to \underline{s} or \bar{s} . The wage and market tightness functions will satisfy the differential equations

$$\frac{w'(s)}{w(s)} = \begin{cases} \frac{d}{ds} \log y(\lambda(s), s, 0) & s < s^* \\ \frac{d}{ds} \log y(\lambda(s), s, \bar{K}) & s > s^* \end{cases} \quad (24)$$

$$\lambda'(s) = \begin{cases} \frac{y(\lambda(s), s, 0)^{1-\rho} F'(s)}{\beta(\lambda(s))^\rho Y} w(s)^\rho & s < s^* \\ \frac{y(\lambda(s), s, \bar{K})^{1-\rho} F'(s)}{\beta(\lambda(s))^\rho Y} \left(w(s) + \kappa [1 - \alpha(\lambda(s))] r y(\lambda(s), s, \bar{K}) \right)^\rho & s > s^* \end{cases} \quad (25)$$

where $\lambda(\underline{s}) = 0$ and $\lambda(\bar{s}) = 1$. If s^* lies in the interior of $[0, 1]$ then the wage and matching functions will be continuous but not differentiable at s^* , and continuously differentiable for all $s \neq s^*$.

Summarizing, equilibrium in this environment is time and capital allocations t^* and K^* solving (19)-(20), intermediate good bundles $Y^*(j)$ satisfying the final good producer's problem (21), prices $p(j)$ such that zero-profit conditions hold and goods markets clear, and wage and matching functions satisfying the system of differential equations (24)-(25).

3.3.3 Automation, Wages, and Employment

In this section I characterize the short-run and long-run effects of automation, where (in this static setting) I interpret the short-run as describing scenarios where r is sufficiently large that $s^* > \underline{s}$, and the long-run corresponding to the case where $r \rightarrow 0$.

The first property of the model requires no additional proof and is shown in the system of differential equations (24)-(25): in the short-run, the automating technology is only adopted in high-skill jobs. Skilled labor has a comparative advantage at the skilled task, which in turn implies that it is more costly to have skilled workers performing the low-skill task. Automation therefore reduces unit costs by a greater amount when the worker is skilled. Declines in r will tend to lower the automation threshold s^* and, with $w(\underline{s}) > 0$, it is evident that for sufficiently low rental rates we will have $s^* = \underline{s}$ and all jobs will adopt the technology.

A second property is that in the short-run where $s^* > 0$, automation will tend to polarize the wage and employment distributions:

Theorem 1 (short-run polarization). *For $s^* \in (0, 1)$, low-skill employment $\int_0^{\lambda(s^*)} L(j) dj$ is greater under automation. If ρ is sufficiently small and whenever $\rho < 1$, there will exist $j' > \lambda(s^*)$ such that $\int_{j'}^{\infty} L(j) dj / \int_{\lambda^{-1}(s^*)}^{j'} L(j) dj$ is greater under automation. Moreover for $s < s^*$ we will have $w'(s)/w(s)$ strictly smaller under automation, whereas for $s \in (s', \bar{s}]$ for s' sufficiently large we will have $w'(s')/w(s')$ greater under automation.*

For low values of ρ disemployment effects will be strongest among automated jobs 'close' to the threshold skill level s^* , while a smaller share in the unskilled task and greater reductions in unit costs result in better employment outcomes for skilled jobs close to \bar{s} . As ρ becomes

large, the short-run effects on employment are more difficult to characterize because changes to the wage function will have a larger impact on labor demand. With respect to wages, the immediate effect of automation is that low-skill workers shift towards lower-skill jobs, reducing $w'(s)/w(s)$ in the bottom part of the wage distribution. Workers at the upper end of the wage distribution see an increase in $w'(s)/w(s)$ as, even if $\lambda(s)$ shifts downwards (i.e. employment falls in skilled jobs), this will be more than compensated for by the increased return to skill at automated jobs. Assumption 1 and continuity then predict the existence of a region for which $w'(s)/w(s)$ is increasing, with the lower bound of this regional potentially equal to s^* but in all cases smaller than \bar{s} .

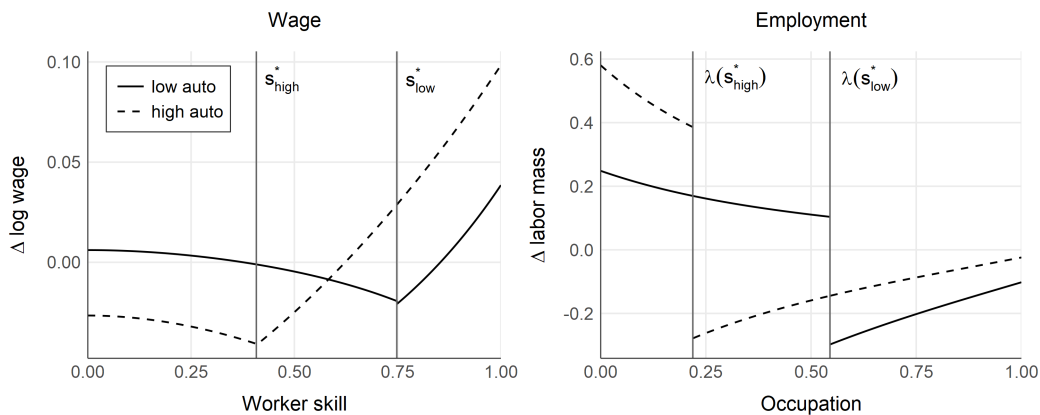


Figure 3.3: Short-Run Effects of Automation

NOTE. Linear task shares where $\gamma_l(s) = \exp(.5s)$, $\gamma_h(s) = \exp(s)$, $\kappa = .5$, and $\rho = .5$. Changes are relative to the case where $s^* = \bar{s}$; wages are demeaned prior to calculating change.

Long-run effects of automation will depend critically on the behavior of the rental rate r , the magnitude of the elasticity of substitution ρ , and on how $t^*(j, s, K)$ varies with s . I impose the condition, consistent with the empirical results shown previously, that t^* is increasing in s .

Assumption 7: $t^*(j, s, K)$ is increasing in s .

A general characterization is not possible due to the lack of a closed-form solution to (24) and (25), but results can be derived for the two cases $\rho \in \{0, 1\}$, which are helpful for developing intuition as to the implications of automation in this environment. First I define $\bar{r} = \sup\{r | s^* = 0\}$ to be the highest rental rate consistent with ‘full’ automation. In the

comparatively simple case where $\rho = 0$, declines in the rental rate below \bar{r} will have no effect on labor demand or on wages and we can show that:

Theorem 2 (long-run effects of automation: $\rho = 0$). *If $r \leq \bar{r}$ and $\rho = 0$, then relative to the case where $r \rightarrow +\infty$, $w'(s)/w(s)$ will be greater for all s and for any $j' \in (0, 1)$ we will have $\int_0^{j'} L(j)ds$ smaller. In this case the wage and matching functions will be independent of $r \in [0, \bar{r}]$.*

The perfect complements case yields an intuitive outcome: low-skill automation reduces employment at low-skill jobs and increases wage inequality. This is consistent with standard notions of labor-substituting automation, and notwithstanding the distinction between short- and long-run effects of automation, this may be said to be the case where the model's predictions are closest to those of the one task-one job environment studied in the extant literature.

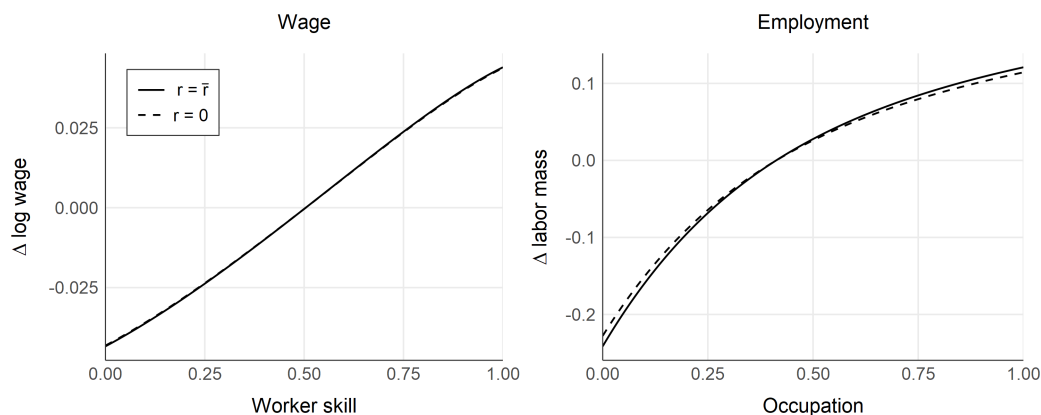


Figure 3.4: Long-Run Effects of Automation, $\rho = 0$

NOTE. Linear task shares where $\gamma_l(s) = \exp(.5s)$, $\gamma_h(s) = \exp(s)$, $\kappa = .5$, and $\rho \approx 0$. Changes are relative to the case where $s^* = \bar{s}$; wages are demeaned prior to calculating change.

Long-run effects of automation become considerably more complicated as we move away from the limiting case where $\rho = 0$, because there are two offsetting effects. On the one hand automation increases the intensity of the skilled task and therefore tends to raise wages of skilled workers, leading to substitution away from skilled occupations. On the other hand the cost of automation is smaller relative to wage costs for higher s , resulting in substitution in the opposite direction. The result is that while wage inequality will tend to increase overall,

the implications for employment are ambiguous.

Theorem 3 (long-run effects of automation: $\rho = 1$). *If $r \leq \bar{r}$ and $\rho = 1$, then for any s we will have $w(s)/w(\underline{s})$ greater under automation. If $r = 0$ then there will exist a $j' > 0$ such that $\int_0^{j'} L(j)ds$ is also greater under automation. Finally, comparing any two rental rates $r' < r''$, we will have $w'(s)/w(s)$ smaller under r' and, for any $j' \in (0, 1)$ we will have $\int_0^{j'} L(j)ds$ larger.*

When $r = \bar{r}$ the effect of automation on employment is unclear, but further declines in r will tend to both reduce wage inequality and increase low-skill employment, and as r goes to zero the overall effect will be to raise employment in low-skill occupations. Hence technological disemployment will depend critically on the long-run behavior of r and the value of ρ .

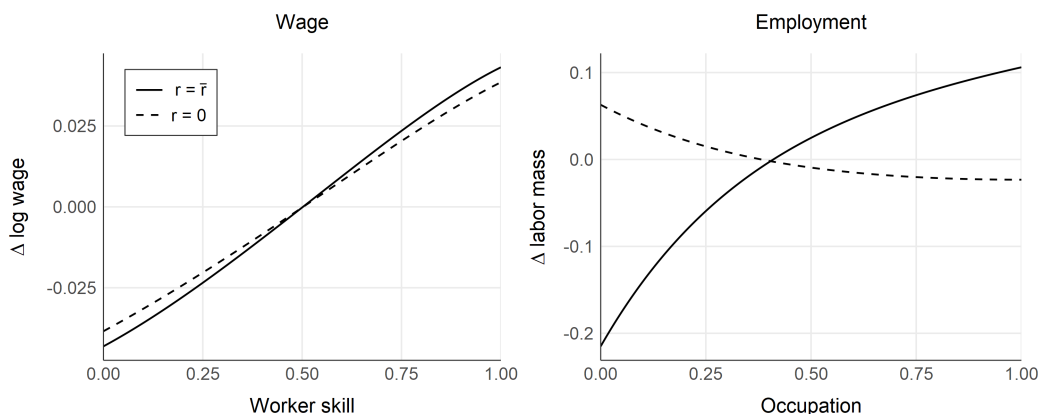


Figure 3.5: Long-Run Effects of Automation, $\rho = 1$

NOTE. Linear task shares where $\gamma_l(s) = \exp(.5s)$, $\gamma_h(s) = \exp(s)$, $\kappa = .5$, and $\rho = 1$. Changes are relative to the case where $s^* = \bar{s}$; wages are demeaned prior to calculating change.

For values of ρ greater than one characterization of the wage and matching functions becomes difficult. In general, larger values of ρ and smaller values of r will be associated with greater compression of the wage distribution and more employment in low-skill occupations. The extent to which automation affects wages rather than prices will also depend on the functional forms of γ_l and γ_h , which together with task shares α determine how easy it is to substitute one skill type for another at a given job, and hence the response of occupational labor supply.

3.3.4 Continuous Model

The model developed thus far is of limited quantitative usefulness, as empirical adoption patterns are not binary at the occupation level and there is no observed analogue of the automation threshold s^* . Intra-occupational heterogeneity must be incorporated if the model is to yield quantitatively meaningful predictions. Adding heterogeneity will tend to weaken the theoretical results shown in the previous section, but they will continue to hold *approximately* so long as the proportion of jobs automated is increasing in j . In this section I allow for a simple form of intra-occupation heterogeneity, and derive the equilibrium for a continuous model that will form the basis for the quantitative results shown below.

I assume that producers face idiosyncratic shocks ϵ to their rental cost, with total costs equal to ϵr per unit of capital. I further assume that these shocks are realized only after hiring labor, a simplification which has the implication that all j -firms will continue to hire the same worker type. I assume that ϵ is drawn from a continuously differentiable distribution $G(\epsilon)$, and that producers are constrained to a unit of output, so that $G(\epsilon)$ also gives the proportion of output attributable to producers with costs ϵ or lower.⁸⁰ I assume that workers are paid wages upon hiring and hence, firms have no incentive to leave the market upon discovering their rental costs as they can always choose not to automate and still receive positive revenue. Capital costs paid by automating firms will then be

$$\epsilon r K^*(j, s, z) = \kappa \frac{1 - \alpha(j)}{z} y^*(j, s, \bar{K}) ,$$

and producers will automate whenever

$$w(s) > \epsilon r \gamma_l(s) .$$

I assume that G has full support over \mathbb{R}^{++} , in which case for each s there will exist a $\epsilon^*(s)$ such that producers employing s automate whenever $\epsilon > \epsilon^*(s)$. The policy functions and optimal assignment are similar to before and so I provide them in the appendix but omit them here.

⁸⁰Because firms will now enjoy greater or lesser rents depending on the value of ϵ , some restriction on scale is required for the equilibrium to be well-defined

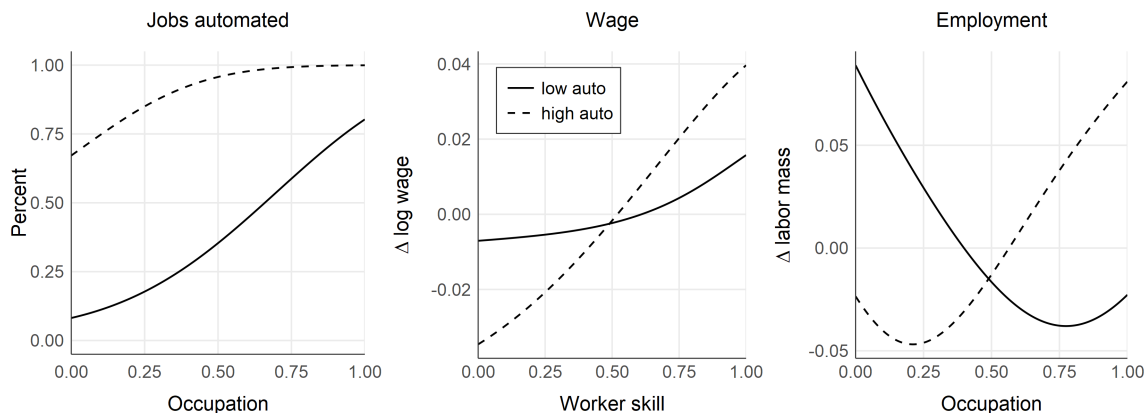


Figure 3.6: Short-Run Effects of Automation, Continuous Model

NOTE. Linear task shares where $\gamma_l(s) = \exp(.5s)$, $\gamma_h(s) = \exp(s)$, $\kappa = .5$, $\rho = .5$, and ϵ log-linear with mean 0 and standard deviation .1. Changes are relative to the case where $s^* = \bar{s}$; wages are demeaned prior to calculating change.

Although equilibrium in the continuous model is more difficult to characterize, two general predictions are preserved from static case. The first prediction is that $\epsilon^*(s)$ is increasing: a larger proportion of jobs will be automated in high-skill occupations. The reason for this is the same as before, that skilled workers have a comparative advantage at skilled tasks and therefore it is more costly when they must spend time on the unskilled task. The second prediction is that, holding r fixed, an increase in the automation threshold ϵ^* will tend to reduce occupational employment; whereas holding ϵ^* fixed, a decrease in r will tend to increase employment in occupations with a high ϵ^* . A lower cost of capital will tend to reduce expected costs in highly automated occupations and incentivize job creation, whereas an increase in the proportion of jobs being automated will result in less demand for labor. Hence as before, the effect of a change in r on j -employment will depend on the relative magnitude of these two effects.

3.4 Qualitative Evaluation

In this section I test the main qualitative predictions of model developed above, one concerning the cross-sectional distribution of technology adoption, and the other the effects of declining capital costs on employment. Although the model also delivers predictions regarding wages, these are difficult to test retrospectively, because in the first place they cannot be

precisely characterized without restricting the model parameters, and in the second, wage data in the BIBB surveys are highly aggregated prior to the 2005/06 survey. I therefore defer consideration of the model’s wage predictions until the quantitative section of this paper.

3.4.1 Time Path of Automation

The first qualitative prediction of the model is that automation of a low-skill task is “top-down”, and at any given point in time is more likely to take place in high-skill occupations. In terms of model primitives:

Prediction #1: the automated proportion of jobs $G(\epsilon^*(s))$ is increasing in s .

This prediction is straightforward to test in the context of the BIBB surveys. Numerous case studies suggest that computers reduce the need for labor at low-skill tasks, for example by facilitating automation of manual tasks or allowing for faster execution of simple calculation and data manipulation tasks. To analyze computer adoption patterns, I calculate for each KLDB occupation the percentage of workers using a computer on the job and the mean log wage associated with the occupation. Taking a similar approach to [Acemoglu and Autor \(2011\)](#), I use wage data to construct a time-invariant ranking of occupations into percentiles, allowing for a more straightforward comparison across survey years. Results are shown in figure 3.7, and are consistent with model predictions. Between 1979 and 1999, PC adoption occurs primarily in high-paying and presumably high-skill occupations. By 1999 adoption rates are close to 100% in the highest-paying quartile of jobs, and only during the years after 1999 is there substantial PC adoption in the lower half of the occupation wage distribution.

One concern with these results is that, as a general purpose technology, PCs may substitute for labor at a variety of tasks, and do not necessarily reflect a distinct technological innovation as assumed in the model. I therefore consider two use cases for PCs - CNC machining in production jobs, and word processing for clerical tasks - for which data is available during the survey years 1986-1999. These applications are not relevant for all jobs, and so I therefore omit observations for which workers do not report using machinery (computer-controlled or not) and typing equipment (word processors or typewriters). Figure 3.8 shows similar adoption patterns to that for PCs as a whole, with early-stage use concentrated in

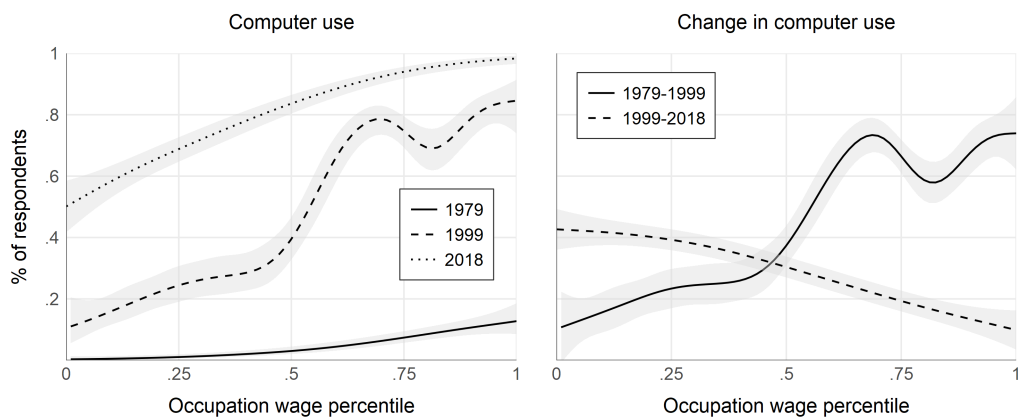


Figure 3.7: Computer Use by Occupation Wage Percentile

NOTE. Log wage and PC use averaged by 1988 KLDB occupation for 1979-1999, and by 1992 KLDB occupation for 1999 (left panel) and 1999-2018 (right panel). Percentiles are time-invariant and reflect 1979 and 1999 wages. Shaded regions indicate 95% confidence intervals.

high-paying jobs. A caveat is that most adoption of word processors occurs between two panels in the 1990's, and consequently I do not observe intermediate levels of adoption.



Figure 3.8: Technology Penetration by Wage Percentile

NOTE. First panel: use of CNC machines as percent of workers who report using heavy machinery. Second panel: use of word processors as percent of workers who report using typing equipment. Occupation-mean wages and weights calculated from equipment-using subsample. Shaded regions indicate 95% confidence intervals.

Summarizing these results, PC adoption over 1979-2018 exhibited a top-down pattern consistent with model predictions. While intuitive, this result is not explained by models of computerization based on capital-skill complementarity, which abstract from job-level interactions of labor and technology. It is also inconsistent with the simple labor-substitution

framework of Acemoglu and Autor, in which technology adoption is associated with low-skill or routine jobs - which, given the empirical results shown in this paper, would suggest a “bottom-up” pattern of automation. The inability of extant models to predict adoption patterns is important, because these patterns can influence wage and employment trends *irrespective of the long-run effects of the technology*. In other words, the skill-bias of technological change at any point in time is likely to reflect both adoption patterns and the fixed characteristics of the technology, and disentangling these two factors is critical for predicting the long-run effects on labor markets.

3.4.2 Employment Effects of Automation

The second prediction of the model developed in this paper is that automation reduces contemporaneous employment, but will tend to increase employment as the cost of the technology declines further.

Prediction #2: employment is decreasing in $G(\epsilon^*(s))$ given r , and decreasing in r for occupations with relatively high $G(\epsilon^*(s))$.

I test this prediction through a difference-in-difference approach relating occupational growth rates to levels of, and changes to, occupational PC use. A key limitation is that the BIBB surveys are not designed to measure occupational employment. Although the surveys are randomly sampled, non-response rates are not random and in addition, sample sizes are not large enough to ensure precise measures at the 3-digit occupational level. Survey weights correct for demographics but not occupation, and I find that in practice the weights introduce additional noise when calculating occupation growth rates. I therefore opt for the simplest approach and use raw counts to calculate occupational employment.⁸¹ Results for the pooled panel regressions are shown in table C.5, and indicate that high levels of PC use are associated with employment growth, while increases in the rate of PC are generally accompanied by declining employment. This pattern is robust to controlling for occupation-mean wage. Inclusion of occupation fixed effects results in a somewhat stronger pattern for the years 1979-1999, but over the 1999-2018 period PC use contains no additional information over

⁸¹Future plans for this project include the use of administrative employment data to calculate occupational wage and employment data.

the occupational effects, which may reflect the fact that there is less cross-sectional and time variation in PC use after 2006.

Table 3.5: Regression: Change in Log Occupational Employment Share

Independent variable	<i>Regression coefficients</i>							
	Years 1979-1999				Years 1999-2018			
PC use	.268 (.077)	.226 (.085)	.497 (.174)	.503 (.172)	.472 (.089)	.296 (.116)	-.043 (.293)	-.117 (.300)
Δ PC use	-.462 (.186)	-.472 (.185)	-.520 (.192)	-.519 (.192)	-.378 (.133)	-.296 (.131)	-.097 (.177)	-.123 (.177)
Log(wage)		.079 (.085)		-.064 (.272)		.223 (.084)		.448 (.215)
Occup. FE			X	X			X	X
Observations	830	830	830	830	874	874	874	874

NOTE. Difference-in-difference regression with occupational employment share as the dependent variable. Employment shares calculated from raw survey counts. All regressions include year fixed effects.

To further assess the robustness of these results, in table 3.6 I show the estimated coefficients when considering only consecutive survey panels. The predicted employment patterns are present in four out of six regressions. Results are insignificant for the remaining two regressions, one of which (1992-1999) is compromised by changes to the wording and format of the survey question concerning PC use. As in the previous table, the magnitude of the coefficients is large: a 10% point increase in PC use is associated initially with a 2-7% decline in occupational employment share, while a pre-existing PC use rate of 10% is associated with a comparable to somewhat smaller increase in share.

Table 3.6: Regression: Change in Log Occupational Employment Share, By Year

Indep. Var.	1979-86	1986-92	1992-99	1999-06	2006-12	2012-18
PC use	.681 (.183)	.205 (.099)	.128 (.099)	.704 (.135)	-.005 (.120)	.530 (.170)
Δ PC use	-1.396 (.392)	-.470 (.195)	.097 (.231)	-.380 (.174)	.135 (.222)	-.635 (.244)
Observations	277	277	276	291	292	291

NOTE. Difference-in-difference regression with occupational employment share as the dependent variable. Employment shares calculated from raw survey counts.

These results suggest an intuitive explanation for why PCs have, in various cases, both

complemented and substituted for occupational labor. Past studies such as [Autor, Levy, and Murnane \(2002\)](#) have made note of these disparate effects, and attributed them to fundamental (i.e. unexplained) differences in how technology interacts with different skill types, but the patterns shown in table C.5 are present even if one restricts attention to skilled occupations.⁸² There is no mystery if only portions of jobs are automated, because this naturally introduces a distinction between the marginal (labor-substituting) and the long-run (labor-complementing) effects of automation.

3.5 Quantitative Analysis

In this section I estimate the continuous model of partial automation developed above, and derive long-run predictions for wages and employment. I begin with an overview of the estimation procedure. Next I consider the model’s long-run predictions regarding the wage and employment distributions. I close the section with results characterizing the capital-labor and labor-labor elasticities of substitution in this environment.

3.5.1 Estimation

Estimation is performed on the 2017-18 survey panel. As a first step, I reduce the dimensionality of the task space from 6 task variables (i.e. those used in the empirical analysis) to 2 variables, using a factor analysis approach.⁸³ Of the two composite variables, one loads principally on the routine task characteristics and the other on the non-routine characteristics. For each observation I divide factor scores by their sum and take the resulting values as measures of workers’ time allocation t^* and $1 - t^*$.

I assume that the task production functions $\gamma_l(s)$ and $\gamma_h(s)$ are exponential functions taking the form $\gamma_k(s) = \exp(G(s)\gamma_k)$ with the normalization $\gamma_l = .5$. Once a value for γ_h is fixed, task shares can be estimated by solving the worker’s time allocation policy function $t^*(j, s, K)$ for $\alpha(j)$, taking the expectation over all workers in j , substituting empirical values

⁸²Tabulated results for high-wage occupations are provided in the appendix, and are qualitatively and quantitatively similar to those for the full sample.

⁸³Factor analysis is appropriate in this case because of the assumption of an underlying two-task structure; by preserving only the off-diagonal elements of the covariance matrix, factor analysis effectively ignores the idiosyncratic variance associated with individual survey questions.

for t^* and the percentage of workers using PCs $G(\epsilon^*(s))$, and then numerically solving the resulting quadratic. The empirical values of t^* and $G(\epsilon^*(s))$ are averaged by occupation wage percentile and are generally increasing but, in order to ensure that equilibrium monotonicity conditions are met, I interpolate the empirical values subject to a non-decreasing constraint. The return to skill function $G(s)$ is then estimated by imposing the normalization $\lambda(s) = s$, solving the differential equation describing the wage function for $G'(s)$, and solving this system using empirical wages. The empirical and predicted distributions are shown for reference in figure 3.9.

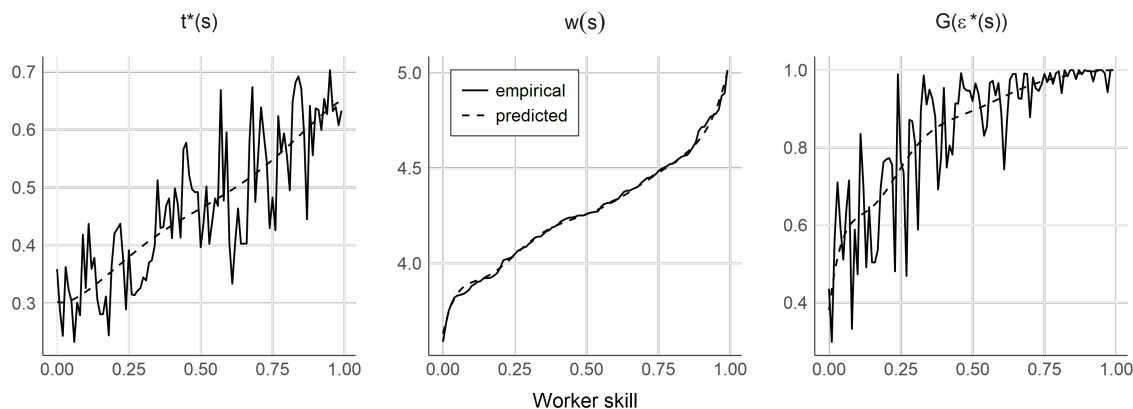


Figure 3.9: Model Goodness-of-Fit

The aggregate model parameters are estimated by using an indirect inference approach. The distribution $G(\epsilon)$ is assumed to be log-normal with standard deviation σ , and values of $\{r, \kappa, \sigma, \rho, \gamma_h/\gamma_l\}$ are obtained by minimizing the distance between empirical and model-predicted outcomes. The technological parameters r and σ are chosen to minimize the (squared) distance between the empirical and model-predicted distributions of PC use by wage percentile $G(\epsilon^*(s))$. To estimate κ , I match the empirical and model-predicted coefficients from a diff-in-diff regression of occupation-mean skilled task share $\mathbb{E}[t^*|s]$ on PC use $G(\epsilon^*(s))$. The elasticity of demand ρ can only be obtained by comparing outcomes over time, so I implement a second diff-in-diff regression of occupational employment shares on (1) PC use and (2) the time trend in PC use, and minimize the distance between the predicted and empirical coefficients on PC use. For each of the two regressions, I use data from the 2006-2018 surveys as they are identical in terms of questions and coding. Finally, the task

skill differential γ_h is not separately identified from task shares $\alpha(j)$ in cross-sectional data, but will determine the wage response to automation. I therefore estimate γ_h by matching predicted and observed changes to the 90/10 wage ratio over 1979-2018.⁸⁴ Estimates are given in table B.4.

Table 3.7: Aggregate Parameter Estimates

Parameter	Definition	Value	Obj. function	Survey years
r	rental rate	5.40	$(G^*(s) - \hat{G}^*(s))^2$	2018
σ	std. dev. $G(\epsilon)$.157	$(G^*(s) - \hat{G}^*(s))^2$	2018
ρ	elasticity of subs.	1.72	$(\beta_{G^*}^L - \hat{\beta}_{G^*}^L)^2$	2005-18
κ	automation feas.	.44	$(\beta_{G^*}^{t*} - \hat{\beta}_{G^*}^{t*})^2$	2005-18
γ_h/γ_l	task skill diff.	2.82	$(\Delta WR_{10}^{90} - \Delta \hat{W}R_{10}^{90})^2$	1979, 2018

The most problematic aspect of model estimation is identification of ρ and γ_h . Changes over time to employment shares and wages are likely to reflect a number of factors, some of which may also be associated with PC use. The value of 1.72 is high relative to the “best guess” range of 1.4-1.5 proposed by [Johnson \(1997\)](#) for demand elasticity of substitution across skill types, but here the estimation procedure is influenced by other parameter values (in particular κ and γ) and interpretation of this parameter is therefore difficult. The value of γ_h determines the elasticity of labor supply, and is similarly unconstrained by past studies. I therefore consider alternative values of ρ and γ_h later in this section. For similar reasons the estimation of κ is also potentially problematic, but as κ is primarily a scaling parameter, over- or under-estimation should not bias comparison between different points in time (i.e. different values of r) that will be the focus of the analysis in this section.

3.5.2 Employment, Wages, and Technology Adoption

The model-predicted wage and employment distributions are shown in figure 3.10. The second panel shows how automation exerts a staggered effect on labor shares: employment declines initially in high-skill occupations, followed by middle-skill and finally low-skill oc-

⁸⁴Wage data for the 1979 survey are aggregated into bins, and so prior to calculations I perform similar aggregation on the wage data from the 2018 survey after adjusting for purchasing power using data on [purchasing power comparisons of historical monetary amounts](#) obtained from Deutsche Bundesbank. Wage ratios are calculated from occupation-mean log wage, and the model-predicted change is obtained by finding the rental rate for which aggregate PC adoption is equal to the observed 1979 value.

cupations. Wages initially polarize, with the 90/50 ratio increasing at a faster rate than the 50/10 rate. As technology adoption approaches 100% for all j , polarization gives way to a general increase in wage dispersion. The main implication of partial automation, however, is the trajectory of wages and employment after the technology has been everywhere adopted. As r approaches zero, employment increases in low-skill jobs and wage inequality declines. The model predicts that employment in low-skill occupations ultimately increases, although automation tends also to increase the skill-intensity of these jobs in that workers spend a greater proportion of their time on the skilled task.

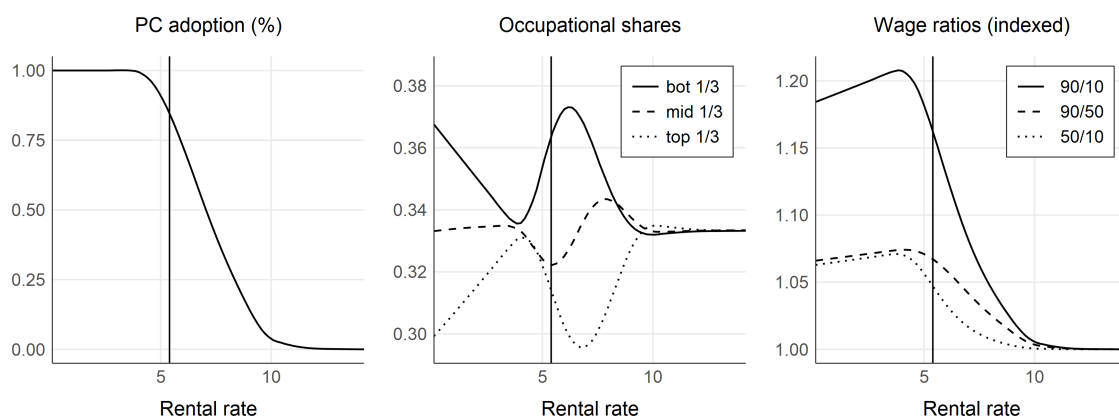


Figure 3.10: Technology Adoption and Labor Outcome

NOTE. Smoothed model output over a grid of rental rates r . Occupational groups correspond to the j -intervals $[0, 1/3)$, $[1/3, 2/3)$, and $[2/3, 1]$.

The distributional effects of automation are shown in greater detail in figure 3.11. Wage variance peaks as PC adoption approaches 100% (i.e. r approaches \bar{r}), and then declines slightly with further decreases in r . The small magnitude of this decline reflects the fact that the wage effects of automation are largely associated with the intensive task margin: automation increases the return to skill within a given automation because it results in workers spending a greater proportion of their time on the skilled task. The extensive margin - changes to the sorting of workers across jobs - is less important from a wage standpoint, with the caveat that job transitions in this environment are friction-less. If there are costs to switching jobs, for example due to non-transferable human capital or to search frictions, then the extensive margin is likely to exert a stronger effect on wages.

Summarizing these results, the general prediction of the quantitative model is that au-

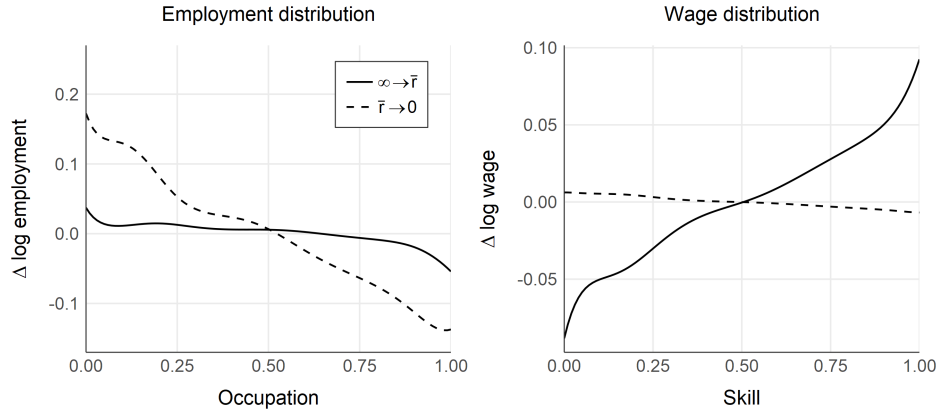


Figure 3.11: Short-Run and Long-Run Distributional Effects

NOTE. The value \bar{r} is defined as the highest value of r for which $G(\epsilon^*(s)) \geq .99$ for all s .

tomation of the low-skill task results in a persistent increase in wage inequality but non-persistent employment effects at the occupation level. In many respects this is a negative result. The undesirable outcome - wage inequality - is not expected to abate in the long-run, while the short-run nature of occupational disemployment suggests that policies facilitating job transitions may be counterproductive. Overall, figures 3.10 and 3.11 indicate that long-run effects of automation are important, but it is clear that the implications of this result will depend on the nature and magnitude of the labor market frictions affecting occupational choice.

3.5.3 Capital-Labor Substitution

The capital-labor elasticity of substitution $\partial \log L / \partial \log r$ is a common object of study in the skill-bias literature. A distinction is often made between technologies that are complementary with skill (i.e. “capital-skill complementarity”) and those that substitute for low-skilled labor, while Autor, Levy, and Murnane (2003) make the case that both of characteristics are true of computers. A major departure of the model studied in this paper is that whether capital and labor are gross substitutes is not the result of assumptions imposed on technology, but is endogenous to occupational task structure and the rental rate. In figure 3.12 I plot cross-price elasticity labor demand over skill types and rental rates. In the early stages of computerization, capital *substitutes* for skilled labor, although it tends also to increase wage inequality as skilled workers devote a greater share of their time to the skilled task, and hence

the return to skill rises. For intermediate values of the rental rate, capital complements labor in occupations at the extremes of the distribution, and substitutes for those in the middle. And in the final stages of adoption, capital substitutes for low-skill labor.

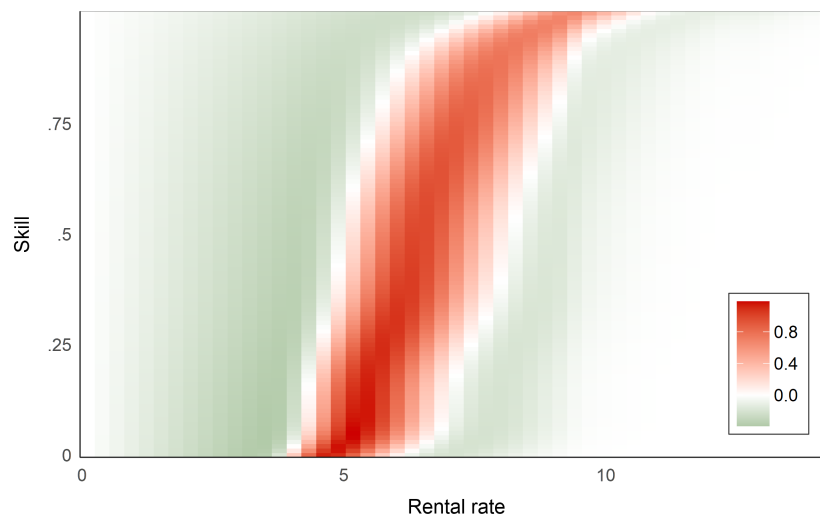


Figure 3.12: Cross-Price Elasticity of Labor Demand: Rental Rate

NOTE. Partial derivative of $\log L(j, s, K)$ with respect to $\log r$, numerically calculated over rental rates and worker skill s .

What determines this elasticity? It will generally be positive when changes to r have a large effect on $G(\epsilon^*(s))$ - when adoption is responsive to costs - and negative when $G(\epsilon^*(s))$ is large and relatively insensitive to the rental rate, which will be the case once the technology has been adopted in most jobs within an occupation. Note that for any given skill level, there exist rental rates for which capital is a gross substitute *and* rental rates for which capital complements labor. This result illustrates a main implication of task-level automation, which is that labor outcomes will depend critically on the time frame under consideration, and therefore the time-path of capital costs will be important for predicting labor outcomes in the long-run. These outcomes will also depend on the distribution of tasks across and within occupations - a point that is likely to hold with even more force under more general specifications of the model, in which more than two tasks are allowed.

3.5.4 Supply and Demand Elasticities

Two key elasticities in the model are the demand elasticity of substitution ρ , and the elasticity of occupational labor supply that, given the functional assumptions on match production, will depend on the ratio γ_h/γ_l . Together, these elasticities determine the response of wages and employment to automation, and it is for this reason that they are difficult to identify in practice. Historical changes to wages and employment are likely to reflect a multitude of other factors and therefore to be unreliable as measures of the effects of automation. In this section I consider the sensitivity of model predictions to alternative values of these parameters.

I begin with the demand elasticity ρ , which has some parallel to the labor-labor elasticity studied by previous authors, and typically formalized as the elasticity of substitution across skill types (typically education). Recent estimates of this elasticity include those by [Ciccone and Peri \(1.5, 2005\)](#) and [Autor, Katz, and Kearney \(1.57, 2008\)](#), and values in the neighborhood of 1.5 are common in this literature; but [Gechert et al. \(2021\)](#) suggest that published estimates are biased, and that the actual value of the labor-labor elasticity is as low as 1. In this paper, the elasticity of substitution across skill types is not a technological parameters but will depend on γ_h/γ_l as well as ρ , and so it is *ex ante* unclear what a plausible range for ρ would be. I therefore choose the somewhat arbitrary boundary points $\rho \in \{1, 2\}$, for which short-run and long-run employment predictions are shown in figure 3.13. When ρ is




Figure 3.13: Varying the Elasticity of Labor Demand

NOTE. The value \bar{r} is defined as the highest value of r for which $G(\epsilon^*(s)) \geq .99$ for all s . Occupational shares are re-estimated conditional on ρ .

larger, dis-employment from automation tends to be quickly reversed, as affected occupations enjoy greater benefits from further declines in r . For this reason, at the point \bar{r} where adoption is complete, occupational labor shares are largely unchanged. As the rental rate continues to fall, however, low-skill occupations see growth, and this post-adoption effect is much stronger than in the case where ρ takes on a smaller value. Therefore a wide range of long-run employment outcomes are defensible, given uncertainty as to the true value of ρ . Wage predictions, on the other hand, are unaffected by this parameter and instead reflect the non-parametric skill productivity term $H(s)$ and the task skill differential γ_h/γ_l , discussed next.

Turning to the task skill ratio γ_h/γ_l , a larger value of this ratio will imply that skilled workers have a stronger comparative advantage at high- j occupations, and consequently occupational labor supply will be less elastic. For this reason wages will exhibit a stronger response to automation. To show this I plot results in figure 3.14 for two alternative parameter specifications where $\gamma_h/\gamma_l = 2$ and $\gamma_h/\gamma_l = 4$. The choice of these points is again somewhat arbitrary, but is sufficient for a demonstration of the sensitivity of model predictions. The first column of figure 3.14 shows occupational employment shares, from which we can see that a larger value of γ_h is associated with greater divergence over time in the share of high-skill and low-skill occupations. This in turn reflects the greater increase in the relative wages of skilled workers, indicated in the second column of the figure. Skilled workers experience a larger increase in wages in this case because the comparative advantage profile is steeper, and therefore as automation increases the portion of time spent on the skilled task, the return to skill increases by a greater amount. The long-run decline in wage inequality is somewhat greater in this case: from its peak, the 90/10 wage ratio falls by approximately 15% as capital costs fall towards zero. Uncertainty about the value of γ_h therefore suggests a wide range of plausible values for long-run employment shares, as well as some margin of error for wages, although in all cases the estimated model predicts only weak reversals of wage inequality in the long-run.

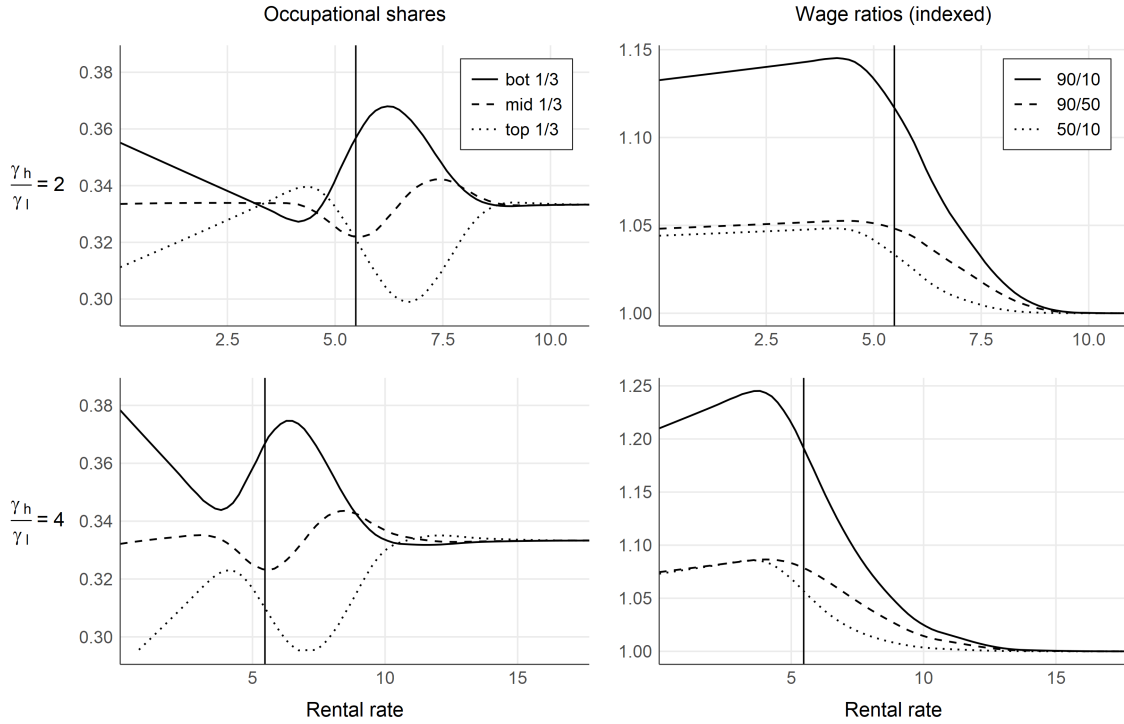


Figure 3.14: Varying the Elasticity of Labor Supply

NOTE. Smoothed model output over a grid of rental rates r . All model parameters other than ρ are re-estimated conditional on γ .

3.6 Conclusion

In this chapter I have presented theoretical and quantitative results on the labor market effects of automation, when technology does not eliminate jobs but replaces labor at particular tasks within jobs. I show descriptively using German survey data that the vast majority of jobs contain both routine and non-routine task content, and that computerization over the 1979-2018 period was associated with intra-occupational changes in the frequency of routine content. The theoretical model developed in this paper predicts that in the short-run, automation of a low-skill task reduces employment as predicted by standard models; but in the long-run, as technology costs continue to fall, employment in low-skill occupations recovers and it is possible for long-run job growth to fully offset short-term losses. I show that the experience of West Germany is consistent with short-run model predictions concerning occupational computerization patterns and employment dynamics. A structurally estimated version of the model predicts that while increases in wage inequality associated

with automation will be persistent, employment in middle-skill and low-skill jobs will recover as information technology costs continue to decline.

These results have two broad implications. First, the “traditional” effect of labor-substituting technology - dis-employment - is not an inevitable long-run outcome when automation is gradual, and jobs are not fully but only partially automated. More catastrophic predictions such as those in [Frey and Osborne \(2017\)](#) are generally based on the assumption of full automation, for which the empirical and historical support is weak. Second, because job outcomes in the immediate aftermath of automation are substantially different from job outcomes in the long-run, policies focused on worker retraining and occupational upskilling will tend to be more costly over a long time-frame. The extent to which this is true will depend on how large and how fast are the declines in technology costs, and on the elasticity of labor supply and demand across occupations. On the other hand, policies intended to dampen job losses, such as employer subsidies, are likely to be more effective and less costly than they would be in the case of full automation.

A key shortcoming of this study is that I abstract from frictions affecting occupational labor supply, such as costs associated with retraining and re-schooling, barriers to entry such as licensure, and search frictions associated with finding a new job and, in many cases, a new employer. The effects of such frictions are unclear, as they will affect the response of both wages and employment to automation; but it is reasonable to expect that they would lead to greater variability of wages in the short-run. More broadly, the results in this paper reinforce the notion that factors influencing occupation transitions are key to understanding the quantitative effects of, and the policy trade-offs associated with automation.

References

- [1] Abowd, John, Francis Kramarz, Paul Lengeremann, Kevin L. McKinney, and Sébastien Roux (2012). “Persistent inter-industry wage differences: rent sharing and opportunity costs.” *IZA Journal of Labor Economics* 1, no. 1: 7.
- [2] Abowd, John, Francis Kramarz, and David Margolis (1999). “High wage workers and high wage firms.” *Econometrica* 67, no. 2: 251-333.
- [3] Acemoglu, Daron, and David Autor. “Skills, tasks and technologies: Implications for employment and earnings.” In *Handbook of labor economics*, vol. 4, pp. 1043-1171. Elsevier, 2011.
- [4] Acemoglu, Daron, and Pascual Restrepo (2018). “The race between man and machine: Implications of technology for growth, factor shares, and employment.” *American Economic Review* 108, no. 6: 1488-1542.
- [5] Acemoglu, Daron, and Pascual Restrepo (2020). “Robots and jobs: Evidence from US labor markets.” *Journal of Political Economy* 128, no. 6: 2188-2244.
- [6] Aghion, Philippe, Benjamin Jones, and Charles Jones (2019). 9. *Artificial Intelligence and Economic Growth*. University of Chicago Press.
- [7] Akerman, Anders, Elhanan Helpman, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen Redding (2013). “Sources of wage inequality.” *American Economic Review* 103, no. 3: 214-19.
- [8] Ales, Laurence, Musab Kurnaz, and Christopher Sleet (2015). “Technical change, wage inequality, and taxes.” *American Economic Review* 105, no. 10: 3061-3101.
- [9] Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser (2018). “Firms and the decline in earnings inequality in brazil.” *American Economic Journal: Macroeconomics* 10, no. 1: 149-89.
- [10] Andrews, Martyn, Len Gill, Thorsten Schank, and Richard Upward (2008). “High wage workers and low wage firms: negative assortative matching or limited mobility bias?.”

Journal of the Royal Statistical Society: Series A (Statistics in Society) 171, no. 3: 673-697.

- [11] Arntz, Melanie, Terry Gregory, and Ulrich Zierahn (2017). “Revisiting the risk of automation.” *Economics Letters* 159 (2017): 157-160.
- [12] Atalay, Enghin, Phai Phongthientham, Sebastian Sotelo, and Daniel Tannenbaum (2020). “The evolution of work in the United States.” *American Economic Journal: Applied Economics* 12, no. 2: 1-34.
- [13] Autor, David, and David Dorn (2013). “The growth of low-skill service jobs and the polarization of the US labor market.” *American Economic Review* 103, no. 5: 1553-97.
- [14] Autor, David, David Dorn, and Gordon Hanson (2013). “The China syndrome: Local labor market effects of import competition in the United States.” *The American Economic Review* 103, no. 6: 2121-68.
- [15] Autor, David, David Dorn, Lawrence Katz, Christina Patterson, and John Van Reenen (2020). “The fall of the labor share and the rise of superstar firms.” *The Quarterly Journal of Economics* 135, no. 2: 645-709.
- [16] Autor, David, and Michael Handel (2013). “Putting tasks to the test: Human capital, job tasks, and wages.” *Journal of Labor Economics* 31, no. S1: S59-S96.
- [17] Autor, David, Lawrence Katz, and Melissa Kearney (2008). “Trends in US wage inequality: Revising the revisionists.” *The Review of Economics and Statistics* 90, no. 2: 300-323.
- [18] Autor, David, Frank Levy, and Richard Murnane (2002). “Upstairs, downstairs: computers and skills on two floors of a large bank.” *ILR Review* 55, no. 3: 432-447.
- [19] Autor, David, Frank Levy, and Richard Murnane (2003). “The skill content of recent technological change: An empirical exploration.” *The Quarterly Journal of Economics* 118, no. 4: 1279-1333.

- [20] Autor, David, and Anna Salomons (2018). “Is automation labor-displacing? Productivity growth, employment, and the labor share.” No. w24871. National Bureau of Economic Research.
- [21] Bachmann, Ronald, Merve Cim, and Colin Green (2019). “Long-Run Patterns of Labour Market Polarization: Evidence from German Micro Data.” *British Journal of Industrial Relations* 57.2: 350-376.
- [22] Bagger, Jesper, and Rasmus Lentz (2019). “An empirical model of wage dispersion with sorting.” *The Review of Economic Studies* 86, no. 1: 153-190.
- [23] Bagger, Jesper, Kenneth L. Sorensen, and Rune Vejlin (2013). “Wage sorting trends.” *Economics Letters* 118, no. 1: 63-67.
- [24] Bartel, Ann, Casey Ichniowski, and Kathryn Shaw (2007). “How does information technology affect productivity? Plant-level comparisons of product innovation, process improvement, and worker skills.” *The quarterly journal of Economics* 122, no. 4: 1721-1758.
- [25] Barth, Erling, Alex Bryson, James Davis, and Richard Freeman (2016). “It’s where you work: Increases in the dispersion of earnings across establishments and individuals in the United States.” *Journal of Labor Economics* 34, no. S2: S67-S97.
- [26] Becker, Gary (1973). “A theory of marriage: Part I.” *Journal of Political economy* 81, no. 4: 813-846.
- [27] Bellman, Lisa, Ben Lochner, Stefan Seth, and Stefanie Wolter (2020). “AKM effects for German labour market data.” *FDZ-Methodenreport, IAB FDZ*.
- [28] Bessen, James (2016). “How computer automation affects occupations: Technology, jobs, and skills.” *Boston Univ. school of law, law and economics research paper* 15-49.
- [29] Blossfeld, Hans-Peter (1985). “Berufseintritt und Berufsverlauf.” *Mitteilungen aus der Arbeitsmarkt-und Berufsforschung* 18, no. 2: 177-197.

- [30] Dal Bo, Ernesto, Frederico Finan, and Martun Rossi (2013). “Strengthening state capabilities: The role of financial incentives in the call to public service.” *The Quarterly Journal of Economics* 128, no. 3: 1169-1218.
- [31] Bonhomme, Stephane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler (2020). “How Much Should we Trust Estimates of Firm Effects and Worker Sorting?” NBER Working Paper No. 27368, June.
- [32] Bonhomme, Stephane, Thibaut Lamadon, and Elena Manresa (2019). “A distributional framework for matched employer employee data.” *Econometrica* 87, no. 3: 699-739.
- [33] Burda, Michael, and Michael Funke (1993). “German trade unions after unification—Third degree wage discriminating monopolists?.” *Weltwirtschaftliches Archiv* 129, no. 3: 537-560.
- [34] Burstein, Ariel, and Jonathan Vogel (2017). “International trade, technology, and the skill premium.” *Journal of Political Economy* 125, no. 5: 1356-1412.
- [35] Cappelli, Peter (1993). “Are skill requirements rising? Evidence from production and clerical jobs.” *ILR Review* 46, no. 3: 515-530.
- [36] Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline (2018). “Firms and labor market inequality: Evidence and some theory.” *Journal of Labor Economics* 36, no. S1: S13-S70.
- [37] Card, David, and John DiNardo (2002). “Skill-biased technological change and rising wage inequality: Some problems and puzzles.” *Journal of labor economics* 20, no. 4: 733-783.
- [38] Card, David, Jorg Heining, and Patrick Kline (2013). “Workplace heterogeneity and the rise of West German wage inequality.” *The Quarterly Journal of Economics* 128, no. 3: 967-1015.
- [39] Card, David, and Thomas Lemieux (2001). “Can falling supply explain the rising return to college for younger men? A cohort-based analysis.” *The quarterly journal of economics* 116, no. 2: 705-746.

- [40] Ciccone, Antonio, and Giovanni Peri (2005). “Long-run substitutability between more and less educated workers: evidence from US states, 1950–1990.” *Review of Economics and Statistics* 87, no. 4: 652-663.
- [41] Costinot, Arnaud, and Jonathan Vogel (2010). “Matching and inequality in the world economy.” *Journal of Political Economy* 118, no. 4: 747-786.
- [42] Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum (2014). “The rise of the East and the Far East: German labor markets and trade integration.” *Journal of the European Economic Association* 12, no. 6: 1643-1675.
- [43] Dengler, Katharina, and Britta Matthes (2018). “The impacts of digital transformation on the labour market: Substitution potentials of occupations in Germany.” *Technological Forecasting and Social Change* 137: 304-316.
- [44] Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth Troske (2004). “Wage and productivity dispersion in United States manufacturing: The role of computer investment.” *Journal of Labor Economics* 22, no. 2: 397-429.
- [45] Dustmann, Christian, Johannes Ludsteck, and Uta Schonberg (2009). “Revisiting the German wage structure.” *The Quarterly journal of economics* 124.2: 843-881.
- [46] Ebenstein, Avraham, Ann Harrison, Margaret McMillan, and Shannon Phillips (2014). “Estimating the impact of trade and offshoring on American workers using the current population surveys.” *Review of Economics and Statistics* 96, no. 4: 581-595.
- [47] Eeckhout, Jan, and Philipp Kircher (2010). “Sorting and decentralized price competition.” *Econometrica* 78, no. 2: 539-574.
- [48] Eeckhout, Jan, and Philipp Kircher (2011). “Identifying sorting—in theory.” *The Review of Economic Studies* 78, no. 3: 872-906.
- [49] Feenstra, Robert, and Gordon Hanson (1999). “The impact of outsourcing and high-technology capital on wages: estimates for the United States, 1979–1990.” *The quarterly journal of economics* 114, no. 3: 907-940.

- [50] Fernandez, Roberto (2001). "Skill-biased technological change and wage inequality: Evidence from a plant retooling." *American Journal of Sociology* 107, no. 2: 273-320.
- [51] Firpo, Sergio, Nicole Fortin, and Thomas Lemieux (2011). "Occupational tasks and changes in the wage structure." No. 5542. IZA Discussion Papers.
- [52] Fortin, Nicole, and Thomas Lemieux (1997). "Institutional changes and rising wage inequality: Is there a linkage?." *Journal of Economic Perspectives* 11, no. 2: 75-96.
- [53] Frey, Carl Benedikt, and Michael Osborne (2017). "The future of employment: How susceptible are jobs to computerisation?." *Technological forecasting and social change* 114: 254-280.
- [54] Garicano, Luis. "Hierarchies and the Organization of Knowledge in Production." *Journal of political economy* 108, no. 5 (2000): 874-904.
- [55] Garz, Marcel (2013). "Employment and wages in Germany since the 2004 deregulation of the temporary agency industry." *International Labour Review* 152, no. 2: 307-326.
- [56] Gathmann, Christina, and Uta Schonberg (2010). "How general is human capital? A task-based approach." *Journal of Labor Economics* 28.1: 1-49.
- [57] Gechert, Sebastian, Tomas Havranek, Zuzana Irsova, and Dominika Kolcunova (2021). "Measuring capital-labor substitution: The importance of method choices and publication bias." *Review of Economic Dynamics* (2021).
- [58] Gibbons, Robert, and Lawrence Katz (1992). "Does unmeasured ability explain inter-industry wage differentials?." *The Review of Economic Studies* 59, no. 3: 515-535.
- [59] Goldschmidt, Deborah, and Johannes Schmieider (2017). "The rise of domestic outsourcing and the evolution of the German wage structure." *The Quarterly Journal of Economics* 132, no. 3: 1165-1217.
- [60] Goos, Maarten, and Alan Manning (2007). "Lousy and lovely jobs: The rising polarization of work in Britain." *The review of economics and statistics* 89.1: 118-133.

- [61] Goos, Maarten, Alan Manning, and Anna Salomons (2014). “Explaining job polarization: Routine-biased technological change and offshoring.” *American Economic Review* 104, no. 8: 2509-26.
- [62] Haanwinckel, Daniel (2021). “Supply, Demand, Institutions, and Firms: A Theory of Labor Market Sorting and the Wage Distribution.” Unpublished manuscript.
- [63] Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii (2017). “Identifying equilibrium models of labor market sorting.” *Econometrica* 85, no. 1: 29-65.
- [64] Hakanson, Christina, Erik Lindqvist, and Jonas Vlachos (2021). “Firms and Skills The Evolution of Worker Sorting.” *Journal of Human Resources* 56, no. 2: 512-538.
- [65] Hall, Anja, and Beermann, Beate (2021). BIBB/BAuA-Employment Survey 2006. GESIS Data Archive, Cologne. ZA4820 Data file Version 1.0.1, <https://doi.org/10.4232/1.11072>.
- [66] Haltiwanger, John, and James Spletzer (2020). “Between Firm Changes in Earnings Inequality: The Dominant Role of Industry Effects.” No. w26786. National Bureau of Economic Research.
- [67] Haltiwanger, John, Henry Hyatt, and James Spletzer (2022). “Industries, mega firms, and increasing inequality.” No. 15197. Institute of Labor Economics (IZA).
- [68] Helpman, Elhanan, Oleg Itskhoki, and Stephen Redding (2010). “Inequality and unemployment in a global economy.” *Econometrica* 78, no. 4: 1239-1283.
- [69] Hemous, David, and Morten Olsen (2020). “The rise of the machines: Automation, horizontal innovation and income inequality.” *American Economic Journal: Macroeconomics* (Forthcoming).
- [70] Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang (2014). “The wage effects of offshoring: Evidence from Danish matched worker-firm data.” *American Economic Review* 104, no. 6: 1597-1629.

- [71] Jahn, Elke, and Dario Pozzoli (2013). “The pay gap of temporary agency workers—Does the temp sector experience pay off?” *Labour Economics* 24: 48-57.
- [72] Johnson, George (1997). “Changes in earnings inequality: the role of demand shifts.” *Journal of economic perspectives* 11, no. 2: 41-54.
- [73] Kohlbrecher, Britta, Christian Merkl, and Daniela Nordmeier (2016). “Revisiting the matching function.” *Journal of Economic Dynamics and Control* 69: 350-374.
- [74] Krause, Michael, and Harald Uhlig (2012). “Transitions in the German labor market: Structure and crisis.” *Journal of Monetary Economics* 59, no. 1: 64-79.
- [75] Kremer, Michael, and Eric Maskin (1996). “Wage Inequality and Segregation by Skill,” NBER Working Paper no. 5718.
- [76] Krueger, Alan, and Jorn-Steffen Pischke (1995). “12 A Comparative Analysis of East and West German Labor Markets: Before and After Unification.” Differences and changes in wage structures: 405.
- [77] Krueger, Alan, and Lawrence Summers (1988). “Efficiency wages and the inter-industry wage structure.” *Econometrica: Journal of the Econometric Society*: 259-293.
- [78] Krusell, Per, Lee Ohanian, Jose-Victor Ríos-Rull, and Giovanni Violante (2000). “Capital-skill complementarity and inequality: A macroeconomic analysis.” *Econometrica* 68, no. 5: 1029-1053.
- [79] Lee, Donghoon, and Kenneth Wolpin (2010). “Accounting for wage and employment changes in the US from 1968–2000: A dynamic model of labor market equilibrium.” *Journal of Econometrics* 156, no. 1: 68-85.
- [80] Lee, Gwanhoo, and Weidong Xia (2006). “Organizational Size and IT Innovation Adoption: A Meta-Analysis.” *Information and Management* 43.8: 975-985.
- [81] Lemieux, Thomas (2006). “Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill?” *American Economic Review* 96, no. 3: 461-498.

- [82] Lentz, Rasmus (2010). “Sorting by search intensity.” *Journal of Economic Theory* 145, no. 4: 1436-1452.
- [83] Levy, Frank, and Richard Murnane (1996). “With what skills are computers a complement?.” *The American Economic Review* 86, no. 2 (1996): 258-262.
- [84] Lise, Jeremy, Costas Meghir, and Jean-Marc Robin (2016). “Matching, sorting and wages.” *Review of Economic Dynamics* 19: 63-87.
- [85] Lopes de Melo, Rafael (2018). “Firm wage differentials and labor market sorting: Reconciling theory and evidence.” *Journal of Political Economy* 126, no. 1: 313-346.
- [86] Michaels, Guy, Ashwini Natraj, and John Van Reenen (2014). “Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years.” *Review of Economics and Statistics* 96, no. 1: 60-77.
- [87] OECD (2021). “The Role of Firms in Wage Inequality: Policy Lessons from a Large Scale Cross-Country Study.” OECD Publishing, Paris, <https://doi.org/10.1787/7d9b2208-en>.
- [88] Peretto, Pietro, and John Seater (2013). “Factor-eliminating technical change.” *Journal of Monetary Economics* 60, no. 4: 459-473.
- [89] Petrongolo, Barbara, and Christopher Pissarides (2001). “Looking into the black box: A survey of the matching function.” *Journal of Economic literature* 39, no. 2: 390-431.
- [90] Rohrbach-Schmidt, Daniela, and Michael Tiemann (2013). “Changes in workplace tasks in Germany—evaluating skill and task measures.” *Journal for Labour Market Research* 46.3: 215-237.
- [91] Schimpl-Neimanns, Bernhard. “Mikrodaten-tools: Umsetzung der Berufsklassifikation von Blossfeld auf die Mikrozensus 1973-1998.” (2003): 39.
- [92] Schmidlein, Lisa, Stefan Seth, and Matthias Umkehrer (2019): *Linked-Employer-Employee-Daten des IAB: Linked Employer-Employee Data from the IAB: LIAB Cross-Sectional Model 2 (LIAB QM2) 1993–2017*. FDZ-Datenreport, 06/2019 (en), Nuremberg. DOI: 10.5164/IAB.FDZD.1906.en.v1

- [93] Shimer, Robert, and Lones Smith (2000). “Assortative matching and search.” *Econometrica* 68, no. 2: 343-369.
- [94] Simon, Hipolito (2010). “International differences in wage inequality: a new glance with European matched employer–employee data.” *British Journal of Industrial Relations* 48, no. 2: 310-346.
- [95] Snower, Dennis, and Christian Merkl (2006). “The caring hand that cripples: The East German labor market after reunification.” *American Economic Review* 96, no. 2: 375-382.
- [96] Sokolova, Anna, and Todd Sorensen (2021). “Monopsony in labor markets: A meta-analysis.” *ILR Review* 74, no. 1: 27-55.
- [97] Song, Jae, David Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter (2019). “Firming up inequality.” *The Quarterly journal of economics* 134, no. 1: 1-50.
- [98] Spermann, Alexander (2011). “The New Role of Temporary Agency Work in Germany.” *IZA Discussion Papers* 6180.
- [99] Spitz-Oener, Alexandra (2006). “Technical change, job tasks, and rising educational demands: Looking outside the wage structure.” *Journal of labor economics* 24.2: 235-270.
- [100] Teulings, Coen (1995). “The wage distribution in a model of the assignment of skills to jobs.” *Journal of political Economy* 103, no. 2: 280-315.
- [101] Tomaskovic-Devey, Donald, Anthony Rainey, Dustin Avent-Holt, Nina Bandelj, István Boza, David Cort, Olivier Godechot et al. (2020). “Rising between-workplace inequalities in high-income countries.” *Proceedings of the National Academy of Sciences* 117, no. 17: 9277-9283.
- [102] Torres, Sonia, Pedro Portugal, John Addison, and Paulo Guimaraes (2018). “The sources of wage variation and the direction of assortative matching: Evidence from a three-way high-dimensional fixed effects regression model.” *Labour Economics* 54: 47-60.
- [103] Yamaguchi, Shintaro (2012). “Tasks and heterogeneous human capital.” *Journal of Labor Economics* 30, no. 1: 1-53.

[104] Zeira, Joseph (1998). “Workers, Machines, and Economic Growth,” *Quarterly Journal of Economics*, 113(4): 1091–11

LIAB Linked-Employer-Employee-Data of the IAB. This study uses the LIAB cross-sectional model 2, version 1993-2017, of the Linked-Employer-Employee Data (LIAB) from the IAB. Data access was provided via on-site use at the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) and subsequently through remote data access. AKM wage estimates also provided by IAB, following the methodology of Card et al. (2013) and as detailed in Bellman et al. (2020). DOI: 10.5164/IAB.LIABQM29317.de.en.v1

BIBB/BAuA-Employment Survey 2006. Appendix results make use of task data from the 2006 BIBB/BAuA employment survey, made available by GESIS as a scientific use file. These data were aggregated by 3-digit KLDB 1988 occupation and subsequently merged with LIAB. DOI: 10.4232/1.11072

Labor Market Statistics, Federal Employment Agency [*13211*]. Retrieved from Destatis, Federal Statistical Office, on January 4, 2021.

Consumer Price Index of All Items in Germany, Organization for Economic Cooperation and Development [*DEUCPIALLMINMEI*]. Retrieved from FRED, Federal Reserve Bank of St. Louis, on March 18, 2022.

Appendix A

Supplementary Materials For Chapter 1

A.1 Data Preparation

LIAB Matched Dataset. The LIAB cross-sectional database spans the years 1993-2017 in four waves: 1993-1999, 2000-2006, 2007-2013, and 2014+. Establishments in the IAB establishment survey are linked with workers through the administrative social security data. The number of sampled establishments varies between approximately 4,000 and 15,000 per year, increasing during the 1990's and decreasing thereafter; while the number of employed individuals is between 1.6 and 2.5 million per year, representing a (non-random) 5% sample of the German employed labor force. The IAB establishment survey is a representative sample of establishments with at least one employee subject to social security taxes as of June, with the survey taking place in the third quarter of the following year. The sample is stratified by size, industry, and state, making population weights necessary for inferences about the aggregate economy; smaller establishments, industries, and states are over-sampled. Establishments are aggregated when operating in the same industry and location and under the same ownership, and at the discretion of the interviewer. The same establishment may be interviewed in different waves and assigned a different identification number.

Data cleaning is as follows. For better comparability across individuals, and following past work, I drop part-time workers and apprentices from the sample. These observations, representing 3-5% of the total, are problematic because they introduce variation in earnings due to hours worked and job status, which are not the focus of this paper. In addition, part-time work is measured sporadically after 2011. Workers performing temporary “mini-jobs” are missing from the sample altogether, as these jobs are exempt from social security taxes. Time-consistent industry and occupation codes are provided for different coding systems. I

Table A.1: LIAB Summary Statistics

	1993-1999	1998-2004	2003-2010	2010-2017
Observations	10681613	9228246	9558800	7109948
Persons	3383057	3336946	3129848	2366669
Establishments	11935	24546	26981	23031
Identified sample				
Observations	10,645,769	9,185,412	9,511,130	7,080,688
Persons	3,351,593	3,301,936	3,097,049	2,347,598
Establishments	8,151	18,518	19,989	17,684
% weighted sample	0.978	0.972	0.97	0.977

NOTES. Sample consists of full-time employees (male and female) aged 20-60 in surveyed West German establishments. Person/establishment fixed effects estimated by IAB on the underlying administrative dataset following Card, Heining, Kline (2013). Daily wage in log 1995 euros, censored values imputed.

use 1988 KLDB occupation codes as these are the original codes for the largest portion of the sample period (1993-2010); and 2008 WZ industry codes as these are the most detailed and hence require the least imputation when cross-walking. To address inconsistencies and missing values, I use forward and backward imputation within job spells (occupation) and sets of observations within the same establishment (industry). Non-specific job categories (“disabled”, “rehabilitant”, and codes 971 and above) are set to missing. Forward and backward imputation are also used to populate missing education codes, where educational categories consists of a lower secondary education, a completed apprenticeship, an upper secondary education, a university degree, and a ‘missing’ category that represents 12% of the sample in the early 1990’s and 6% by the late 2010’s. Subsequent to these steps, observations are condensed by person-year by either combining job spells when they occur at the same employer and in the same occupation, or selecting the job spell associated with the greatest amount of income.

Earnings are converted to 1995 Euro values, with daily values below 10 dropped from the sample. Earnings are top-coded at social security thresholds, which can affect upwards of 10% of the sample in any given year. I therefore impute wages using a Tobit approach identical to that used by CHK, with the difference that regressions are not performed separately by age due to sample size limitations. Regressions are performed separately by year, gender, and educational attainment, with dependent variables including age (nominal and

grouped), a dummy variable for establishments with 11 or more workers, another dummy for only a single worker, employment (nominal and squared), the leave-one-out mean wage and censored percentage conditional on the establishment, the leave-one-out mean wage and censored percentage conditional on the person, average years of schooling within the establishment (years imputed following CHK), the percent of workers within the establishment with a college education, and a dummy variable for persons only observed once. Imputed values greater than 1000 are then capped at 1000.

Industry and Occupation Aggregations. Industry and occupational classifications are based off of the WZ 2008 and KLDB 1988 classifications, respectively. For years prior to 2008, the data provider imputes industry using extrapolation when possible and, when not, correspondence tables. This is primarily an issue for observations prior to 1999, which were recorded with a much older classification system (WZ 1973). Occupation codes after 2010 are recorded as KLDB 2010 values, for which there exists no direct correspondence with the earlier system. The data provider imputes KLDB 1988 occupation for these years but with substantial inaccuracies; comparing matches observed both before and after 2010, roughly 1/3 experience a change in occupational code, generally involving a ‘move’ to a similar occupation or from a detailed occupation to a broad N.E.C. category.

To minimize errors from recoding, I conduct the empirical analysis in this paper using aggregated industry and occupational groups. Doing so introduces a second problem: the KLDB 1988 system lacks hierarchical categories, and aggregate WZ 2008 classifications (which follow the NACE system) do an unsatisfactory job of preserving industry wage differentials. In both cases I find it preferable to aggregate classifications manually by combining neighboring industries (occupations) that exhibit similar mean establishment (person) wage effects. This is necessarily an arbitrary approach, but it is successful in producing industry and occupation classifications that are comparable over time while preserving as much as possible of the underlying wage structure.

Table A.2: Aggregate Industry Classifications

Industry	Mean person wage effect	Mean estab wage effect	WZ 2008 industry codes
Food mfg.	4.38	-.15	11-32, 161, 101-103, 107
Crafts mfg.	4.53	.01	310-332, 370-439
Materials mfg.	4.54	.04	104-106, 108-152, 162-182, 221-239, 251, 255-259, 292, 331-332, 370-390, 411-429, 51-99, 241-245, 252-254, 264-275, 281-289
Durables mfg.	4.67	.12	191-212, 261-263, 279, 290-291, 293-309
High-tech mfg.	4.71	.18	451-454, 471-532
Wholesale	4.62	.01	561, 563, 781-783
Retail/transport	4.52	-.05	472-473, 476-478, 551-559, 562, 801-822, 829, 920, 931-932, 960
Hospitality/temp	4.34	-.34	561, 563, 781-783
Personal svc.	4.40	-.18	472-473, 476-478, 551-559, 562, 801-822, 829, 920, 931-932, 960
Professional svc.	4.53	-.13	691-692, 741-774, 791-822, 829, 855-856, 862-889, 951-952
Commercial svc.	4.68	.00	661-683, 711, 731, 823, 841-854, 860-861, 900, 910, 941-949
Information svc.	4.82	.09	581-653, 701-702, 712-722, 732

NOTES. Mean wage effects calculated by panel, then averaged across panels.

Table A.3: Aggregate Occupation Classifications

Occupation	Mean person	Mean estab	KLDB 1988 occupation codes
	wage effect	wage effect	
Goods handlers	4.35	-.02	111, 121, 143-152, 181-184, 212, 242, 321-346, 321, 331-344, 361-376, 402-403, 412, 432-433, 492, 531
Machinists	4.47	.07	82-101, 112, 131-135, 161-162, 164, 231-241, 322-323, 422-432, 441-442, 461-466, 71-72, 141-142, 191-211, 213-226, 263, 273-274, 284-291, 312-314, 521, 541-549, 711
Mechanics	4.57	.10	11-21, 51, 163, 175-177, 391-401, 451-453, 470-491, 501-513, 804, 834
Artisans	4.47	-.05	102, 171-174, 251-262, 270-271, 275-282, 301-311, 315, 624, 631, 634-635, 713-716, 725
Fine crafts	4.59	-.01	41-44, 53, 351-358, 411, 792, 856, 901-902, 912-921, 923-937
Unskilled labor	4.28	-.21	
Movers	4.43	-.03	522, 712, 714, 723-724, 731-732, 741-744, 791, 793-801, 805, 814, 851-852, 854-855, 861, 864, 911
Semi-skilled service	4.52	-.07	681, 683, 701-705, 753, 772, 782-783, 811, 822-833, 835-838, 842-844, 853, 857, 862-863, 875-877, 891-893,
Skilled service	4.67	-.02	682, 684-686, 688, 706, 733-734, 773, 784, 781
Sales clerks	4.39	-.10	682, 684-686, 688, 706, 733-734, 773, 784, 781
Office specialists	4.61	.00	
Technicians	4.84	.09	32, 52, 61, 283, 621-623, 625-629, 632-633 691-694, 721-722, 726, 771, 774, 802-803, 922
Engineers	5.07	.11	601-612
Managers	5.03	.05	687, 751-752, 761-763
Doctors	5.04	.02	813, 821, 841, 871-874, 881-883,

NOTES. Mean wage effects calculated by panel, then averaged across panels.

Table A.4: AKM Identification Statistics

	1993-99	1998-04	2003-10	2010-17
<i>Job Entry (%)</i>				
Full-time, reweighted	0.171	0.168	0.150	0.165
Full-time, males-only	0.150	0.143	0.121	0.134
Full sample	0.171	0.168	0.150	0.165
<i>Job Exit (%)</i>				
Full-time, reweighted	0.184	0.179	0.162	0.168
Full-time, males-only	0.166	0.159	0.135	0.141
Full sample	0.184	0.179	0.162	0.168

SOURCE: German linked employer-employee dataset (LIAB). NOTE: All proportions use sample weights. Job entry (exit) is new (ending) job spells as percent of total employment at continuing establishments. Re-weighted rates hold fixed the employment shares of industry-size pairs, for 12 aggregate industry groups and 5 size (employment) groups. Full sample includes part-time workers and apprentices.

A.2 Variance Decompositions




Figure A.1: West German Wage Variance, 1993-2017

SOURCE: German linked employer-employee dataset (LIAB). NOTE: Variance of log daily wage, full-time West German workers.





Figure A.2: West German Wage Sorting (Correlation), 1993-2017

SOURCE: German linked employer-employee dataset (LIAB) and IAB-provided wage effects. NOTE: Solid (dashed) line indicates panel (annual) value of $\text{Corr}(\pi, \phi)$.

Table A.5: Wage Sorting By Demographic (Correlation)

	1993-99	1998-04	2003-10	2010-17
Male	0.067	0.145	0.155	0.249
Female	0.104	0.096	0.041	0.129
Aged 20-30	0.051	0.095	0.061	0.167
Aged 31-40	0.129	0.164	0.142	0.244
Aged 41-50	0.168	0.175	0.165	0.236
Aged 51-60	0.153	0.187	0.148	0.211
Lower secondary ed.	0.219	0.194	0.146	0.266
Apprenticeship	0.095	0.120	0.097	0.203
Upper secondary ed.	0.108	0.147	0.095	0.149
University degree	0.180	0.155	0.153	0.099
Schleswig-Holstein	0.096	0.170	0.107	0.152
Hamburg	0.181	0.192	0.226	0.012
Lower Saxony	0.118	0.105	0.080	0.300
Bremen	0.231	0.209	0.198	0.302
North Rhine-Westphalia	0.111	0.167	0.120	0.210
Hesse	0.151	0.198	0.132	0.207
Rhineland-Palatinate	0.085	0.136	0.105	0.270
Wurttemberg-Baden	0.131	0.187	0.149	0.241
Bavaria	0.141	0.138	0.175	0.224
Saarland	0.112	0.044	0.140	0.209
Berlin	0.221	0.223	0.206	0.248

SOURCE: German linked employer-employee dataset (LIAB) and IAB-provided wage effects. NOTE: Value shown is $\text{Cor}(\pi, \phi)$ where π (ϕ) is the person (establishment) AKM wage effect.

**Figure A.3:** East German Wage Variance, 1993-2017

SOURCE: German linked employer-employee dataset (LIAB). NOTE: Variance of log daily wage, full-time East German workers.

Table A.6: AKM Variance Decomposition, East Germany 1996-2017

	1993-99	1998-04	2003-10	2010-17
$\text{Var}(w)$	0.1664	0.1813	0.2076	0.1925
$\text{Var}(\pi)$	0.0780	0.0966	0.1031	0.1122
$\text{Var}(\phi)$	0.0381	0.0454	0.0582	0.0438
$\text{Var}(x'\beta)$	0.0034	0.0052	0.0050	0.0105
$\text{Var}(\epsilon)$	0.0101	0.0106	0.0121	0.0120
$2 \times \text{Cov}(\pi, \phi)$	0.0294	0.0266	0.0254	0.0362
$2 \times \text{Cov}(\pi, x'\beta)$	0.0006	-0.0040	-0.0018	-0.0218
$2 \times \text{Cov}(\phi, x'\beta)$	0.0016	0.0016	0.0026	-0.0016
Observations	1,914,798	2,602,427	2,123,684	1,560,271
Persons	885,622	888,080	669,038	500,815
Establishments	6,555	8,175	8,039	7,183

Table A.7: AKM Variance Decomposition, Germany 1996-2017

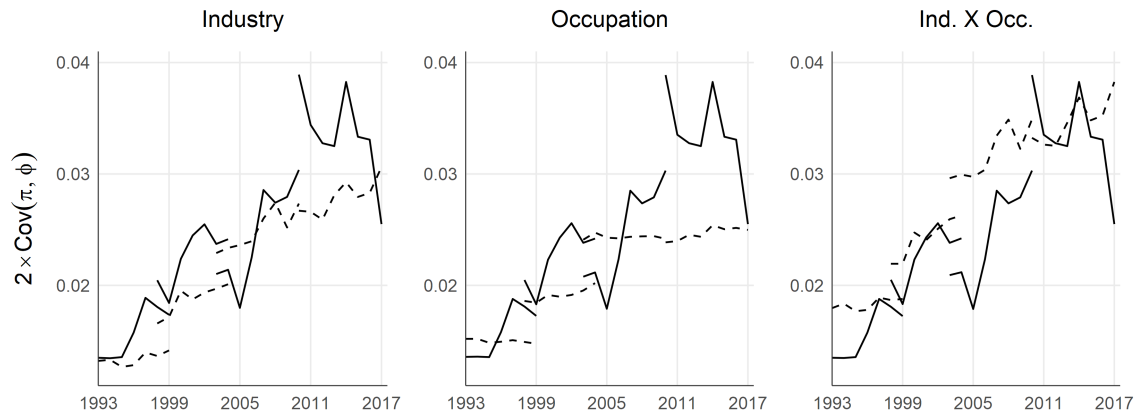
	1996-99	1998-04	2003-10	2010-17
$\text{Var}(w)$	0.1865	0.2125	0.2436	0.2383
$\text{Var}(\pi)$	0.1015	0.1204	0.1357	0.1393
$\text{Var}(\phi)$	0.0416	0.0482	0.0608	0.0448
$\text{Var}(x'\beta)$	0.0035	0.0054	0.0054	0.0127
$\text{Var}(\epsilon)$	0.0119	0.0142	0.0151	0.0173
$2 \times \text{Cov}(\pi, \phi)$	0.0270	0.0296	0.0316	0.0400
$2 \times \text{Cov}(\pi, x'\beta)$	0.0020	-0.0008	-0.0006	-0.0190
$2 \times \text{Cov}(\phi, x'\beta)$	0.0000	0.0014	0.0020	0.0004
Observations	6,939,650	11,787,839	11,634,814	8,640,959
Persons	2,961,442	4,182,882	3,759,847	2,844,816
Establishments	13,061	26,687	28,012	24,863

A.3 Additional Between-Group Results

Table A.8: Between-Group Wage Moments (% Total)

	1993-99	1998-04	2003-10	2010-17
$Cov(\pi, \phi)$				
Industry (46)	0.884	0.853	1.018	0.843
Occupation (75)	0.885	0.812	0.958	0.746
Ind. \times Occ	1.187	1.086	1.300	1.077
Education (5)	0.244	0.275	0.353	0.356
Occ. \times Educ.	0.962	0.876	1.022	0.798
Number of employees (5)	0.737	0.639	0.716	0.463
Ind. \times Size	1.068	1.026	1.246	0.989
$Var(\pi)$				
Industry (12)	0.081	0.102	0.112	0.166
Industry (46)	0.095	0.115	0.124	0.177
Occupation (15)	0.300	0.307	0.303	0.357
Occupation (75)	0.323	0.326	0.319	0.376
Ind. \times Occ (12 \times 15)	0.326	0.334	0.329	0.394
Ind. \times Occ (46 \times 75)	0.380	0.379	0.371	0.438
$Var(\phi)$				
Industry (12)	0.296	0.319	0.323	0.311
Industry (46)	0.339	0.350	0.345	0.331
Occupation (15)	0.160	0.159	0.163	0.152
Occupation (75)	0.213	0.211	0.215	0.188
Ind. \times Occ (12 \times 15)	0.337	0.352	0.356	0.342
Ind. \times Occ (46 \times 75)	0.427	0.433	0.424	0.412

SOURCE: German linked employer-employee dataset (LIAB) and IAB-provided wage effects. NOTE: Value shown is the between-group covariance divided by the total covariance.

**Figure A.4:** Annual Between-Group Wage Sorting, 1993-2017

SOURCE: German linked employer-employee dataset (LIAB). NOTE: Solid (dashed) line indicates total (between-group) value of $2 \times Cov(\pi, \phi)$.




Figure A.5: Wage Effects by Detailed Occupation and Industry

SOURCE: German linked employer-employee dataset (LIAB) and IAB-provided wage effects. NOTE: Weighted average of AKM wage effects by 75 KLDB 1988 occupation groups and 46 WZ 2008 industry groups.




Figure A.6: Wage Sorting and Match Selection (Correlations), 2010-2017

SOURCE: German linked employer-employee dataset (LIAB) and IAB-provided wage effects. NOTE: Time in workforce dated from first payment into social security.

Table A.9: Counterfactual Between-Group Wage Sorting (Correlation)

	1993-99	1998-04	2003-10	2010-17
Total between-group	0.473	0.515	0.551	0.633
Constant group π	0.506	0.534	0.553	0.587
By industry only	0.517	0.535	0.549	0.591
By occupation only	0.462	0.512	0.549	0.632
Constant group ϕ	0.481	0.523	0.557	0.617
By industry only	0.474	0.525	0.573	0.628
By occupation only	0.514	0.513	0.506	0.579
Constant group ω	0.494	0.524	0.531	0.605
By industry only	0.508	0.528	0.537	0.599
By occupation only	0.471	0.511	0.541	0.638
Constant group (π, ω)	0.535	0.545	0.541	0.563
By (occupation, industry)	0.493	0.520	0.538	0.599
By (occupation, both)	0.491	0.524	0.536	0.612
By (both, industry)	0.54	0.543	0.543	0.554

NOTE: Value shown is the counterfactual value of $Cor(\pi, \phi)$. Groups are 180 industry \times occupation pairs. Person effects (π), establishment effects (ϕ), and employment shares (ω) are held constant at their 1993-99 and 2010-17 values, as indicated, with the average correlation reported.

Table A.10: Counterfactual Between-Group Wage Sorting (% Trend)

	1993-99 Values			2010-17 Values		
	1998-04	2003-10	2010-17	1993-99	1998-04	2003-10
Total between-group	0.361	0.831	1.000	0.000	0.361	0.831
Constant group π	0.253	0.590	0.494	0.313	0.639	1.108
By industry only	0.253	0.590	0.518	0.373	0.675	1.096
By occupation only	0.313	0.711	0.771	0.096	0.446	0.916
Constant group ϕ	0.241	0.494	0.964	0.060	0.301	0.554
By industry only	0.277	0.639	1.108	0.012	0.277	0.578
By occupation only	0.169	0.349	0.554	0.265	0.482	0.711
Constant group ω	0.193	0.386	0.434	0.482	0.843	1.072
By industry only	0.217	0.458	0.470	0.518	0.795	1.048
By occupation only	0.337	0.723	0.904	0.060	0.422	0.843
Constant group (π, ω)	0.108	0.229	0.060	0.964	1.181	1.373
By (occupation, industry)	0.157	0.361	0.301	0.627	0.892	1.157
By (occupation, both)	0.145	0.313	0.289	0.663	0.976	1.181
By (both, industry)	0.108	0.265	0.084	0.940	1.133	1.361

NOTE: Value shown is the counterfactual between-group component of $Cov(\pi, \phi)$ divided by the total between-group trend for 1993-2017. Groups are 180 industry \times occupation pairs. Person effects (π), establishment effects (ϕ), and employment shares (ω) are held constant at their 1993-99 and 2010-17 values, as indicated. For results using constant 1993-99 values, a column of zeros is omitted for 1993-99, and likewise a column of ones for 2010-17 with constant 2010-17 values.

Table A.11: Between-Group Variances

	1993-99	1998-04	2003-10	2010-17
<i>Var</i> (π)				
Industry (12)	0.0088	0.0127	0.0156	0.0235
Occupation (15)	0.0326	0.0379	0.0423	0.0504
Ind. \times Occ	0.0355	0.0412	0.0459	0.0556
<i>Var</i> (ϕ)				
Industry (12)	0.0092	0.0121	0.0167	0.0124
Occupation (15)	0.0050	0.0060	0.0084	0.0061
Ind. \times Occ	0.0104	0.0134	0.0184	0.0136

SOURCE: German linked employer-employee dataset (LIAB) and IAB-provided wage effects. NOTE: Value shown is the between-group variance.




Figure A.7: Establishment Effects by Industry and Panel

NOTE: Value shown is the weighted average of the establishment wage effect ϕ .

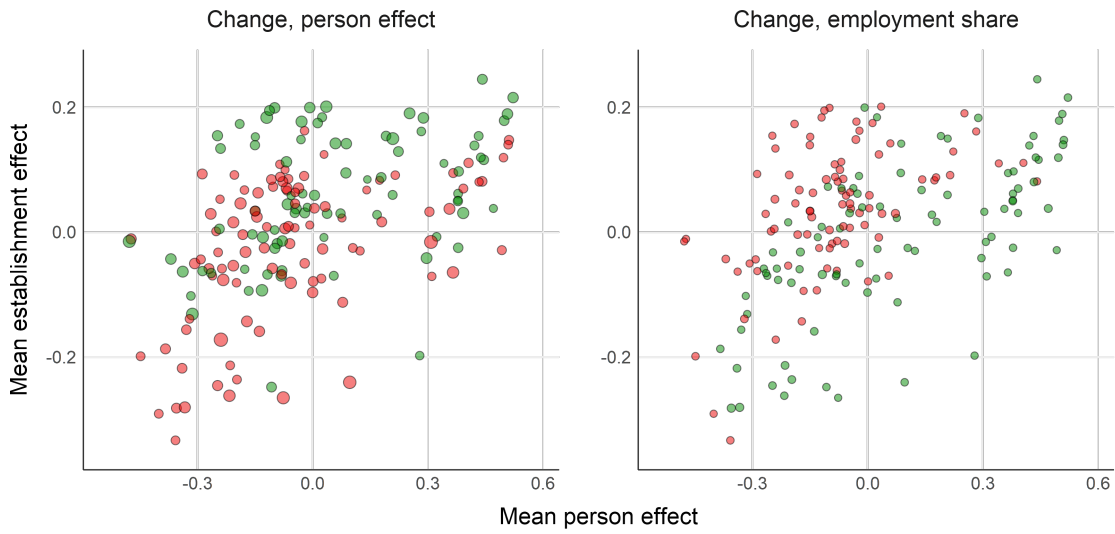


Figure A.8: Changes to Industry-Occupation Mean Person Effects and Shares

NOTE: Axes indicate group mean wage effects, averaged across 1993-1999 and 2010-2017. Green (red) bubbles indicate an increase (decrease) in group mean person effect and employment shares, measured as the difference between 1993-1999 and 2010-2017 values.

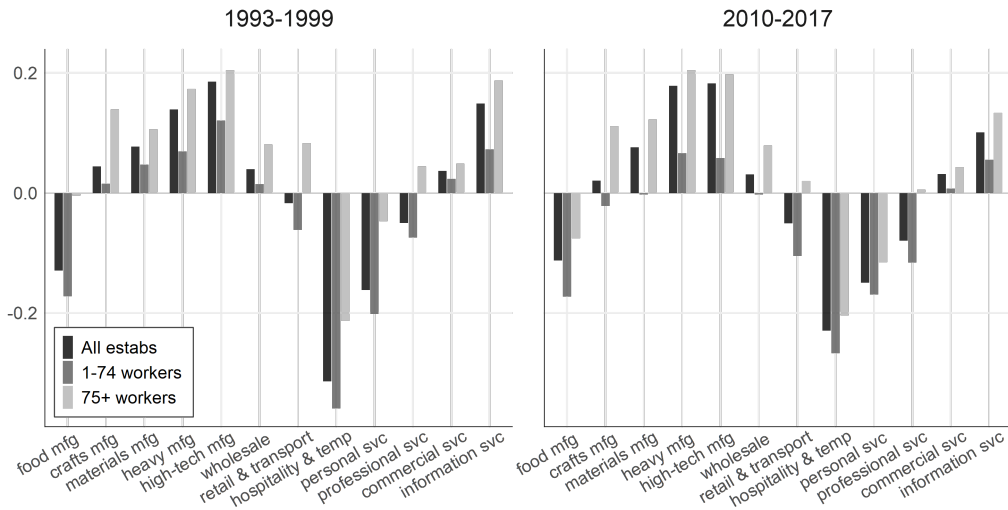


Figure A.9: Establishment Effects by Industry and Size

NOTE: Value shown is the weighted average of the establishment wage effect ϕ . Size categories based on full-time employees.




Figure A.10: Establishment Effects by Industry and Bargaining Agreement

NOTE: Value shown is the weighted average of the establishment wage effect ϕ . Size categories based on full-time employees.

A.4 Additional Results, Person/Establishment Characteristics

Table A.12: Cross-Correlations, Establishment Characteristics

	VAD	Inv	ICT	Multi	Barg	Prod.	Proc.	Profit	Compet.
Log employment	0.979	0.766	0.775	0.745	0.565	0.858	0.926	-0.213	0.233
Log value-Added		0.732	0.837	0.823	0.881	0.939	0.486	-0.169	0.255
Multi-estab. firm (%)			0.391	0.698	0.612	0.662	0.339	-0.031	-0.002
Log investment/worker				0.785	0.834	0.843	0.275	-0.085	0.226
Log ICT inv./worker					0.843	0.838	0.15	0.236	0.049
Product dev. (%)						0.975	0.297	0.052	0.361
Process impr. (%)							0.403	-0.036	0.331
Collective barg. (%)								-0.397	0.017
Positive profits (%)									0.018

Table A.13: Wage Effects and Establishment Characteristics, Confidence Intervals

Variable (X)	Correlations		Partial correlations		
	$\text{Cor}(X, \pi)$	$\text{Cor}(X, \phi)$	$\text{Cor}(\pi_X, \phi)$	$\text{Cor}(\pi, \phi_X)$	$\text{Cor}(\pi_X, \phi_X)$
Log employment	(0.27, 0.52)	(0.68, 0.81)	(0.13, 0.4)	(0.24, 0.49)	(0.27, 0.52)
Log value-added	(0.34, 0.57)	(0.8, 0.88)	(0.03, 0.32)	(0.16, 0.42)	(0.2, 0.46)
Multi-estab. firm (%)	(0.37, 0.59)	(0.42, 0.63)	(0.19, 0.45)	(0.2, 0.46)	(0.26, 0.51)
Log investment/worker	(0.27, 0.52)	(0.82, 0.89)	(0.08, 0.36)	(0.27, 0.52)	(0.31, 0.55)
Log ICT inv./worker	(0.55, 0.72)	(0.75, 0.85)	(-0.12, 0.18)	(-0.11, 0.19)	(-0.09, 0.2)
Product dev. (%)	(0.3, 0.54)	(0.72, 0.83)	(0.09, 0.37)	(0.21, 0.47)	(0.25, 0.5)
Process impr. (%)	(0.28, 0.53)	(0.75, 0.85)	(0.09, 0.37)	(0.23, 0.48)	(0.27, 0.51)
Collective barg. (%)	(-0.08, 0.21)	(0.2, 0.46)	(0.41, 0.62)	(0.45, 0.65)	(0.45, 0.65)
Positive profits (%)	(0.12, 0.39)	(-0.13, 0.17)	(0.45, 0.65)	(0.43, 0.64)	(0.45, 0.65)
Competitive mkt. (%)	(-0.25, 0.04)	(0.12, 0.4)	(0.47, 0.67)	(0.5, 0.68)	(0.5, 0.69)

NOTE: All variables first aggregated at the industry-occupation level *via* a first-stage regression on fixed effects for industry-occupation and year. Terms π_x and ϕ_x indicate residuals from a regression of wage effects on establishment characteristics. All results weighted by employment. Confidence intervals are calculated using Fisher transformation, and are approximate as they do not account for survey error, error due to imputation, or error from the AKM wage effects regression.

Table A.14: BIBB Task Descriptions

Task	Description
Gather information	Gather information, research, document
Use PCs	Work with computers
Manage others	Organize, plan, prepare work processes for others
Buy, sell	Buy, procure, sell
Control machines	Monitor, control machines, systems, technical processes
Weigh, measure	Measure, check, perform quality control
Nurture, care	Nurture, take care of, heal
Clean	Clean, dispose of waste, recycle
Produce goods	Produce goods and wares
Repair, install	Repair, install
Transport goods	Transport, store, ship
PR	Advertising, marketing, public relations
Research, design	Develop, research, create
Teach	Educate, teach
Advise	Advise and inform
Host, serve	Host, accommodate, prepare meals
Guard, protect	Secure, protect, guard, monitor, regulate traffic

Table A.15: Cross-Correlations, Person/Job Characteristics

	Educ.	Ten.	Exp.	Res.	PCs	Org.	Sales	Mach.	Trns.	Serve	Clean
College degree (%)	0.976	-0.179	-0.123	0.67	0.621	0.599	0.054	-0.370	-0.607	-0.053	-0.623
Years of education		-0.163	-0.095	0.636	0.747	0.629	0.115	-0.476	-0.658	-0.071	-0.738
Job tenure			0.702	0.193	0.101	-0.192	-0.486	0.414	-0.084	-0.531	-0.021
Research, design					0.380	0.619	-0.169	0.228	-0.570	-0.140	-0.326
Use PCs						0.467	0.127	-0.541	-0.682	-0.139	-0.896
Manage others							0.382	-0.242	-0.542	0.296	-0.319
Buy, sell								-0.514	-0.064	0.509	0.011
Control machines									0.280	-0.119	0.581
Transport goods										0.018	0.561
Host, serve											0.355

Table A.16: Wage Effects and Person Characteristics, Confidence Intervals

Variable (X)	Correlations		Partial correlations		
	$\text{Cor}(X, \pi)$	$\text{Cor}(X, \phi)$	$\text{Cor}(\pi_X, \phi)$	$\text{Cor}(\pi, \phi_X)$	$\text{Cor}(\pi_X, \phi_X)$
<i>LIAB Person Characteristics</i>					
College degree (%)	(0.77, 0.87)	(0.11, 0.38)	(0.50, 0.68)	(0.22, 0.47)	(0.52, 0.70)
Years of education	(0.83, 0.90)	(0.13, 0.40)	(0.54, 0.72)	(0.19, 0.45)	(0.57, 0.74)
Job tenure	(0.02, 0.30)	(0.71, 0.83)	(0.30, 0.54)	(0.58, 0.74)	(0.59, 0.75)
Years in labor force	(0.05, 0.33)	(0.53, 0.71)	(0.31, 0.55)	(0.44, 0.64)	(0.45, 0.65)
<i>BIBB Task Characteristics</i>					
Research, design	(0.64, 0.79)	(0.33, 0.56)	(0.18, 0.44)	(0.1, 0.38)	(0.22, 0.48)
Use PCs	(0.72, 0.83)	(0.25, 0.5)	(0.27, 0.52)	(0.13, 0.4)	(0.31, 0.55)
Manage others	(0.51, 0.7)	(-0.17, 0.12)	(0.63, 0.78)	(0.45, 0.66)	(0.63, 0.78)
Buy, sell	(-0.2, 0.1)	(-0.61, -0.39)	(0.41, 0.62)	(0.5, 0.69)	(0.5, 0.69)
Control machines	(-0.37, -0.1)	(0.05, 0.33)	(0.51, 0.7)	(0.5, 0.69)	(0.52, 0.71)
Transport goods	(-0.7, -0.52)	(-0.4, -0.12)	(0.37, 0.59)	(0.27, 0.51)	(0.39, 0.61)
Host, serve	(-0.4, -0.13)	(-0.71, -0.54)	(0.26, 0.51)	(0.36, 0.59)	(0.38, 0.6)
Clean	(-0.82, -0.69)	(-0.52, -0.28)	(0.23, 0.49)	(0.12, 0.39)	(0.27, 0.52)

NOTE: All variables first aggregated at the industry-occupation level *via* a first-stage regression. Years of education is imputed following Card et al. (2013). Task values imputed from verbal frequencies. Terms π_x and ϕ_x indicate residuals from a regression of wage effects on job characteristics. All results weighted by employment; see appendix for confidence intervals. Confidence are intervals calculated using Fisher transformation, and are approximate as they do not account for survey error, error due to imputation, or error from the AKM wage effects regression.

Table A.17: Wage Effects and Job Characteristics, Additional Tasks

Variable (X)	Correlations		Partial correlations		
	$\text{Cor}(X, \pi)$	$\text{Cor}(X, \phi)$	$\text{Cor}(\pi_X, \phi)$	$\text{Cor}(\pi, \phi_X)$	$\text{Cor}(\pi_X, \phi_X)$
No control			0.548	0.548	0.548
Produce goods	-0.316	0.138	0.623	0.597	0.629
Repair, install	-0.253	0.076	0.586	0.569	0.588
Weigh, measure	-0.080	0.172	0.563	0.570	0.572
Gather information	0.842	0.350	0.468	0.270	0.500
Teach others	0.598	0.158	0.565	0.459	0.573
Advise others	0.599	-0.014	0.694	0.556	0.694
PR	0.481	-0.245	0.759	0.687	0.783
Nurture, care	-0.020	-0.136	0.545	0.550	0.550
Guard, protect	-0.137	0.075	0.563	0.559	0.565

NOTE: All variables first aggregated at the industry-occupation level *via* a first-stage regression. Years of education is imputed following Card et al. (2013). Task values imputed from verbal frequencies. Terms π_x and ϕ_x indicate residuals from a regression of wage effects on job characteristics. All results weighted by employment; see appendix for confidence intervals.

Table A.18: Wage Effects and Job Characteristics, Non-Imputed

Variable (X)	Correlations		Partial correlations		
	$\text{Cor}(X, \pi)$	$\text{Cor}(X, \phi)$	$\text{Cor}(\pi_X, \phi)$	$\text{Cor}(\pi, \phi_X)$	$\text{Cor}(\pi_X, \phi_X)$
No control			0.541	0.541	0.541
Research, design	0.724	0.422	0.341	0.259	0.376
Use PCs	0.703	0.402	0.364	0.282	0.397
Manage others	0.577	-0.008	0.668	0.545	0.668
Buy, sell	0.040	-0.401	0.557	0.608	0.608
Control machines	-0.245	0.226	0.615	0.612	0.631
Transport goods	-0.580	-0.245	0.489	0.411	0.505
Host, serve	-0.216	-0.511	0.441	0.501	0.513
Clean	-0.707	-0.372	0.393	0.299	0.423

NOTE: All variables first averaged at the industry-occupation level. Years of education is imputed following Card et al. (2013). Task values imputed from verbal frequencies. Terms π_x and ϕ_x indicate residuals from a regression of wage effects on job characteristics. All results weighted by employment; see appendix for confidence intervals.

Appendix B

Supplementary Materials For Chapter 2

B.1 Proofs of Main Theoretical Results

UI and Efficiency (Informal). In this section I take as fixed worker assignment and suppress j -notation. In addition I set $\psi = 1$, which simplifies but does not substantially change the proof.

Writing the worker's flow value of unemployment as

$$\begin{aligned} U(s) &= \frac{(\rho + \delta) \frac{B(i,s)}{w(i,s)} + \zeta \theta(i,s)^\eta}{\rho + \delta + \zeta \theta(i,s)^\eta} w(i,s) \\ &\equiv G(i,s) w(i,s) , \end{aligned}$$

and defining the elasticity of G with respect to θ as $\epsilon_G(i,s)$ (which may depend on $\theta(i,s)$), we can show that total (i,s) -vacancies satisfy

$$V(i,j) = \frac{\alpha(i,s)^\sigma [m(s)\zeta\theta(i,s)^{\eta-1}]^{\sigma-1} \left[\frac{1}{C(i,s)} \frac{\epsilon_G(i,s)}{\epsilon_G(i,s)+\psi(1-\eta)} \right]^\sigma}{\left(\int \alpha_i(k) \left[\alpha(i,s)m(s)\zeta\theta(i,k)^{\eta-1} \frac{1}{C(i,k)} \frac{\epsilon_G(i,s)}{\epsilon_G(i,s)+\psi(1-\eta)} \right]^{\sigma-1} dk \right)^{\frac{\sigma}{\sigma-1}}} Y_i^* .$$

Fixing any two labor types s and s' , the vacancy ratio $\frac{V(i,s')}{V(i,s)}$ will be equal to

$$\frac{V(i,s')}{V(i,s)} = \left(\frac{\alpha(i,s')C(i,s)}{\alpha(i,s)C(i,s')} \right)^\sigma \left(\frac{[\theta(i,s')]^{\eta-1}]^{\sigma-1} \left[\frac{\epsilon_G(i,s')}{\epsilon_G(i,s)+\psi(1-\eta)} \right]^\sigma}{[\theta(i,s)]^{\eta-1}]^{\sigma-1} \left[\frac{\epsilon_G(i,s)}{\epsilon_G(i,s)+\psi(1-\eta)} \right]^\sigma} \right) .$$

Efficiency requires that the term in parenthesis be independent of i , which will only be the case when $\epsilon_G(i,s)$ is a constant.

To see this, consider a central planner that directly allocates vacancies and job applicants across submarkets in order to maximize the utility of the representative household. I simplify by assuming $C(i, j) = C(i)$. Writing the planner's Bellman equation:

$$\begin{aligned} \rho S^V(N, Y) = & \left(\sum_i \beta(i) Y(i)^{\frac{\tau-1}{\tau}} \right)^{\frac{\tau}{\tau-1}} + \\ & \max_{\phi, v} \left[\left(\delta[\kappa(s) - N(s)] - \sum_i \zeta v(i, s)^\eta (N(s)\phi(i, s))^{1-\eta} \right) \frac{\partial V}{\partial N(s)} - \sum_i \int C(i)v(i, s) + \right. \\ & \left. \left(\left[\int \alpha(i, s) \left(\zeta v(i, s)^\eta (N(s)\phi(i, s))^{1-\eta} m(s) \right)^{\frac{\sigma}{\sigma-1}} ds \right]^{\frac{\sigma-1}{\sigma}} - \delta Y(i) \right) \frac{\partial V}{\partial Y(i)} \right]. \end{aligned}$$

As the planner is free to shift workers across i , we can define a shadow price $\omega(s)$, and by taking the first-order conditions of the planner's problem we may derive the following condition for market tightness:

$$\theta^P(i, s) = \frac{\eta}{1 - \eta} \frac{\omega(s)}{C(i)N(s)}.$$

At the planner's solution, market tightness is multiplicatively separable in industry and worker type. But in the market equilibrium, θ will only be separable if ϵ_G is a constant. More explicitly, if we define $H(\theta) = \frac{G(i, s)}{\zeta \theta^{\eta-1} \epsilon_G(i, s)}$ we can show that

$$\theta^M(i, s) = H^{-1} \left(\frac{U(s)}{(1 - \eta)(\rho + \delta)C(i)} \right).$$

For this to be separable it must be that H has a constant elasticity with respect to market tightness, which in turn implies the same for G and hence that ϵ_G is constant.

Supposing then that $\epsilon_G(i, s) = D$ for some constant D , the market equilibrium will result in market tightness equal to

$$\theta^M(i, s) = \frac{\sum_k V_k^M(s) C(k)^{\frac{1}{D+1-\eta}}}{N(s) C(i)^{\frac{1}{D+1-\eta}}},$$

and search probabilities will take the form

$$\frac{\phi^M(i', s)}{\phi^M(i, s)} = \frac{V(i')^M(s)C(i')^{\frac{1}{D+1-\eta}}}{V(i)^M(s)C(i)^{\frac{1}{D+1-\eta}}}.$$

From here it is straightforward to show that by setting $D = \eta - 1 + 1/\kappa$ we arrive at the functional form for $B(i, s)$ given in assumption 3 of the main text. But from the planner's problem we have that

$$\frac{\phi^P(i', s)}{\phi^P(i, s)} = \frac{V(i')^P(s)C(i')}{V(i)^P(s)C(i)}.$$

Therefore we must have $D = \eta$ at the efficient allocation, or correspondingly $\kappa = 1$.

Occupational Assignment. This proof is an adaptation of that given by Costinot and Vogel (2010) to the case of search frictions and firm heterogeneity. I begin with the following notation:

- let $\omega(s)$ denote the set of jobs chosen by skill type s in at least one industry: $\omega(s) = \{j \mid \exists i \text{ s.t. } \phi^*(i, j, s) > 0\}$.
- let $\mu(j)$ denote the set of skill types s that choose job j in at least one industry: $\mu(j) = \{s \mid \exists i \text{ s.t. } \phi^*(i, j, s) > 0\}$.

Lemma 1 establishes the correspondence between skill and occupation, and lemmas 2 and 3 the differential equations characterizing optimal assignment.

Lemma 1. *There exists a continuous and strictly increasing function $\lambda : [\underline{s}, \bar{s}] \rightarrow [0, 1]$, independent of i , such that $\phi^*(i, j, s) > 0$ if and only if $\lambda(s) = j$, and where $\lambda(\underline{s}) = 0$ and $\lambda(\bar{s}) = 1$.*

Proof. That $\omega(s)$ is non-empty follows from $N(s) > 0$, which will be true if $\delta > 0$ and $\nu(s) > 0$, which I assume. Because all workers are assumed to search we must have $\phi^*(i, j, s) > 0$ for at least one (i, j) pair.

Regarding non-emptiness of $\mu(j)$ suppose that $\mu(j)$ is empty for some j . By assumption we have $\alpha_i(j) > 0$, and so it must be that if $\omega(s)$ is non-empty and prices are strictly positive, then

$V(i, j) > 0$ for all i and j provided that firms rationally expect that $\phi(i, j, s) > 0$ for at least one s . At the same time $\omega(s)$ non-empty implies that for any s , there exists at least one i' and j' for which $\phi^*(i', j', s) > 0$. Now if $\phi^*(i, j, s) = 0$ for all i and all s , then $\frac{\int \phi^*(i, j, s)N(s)ds}{\int \phi^*(i', j', s)N(s)ds} = 0$. However, defining $\bar{m}(i, j) > 0$ to be the firm's expected worker productivity, from workers' first-order condition we have that

$$\frac{\int \phi^*(i, j, s)N(s)ds}{\int \phi^*(i', j', s)N(s)ds} = \frac{V(i, j) \left[(C(i, j) \frac{m(j, s)}{\bar{m}(i, j)})^{\kappa\psi} A(i)^{\kappa(1-\psi)} \right]^{\frac{1}{\eta + \kappa\psi(1-\eta)}}}{V(i', j') \left[(C(i', j') \frac{m(j', s)}{\bar{m}(i', j')})^{\kappa\psi} A(i')^{\kappa(1-\psi)} \right]^{\frac{1}{\eta + \kappa\psi(1-\eta)}}} > 0,$$

a contradiction.

Third, $\mu(j)$ is non-decreasing. Suppose otherwise. From the worker's first-order condition, we must have

$$\begin{aligned} 0 &\geq \theta(i, j)^{\eta + \kappa\psi(1-\eta)} \left(\frac{\kappa\psi(1-\eta)}{\eta} (\rho + \delta) \frac{C(i, j)}{\bar{m}(i, j)} \right)^{\kappa\psi} (\zeta A(i))^{\kappa(1-\psi)} m(j, s)^{\kappa\psi} - U(s) \\ &\equiv R(i, j) m(j, s)^{\kappa\psi} - U(s), \end{aligned}$$

with equality if $\phi^*(i, j, s) > 0$. Supposing that there exist two industries i and i' , two jobs $j^+ > j^-$, and two worker types $s^+ > s^-$ such that $\phi^*(i, j^+, s^-) > 0$ and $\phi^*(i', j^-, s^+) > 0$, it must be that

$$\begin{aligned} 0 &= R(i, j^-) m(j^-, s^+)^{\psi} - U(s^+) \\ &\geq R(i, j^-) m(j^-, s^-)^{\psi} - U(s^-) \\ &= \frac{m(j^-, s^-)^{\psi}}{m(j^-, s^+)^{\psi}} U(s^+) - U(s^-) \\ &> \frac{m(j^+, s^-)^{\psi}}{m(j^+, s^+)^{\psi}} U(s^+) - U(s^-) \\ &= -\frac{U(s^-)}{m(j^+, s^+)^{\psi} R(i', j^+)} \left(R(i', j^+) m(j^+, s^+)^{\psi} - U(s^+) \right) \\ &\geq 0, \end{aligned}$$

a contradiction for $\psi > 0$. Note that the result holds across industries due to assumption 1, without which there would be a non-separable term depending on (i, j, s) .

Fourth, ω and μ are single-valued almost everywhere. The proof is unchanged from CV and so

I provide only the intuition: if ω (or μ) has positive measure over a domain with positive measure, then from the previous result the range of the correspondence will have measure greater than the measure of $[0, 1]$ (or $[\underline{s}, \bar{s}]$), a contradiction.

Fifth, $\mu(j)$ is single-valued. If this is not the case, then from step 3 there exists a non-degenerate interval $[s, s']$ in which all workers choose job j . Step 4 implies that there exists another job j' that is chosen by a single worker type. From the firm's first-order condition and market clearing it must be that

$$\sum_i \left(\int \phi^*(i, j, s) N(s) ds \right) \zeta \theta(i, j)^\eta = \sum_i \frac{\left[\alpha(i, j) \frac{\zeta \theta(i, j)^{\eta-1}}{C(i, j)} \right]^\sigma [\bar{m}(i, j)]^{\sigma-1}}{\left(\int \alpha(i, k) \left[\alpha(i, k) \frac{\zeta \theta(i, k)^{\eta-1}}{C(i, k)} \bar{m}(i, k) \right]^{\sigma-1} dk \right)^{\frac{\sigma}{\sigma-1}}} Y(i)^* .$$

But then

$$\frac{\sum_i \frac{\left[\alpha(i, j') \frac{\zeta \theta(i, j')^{\eta-1}}{C(i, j')} \right]^\sigma [\bar{m}(i, j')]^{\sigma-1}}{\left(\int \alpha(i, k) \left[\alpha(i, k) \frac{\zeta \theta(i, k)^{\eta-1}}{C(i, k)} \bar{m}(i, k) \right]^{\sigma-1} dk \right)^{\frac{\sigma}{\sigma-1}}}}{\sum_i \frac{\alpha(i, j) \left[\alpha(i, j) \frac{\zeta \theta(i, j)^{\eta-1}}{C(i, j)} \right]^\sigma [\bar{m}(i, j)]^{\sigma-1}}{\left(\int \alpha(i, k) \left[\alpha(i, k) \frac{\zeta \theta(i, k)^{\eta-1}}{C(i, k)} \bar{m}(i, k) \right]^{\sigma-1} dk \right)^{\frac{\sigma}{\sigma-1}}}} = 0 ,$$

which violates the assumptions that $\alpha(i, j)$ is strictly positive and continuous (and therefore finite), that $m(j, s) > 0$, and that $C(i, j)$ is finite.

From the last step we have $\mu(j)$ single-valued; from the third step, weakly increasing; from the first step, continuous and such that $\mu(0) = \underline{s}$ and $\mu(1) = \bar{s}$; and from the fourth step, μ is strictly increasing. Hence we have a continuous, strictly increasing bijection $\lambda(s) = \omega(s) = \{j \mid \exists i \text{ s.t. } \phi^*(i, j, s) = 1\} = \mu^{-1}(s)$. \square

Lemma 2. *Reservation values satisfy the equation*

$$\frac{d \log U(s)}{ds} = \kappa \psi \frac{m_s(s, \lambda(s))}{m(s, \lambda(s))} .$$

Proof. From lemma 1, and with $R(i, j)$ defined as above, for any $\phi^*(i, j, s) > 0$ we must have

$$0 \geq R(i, j) m(j, s)^{\kappa \psi} - U(s) ,$$

and following CV the following two inequalities must hold:

$$\begin{aligned} R(i, \lambda(s)) - \frac{U(s)}{m(\lambda(s), s)^{\kappa\psi}} &\geq R(i, \lambda(s)) - \frac{U(s+ds)}{m(\lambda(s), s+ds)^{\kappa\psi}} \\ R(i, \lambda(s+ds)) - \frac{U(s+ds)}{m(\lambda(s+ds), s+ds)^{\kappa\psi}} &\geq R(i, \lambda(s+ds)) - \frac{U(s)}{m(\lambda(s+ds), s)^{\kappa\psi}}, \end{aligned}$$

and therefore

$$\begin{aligned} R(i, \lambda(s)) &\left[m(\lambda(s), s+ds)^{\kappa\psi} - m(\lambda(s), s)^{\kappa\psi} \right] \\ &\leq U(s+ds) - U(s) \\ &\leq R(i, \lambda(s+ds)) \left[m(\lambda(s+ds), s+ds)^{\kappa\psi} - m(\lambda(s+ds), s)^{\kappa\psi} \right]. \end{aligned}$$

Regarding continuity in s : U is continuous given continuity of θ , which in turn follows from continuity of α , ν , and λ . If θ is continuous then continuity of $R(i, j)$ also follows, and hence we can divide by ds and take the limit of the previous inequalities to show that

$$\begin{aligned} U'(s) &= R(i, \lambda(s))\kappa\psi m_s(\lambda(s), s)m(\lambda(s), s)^{\kappa\psi-1} \\ &= U(s)\kappa\psi \frac{m_s(\lambda(s), s)}{m(\lambda(s), s)}, \end{aligned}$$

and the result follows. □

Lemma 3. *The matching function satisfies*

$$\begin{aligned} \frac{d\lambda(s)}{ds} &= \frac{m(\lambda(s), s)N(s)U(s)}{\sum_i \left[\left(\frac{\kappa\psi(1-\eta)}{\eta+\kappa\psi(1-\eta)} p(i, \lambda(s)) m(\lambda(s), s) \right)^\psi A(i)^{(1-\psi)} \right]^\kappa y(i, \lambda(s))} \\ \frac{d\lambda(s)}{ds} &= \frac{\left[\frac{\eta+\psi(1-\eta)}{\psi(1-\eta)} \right]^\psi m(\lambda(s), s)U(s)N(s)}{\sum_i \left[m(\lambda(s), s)p_i(\lambda(s)) \right]^\psi A_i^{1-\psi} \frac{[\alpha_i(\lambda(s))/p_i(\lambda(s))]^\sigma}{\left[\int \frac{\alpha_i(\lambda(k))^\sigma}{p_i(\lambda(k))^{\sigma-1}} dk \right]^{\frac{\sigma}{\sigma-1}}} \frac{[\beta_i/P_i]^\tau}{\left[\sum_k \frac{\beta_k^\tau}{P_k^{\tau-1}} \right]^{\frac{\tau}{\tau-1}}} Y}, \end{aligned}$$

where $\lambda(\underline{s}) = 0$ and $\lambda(\bar{s}) = 1$.

Proof. Total demand for j -output is given by

$$\sum_i y^D(i, j) = \sum_i y(i, j),$$

while from market-clearing and lemma 1 we have supply of j -output as

$$y^S(i, j) = \sum_i m(\lambda(s), s) \phi^*(i, j, s) \zeta \theta(i, \lambda(s))^\eta N(s) \delta[j - \lambda(s)] ,$$

with δ the Dirac function. Combining these equations:

$$\sum_i y^D(i, j) = \int m(j, \lambda^{-1}(j')) \sum_i [\phi^*(i, j, \lambda^{-1}(j')) \zeta \theta(i, j)^\eta N(\lambda^{-1}(j'))] \delta[\lambda(s) - j'] \frac{1}{\lambda'(\lambda^{-1}(j'))} dj' .$$

This simplifies to

$$\lambda'(s) = \frac{m(\lambda(s), s) N(s) \sum_i \phi^*(i, \lambda(s), s) \zeta \theta(i, \lambda(s))^\eta}{\sum_i y(i, \lambda(s))} ,$$

where

$$\phi^*(i, \lambda(s), s) = \frac{\frac{y(i, \lambda(s))}{m(\lambda(s), s)} \left[m(\lambda(s), s) p(i, \lambda(s)) \right]^{\kappa\psi} A(i)^{\kappa(1-\psi)}}{\sum_k \frac{y(k, \lambda(s))}{m(\lambda(s), s)} \left[m(\lambda(s), s) p(k, \lambda(s)) \right]^{\kappa\psi} A(k)^{\kappa(1-\psi)}} ,$$

$$\theta(i, \lambda(s)) = \left[\frac{U(s)}{\zeta \left(\frac{\kappa\psi(1-\eta)}{\eta + \kappa\psi(1-\eta)} m(\lambda(s), s) p(i, \lambda(s)) \right)^{\kappa\psi} A(i)^{\kappa(1-\psi)}} \right]^{\frac{1}{\eta}} .$$

Substitution into the previous equation gives us the final result. □

B.2 Equilibrium Characterization

Taking the worker's problem and substituting out S^U and S^W , we obtain an equation for the worker's outside option:

$$\rho U(i, j, s) = f(\theta(i, j)) u(w(i, j) m(j, s), A(i))^\kappa .$$

From the firm's first-order condition we then have the wage function

$$w(i, j) = \frac{\kappa\psi(1-\eta)}{\eta + \kappa\psi(1-\eta)} p(i, j) ,$$

which allows us to then solve for the worker's search probability,

$$\frac{\phi^*(i', j', s)}{\phi^*(i, j, s)} = \frac{\frac{p(i', j')y(i', j')}{C(i', j')} \left[(p(i', j')m(j', s))^{\kappa\psi} A(i')^{\kappa(1-\psi)} \right]^{\frac{1}{\eta}}}{\frac{p(i, j)y(i, j)}{C(i, j)} \left[(p(i, j)m(j, s))^{\kappa\psi} A(i)^{\kappa(1-\psi)} \right]^{\frac{1}{\eta}}},$$

and to characterize market tightness:

$$\theta(i, j) = \left(\frac{U(s)}{\zeta [(1 - \epsilon(i, j))m(j, s)p(i, j)]^{\kappa\psi} A(i)^{\kappa(1-\psi)}} \right)^{\frac{1}{\eta}}.$$

Free entry yields solutions for the prices, which are standard, while likewise the optimal bundles y^* and Y^* are straightforward to derive.

The system of equations describing the equilibrium can be written more concisely by defining the skill premium and the firm premium as

$$SP(s) = U(s)^{\frac{1-\eta}{\eta+\kappa\psi(1-\eta)}}$$

$$FP(i, j) = \left(\frac{\kappa\psi(1-\eta)(\rho+\delta)C(i, j)}{\eta\zeta^{\frac{1}{\eta}} A(i)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}}} \right)^{\frac{\eta}{\eta+\kappa\psi(1-\eta)}}.$$

Prices and wages are then given by the equations

$$p(i, j) = \frac{\eta + \kappa\psi(1-\eta)}{\kappa\psi(1-\eta)} \frac{SP(\lambda^{-1}(j))FP(i, j)}{m(j, \lambda^{-1}(j))}$$

$$P(i) = \left(\int \alpha(i, j) \left(\frac{\alpha(i, j)}{p(i, j)} \right)^{\sigma-1} dj \right)^{\frac{-1}{\sigma-1}},$$

while the quantities that characterize labor submarkets and goods markets are

$$\phi^*(i, j, \lambda^{-1}(j)) = \frac{V(i, j)FP(i, j)^{\frac{\kappa\psi}{\eta}} A(i)^{\frac{\kappa(1-\psi)}{\eta}}}{\sum_k V(k, j)FP(k, j)^{\frac{\kappa\psi}{\eta}} A(k)^{\frac{\kappa(1-\psi)}{\eta}}}$$

$$V(i, j) = \frac{1}{\zeta^{\frac{1}{\eta}}} \frac{SP(\lambda^{-1}(j))}{[FP(i, j)^{\kappa\psi} A(i)^{\kappa(1-\psi)}]^{\frac{1-\eta}{\eta}}} \frac{y(i, j)}{m(j, \lambda^{-1}(j))}$$

$$y(i, j) = \left(\frac{\alpha(i, j)}{p(i, j)} \right)^{\sigma} \beta(i)^{\tau} P(i)^{\sigma-\tau}$$

$$\theta(i, j) = \frac{SP(s)^{\frac{1}{1-\eta}}}{\zeta^{\frac{1}{\eta}} FP_i(\lambda(s))^{\frac{\kappa\psi}{\eta}} A_i^{\frac{\kappa(1-\psi)}{\eta}}} .$$

Note that we may also write the system of differential equations describing the optimal assignment in terms of wage premia:

$$\begin{aligned} \frac{SP'(s)}{SP(s)} &= \frac{\kappa\psi(1-\eta)}{\eta + \kappa\psi(1-\eta)} \frac{m_s(\lambda(s), s)}{m(\lambda(s), s)} \\ \lambda'(s) &= \frac{m(\lambda(s), s)N(s)SP(s)^{\frac{\eta}{1-\eta}}}{\sum_k FP(k, j)^{\kappa\psi} A(i)^{\kappa(1-\psi)} y(i, j)} . \end{aligned}$$

B.3 Structural Estimation

Match parameters. The empirical match function $\hat{\lambda}$ and the skill productivity function $G(s)$ are identified by inverting the differential equation for $U(s)$:

$$G'(s)(\gamma_s + \gamma_{sj}\hat{\lambda}(s)) = \frac{\eta + \kappa\psi(1-\eta)}{\kappa\psi(1-\eta)} \frac{SP'(s)}{SP(s)} ,$$

where $\gamma_{sj} = 1$ for the 2010-2017 panel, while in the earlier panels we have from the boundary condition that

$$\gamma_{sj} = 1/\hat{\lambda}(\bar{s}) .$$

The term $H(j)$ is found by solving the equation

$$H(j) = \log \left(\frac{SP(s)}{e^{G(s)(\gamma_s+j)}} \right) ,$$

which may be solved by inverting $\hat{\lambda}$.

Amenities and costs. Given κ , industry amenities are estimated from empirical vacancy-filling rates calculated as total industry hires over vacancies. Given industry vacancies V

and applicants N , the predicted vacancy-filling rate will be

$$\frac{M(V, N)}{V} = SA^{\frac{(1-\psi)(1-\eta)}{\eta}} \frac{\int \left(\frac{\alpha(\lambda(s))}{FP(\lambda(s))SP(s)} \right)^\sigma m(\lambda(s), s)^{\sigma-1} ds}{\int \left(\frac{\alpha(\lambda(s))}{FP(\lambda(s))SP(s)} \right)^\sigma m(\lambda(s), s)^{\sigma-1} \frac{SP(s)}{FP(\lambda(s))^{\frac{\psi(1-\eta)}{\eta}}} ds},$$

where S is a constant term that may be dropped as amenities are only identified up to a multiplicative constant. We may then write amenities as a function of the vacancy-filling rate and a labor-weighted ratio of wage premia:

$$A = \left(\frac{M(V, N)}{V} \frac{\int L(\lambda(s)) \frac{SP(s)}{FP(\lambda(s))^{\frac{\kappa\psi(1-\eta)}{\eta}}} ds}{\int L(\lambda(s)) ds} \right)^{\frac{\eta}{\kappa(1-\psi)(1-\eta)}}.$$

When estimating this equation, SP and FP are averaged for each of the 20 within-industry estimation quantiles, and a weighted average is then taken of the ratio above.

Industry entry costs are obtained using the free entry condition

$$C(i, j) = \frac{FP(i, j)^{\frac{\eta+\kappa\psi(1-\eta)}{\eta}} A(i)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}} \eta \zeta^{\frac{1}{\eta}}}{\kappa\psi(1-\eta)(\rho + \delta)}.$$

Taking the ratio of $C(i', j')$ to $C(i, j)$, we obtain

$$\frac{C(i', j')}{C(i, j)} = \left(\frac{FP(i', j')}{FP(i, j)} \right)^{\frac{\eta+\kappa\psi(1-\eta)}{\eta}} \left(\frac{A(i')}{A(i)} \right)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}},$$

which gives us entry costs up to a multiplicative constant, conditional on amenities, firm premia FP , and the parameters ψ and η which are normalized as described in the text.

Demand functions. The occupational demand function $\alpha(i, j)$ is identified from the intra-industry distribution of person effects $SP(s)$, calculated over 20 quantiles. Denoting sub-

market labor (or alternatively, flow matches) as $L(i, j)$, the model predicts that

$$\frac{w(i, \lambda(s))L(i, \lambda(s))}{\int w(i, \lambda(k))L(i, \lambda(k))dk} = \frac{\alpha(i, \lambda(s))^\sigma \left(\frac{m(\lambda(s), s)}{FP(i, \lambda(s))SP(s)} \right)^{\sigma-1}}{\int \alpha(i, \lambda(k))^\sigma \left(\frac{m(\lambda(k), s)}{FP(i, \lambda(k))SP(k)} \right)^{\sigma-1} dk}.$$

Demand is therefore estimated by solving for α and substituting empirical wage effects, intra-industry labor shares, and the estimated matching function, and finally normalizing to ensure that $\int \alpha(i, j)dj = 1$

To obtain industry shares $\beta(i)$, we may integrate over total wages $w(i, \lambda(s))L(i, \lambda(s))$ to obtain

$$\int w(i, \lambda(s))L_i(\lambda(s))ds = \left(\beta(i) \frac{\psi(1-\eta)}{\eta + \kappa\psi(1-\eta)} \right)^\tau \times \left(\int \alpha(i, \lambda(k))^\sigma \left[\frac{m(\lambda(k), k)}{FP(i, \lambda(k))SP(k)} \right]^{\sigma-1} dk \right)^{\frac{\tau-1}{\sigma-1}}.$$

By solving for β and substituting empirical values as well as α (estimated previously), we may then estimate industry demand, which is normalized through division by $\sum_i \beta(i)$.

Matching function. At each point in time the flow of matches will be equal to

$$M(V, N) = \int \zeta \left(\frac{\delta y^*(i, \lambda(s))}{\zeta \theta(i, \lambda(s))^{\eta-1} m(\lambda(s), s)} \right)^\eta (\phi^*(i, \lambda(s), s)N(s))^{1-\eta}.$$

I do not observe s -unemployment and so I assume that changes over time in N are in equal proportion for all skill types: that is, $N(s) \equiv N(s)U$ where U is aggregate unemployment and $N(s)$ is constant. We can rewrite the previous equation as

$$M(V, N) = SVU^{1-\eta} A(i)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}} \left(\frac{\int L(i, \lambda(s)) \frac{SP(s)}{FP(i, \lambda(s))^{\frac{\psi(1-\eta)}{d}} s}}{\int L(i, \lambda(s)) ds} \right)^{-1},$$

where S is a constant. If wage premia are constant in the short-run, then by taking differences

we obtain an equation in V and U , allowing η and ζ to be estimated from industry vacancies and hires and aggregate unemployment.

However, estimates obtained in this fashion result in values of η that are implausibly large (in the .6-.7 range) relative to estimates in the literature, likely due to the short length of the time series and noise in the estimates of V and hires. Attenuation bias downwards estimates of $1 - \eta$, resulting in upwardly biased estimates of η . I therefore use the value $\eta = .35$ estimated by Kohlbrecher et al. (2016). Matching efficiency ζ is then found by taking the ratio of predicted to actual hires.

Table B.1: Empirical and Predicted Wage Moments

	1993-99	1998-04	2003-10	2010-17
<i>West Germany</i>				
Empirical Variance	0.1572	0.1901	0.2212	0.2221
Var(person)	0.0996	0.1160	0.1310	0.1328
Var(firm)	0.0305	0.0373	0.0516	0.0394
Covariance	0.0071	0.0109	0.0117	0.0167
Between-ind./-occ.	0.0745	0.0914	0.1096	0.1185
Var(person)	0.0408	0.0477	0.0515	0.0622
Var(firm)	0.0161	0.0193	0.0257	0.0203
Covariance	0.0088	0.0122	0.0162	0.0180
Model-predicted	0.0695	0.0861	0.1021	0.1106
Var(person)	0.0426	0.0491	0.0540	0.0641
Var(firm)	0.0102	0.0131	0.0173	0.0131
Covariance	0.0083	0.0119	0.0154	0.0167
<i>East Germany</i>				
Empirical Variance	0.1590	0.1712	0.1955	0.1816
Var(person)	0.0720	0.0889	0.0948	0.1031
Var(firm)	0.0380	0.0451	0.0574	0.0432
Covariance	0.0142	0.0127	0.0123	0.0178
Between-ind./-occ.	0.0914	0.1013	0.1101	0.1188
Var(person)	0.0369	0.0438	0.0453	0.0555
Var(firm)	0.0247	0.0263	0.0324	0.0261
Covariance	0.0149	0.0156	0.0162	0.0186
Model-predicted	0.0814	0.0912	0.0974	0.1121
Var(person)	0.0390	0.0455	0.0478	0.0588
Var(firm)	0.0144	0.0153	0.0175	0.0154
Covariance	0.0140	0.0152	0.0161	0.0189

NOTE: For between-group moments, wage and wage effects are averaged by industry-occupation cells using 3-digit codes.

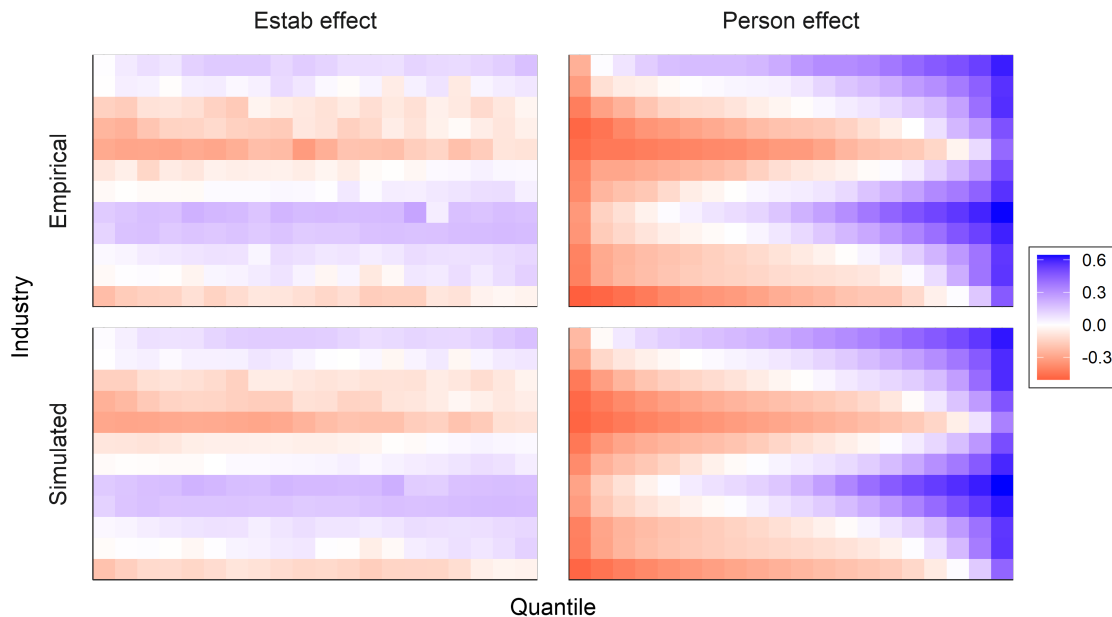


Figure B.1: Empirical and Predicted Wage Distribution

B.4 Quantitative Experiments

Experiment: Homogeneous FP. The experiment performed in the text is to set firm premia equal so that $FP(i, j)$ is constant, without affecting the distribution of vacancies *conditional on occupational assignment* in panel 0 (indicated by subscripts). Vacancies are given by:

$$V(i, j) = \frac{1}{\zeta^{\frac{1}{\eta}}} \frac{SP(\lambda^{-1}(j))}{[FP(i, j)^{\kappa\psi} A(i)^{\kappa(1-\psi)}]^{\frac{1-\eta}{\eta}}} \frac{\left(\frac{\alpha(i, j)m(j, \lambda^{-1}(j))(1-\epsilon(i, j))}{SP(\lambda^{-1}(j))FP(i, j)} \right)^{\sigma} \beta(i)^{\tau} P(i)^{\sigma-\tau}}{m(j, \lambda^{-1}(j))} .$$

By assumption we have $\sigma = \tau$, and therefore consolidating terms that are unchanged or ignored (i.e. SP) we have

$$V(i, j) = \hat{V} \frac{\alpha(i, j)^{\sigma} \beta(i)^{\tau}}{C(i, j)^{\frac{\sigma\eta + \kappa\psi(1-\eta)}{\eta + \kappa\psi(1-\eta)}}} .$$

Now for firm premia to be constant we must have

$$\hat{C}(i, j) = \hat{C} A(i)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}} .$$

Writing $R(i, j) \equiv C(i, j)/A(i)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}}$, in order for occupational shares to be constant we must have

$$\hat{\alpha}(i, j) = \frac{\alpha(i, j)/R_0(i, j)^{\frac{\sigma\eta+\kappa\psi(1-\eta)}{\sigma(\eta+\kappa\psi(1-\eta))}}}{\int \alpha(i, k)/R_0(i, k)^{\frac{\sigma\eta+\kappa\psi(1-\eta)}{\sigma(\eta+\kappa\psi(1-\eta))}} dk}.$$

It then follows that

$$\hat{\beta}(i) = \frac{\beta(i) \left(\int \alpha_0(i, j)/R_0(i, j)^{\frac{\sigma\eta+\kappa\psi(1-\eta)}{\sigma(\eta+\kappa\psi(1-\eta))}} dj \right)}{\sum_k \beta(k) \left(\int \alpha_0(k, j)/R_0(k, j)^{\frac{\sigma\eta+\kappa\psi(1-\eta)}{\sigma(\eta+\kappa\psi(1-\eta))}} dj \right)}.$$

Finally, maintaining the overall level of vacancies requires in the initial panel that

$$\hat{C}_0 = \left(\sum_i \beta_0(i) \left(\int \alpha_0(i, j)/R_0(i, j)^{\frac{\sigma\eta+\kappa\psi(1-\eta)}{\sigma(\eta+\kappa\psi(1-\eta))}} dj \right) \right)^{\frac{-\tau(\eta+\kappa\psi(1-\eta))}{\sigma\eta+\kappa\psi(1-\eta)}},$$

and as there is no simple way to implement a similar condition on subsequent panels, I hold total vacancies approximately constant by assuming that

$$\hat{C} = \left(\frac{\sum_i \beta(i) \left(\int \alpha(i, j) \left(\frac{R(i, j)}{R_0(i, j)} \right)^{\frac{\sigma\eta+\kappa\psi(1-\eta)}{\sigma(\eta+\kappa\psi(1-\eta))}} dj \right)}{\sum_i \beta(i) \left(\int \alpha_0(i, j)/R_0(i, j)^{\frac{\sigma\eta+\kappa\psi(1-\eta)}{\sigma(\eta+\kappa\psi(1-\eta))}} dj \right)} \right)^{\frac{\tau(\eta+\kappa\psi(1-\eta))}{\sigma\eta+\kappa\psi(1-\eta)}},$$

The counterfactual experiment is then to replace estimated parameter values with $\{\hat{C}, \hat{\alpha}, \hat{\beta}\}$ as defined above.

Experiment: Alternative FP. Several experiments involve setting firm premia to alternative (non-constant) values, holding constant the 1993-99 distribution of vacancies conditional on λ . As above we have

$$V(i, j) = \hat{V} \frac{\alpha(i, j)^\sigma \beta(i)^\tau}{C(i, j)^{\frac{\sigma\eta+\kappa\psi(1-\eta)}{\eta+\kappa\psi(1-\eta)}}}.$$

Defining \hat{C} to be the value of entry costs that implements the alternative firm premia (and noting \hat{C} may need to account for differences in the parameters ζ and δ), and with $R(i, j) \equiv C(i, j)/A(i)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}}$ as above, in order for occupational shares to be constant we must have

$$\hat{\alpha}(i, j) = \frac{\alpha(i, j) \left(\frac{R_{A,0}(i, j)}{R_0(i, j)} \right)^{\frac{\sigma\eta + \kappa\psi(1-\eta)}{\sigma(\eta + \kappa\psi(1-\eta))}}}{\int \alpha(i, k) \left(\frac{R_{A,0}(i, k)}{R_0(i, k)} \right)^{\frac{\sigma\eta + \kappa\psi(1-\eta)}{\sigma(\eta + \kappa\psi(1-\eta))}} dk},$$

where A denotes the target economy and 0 the initial panel. Likewise industry shares must be set so that

$$\hat{\beta}(i) = \frac{\beta(i) \left(\int \alpha_0(i, j) \left(\frac{R_{A,0}(i, j)}{R_0(i, j)} \right)^{\frac{\sigma\eta + \kappa\psi(1-\eta)}{\sigma(\eta + \kappa\psi(1-\eta))}} dj \right)}{\sum_k \beta(k) \left(\int \alpha_0(k, j) \left(\frac{R_{A,0}(k, j)}{R_0(k, j)} \right)^{\frac{\sigma\eta + \kappa\psi(1-\eta)}{\sigma(\eta + \kappa\psi(1-\eta))}} dj \right)}.$$

Maintaining the overall level of vacancies requires in the initial panel that

$$\hat{C}_0 = \left(\sum_i \beta_0(i) \left(\int \alpha_0(i, j) \left(\frac{R_{A,0}(i, j)}{R_0(i, j)} \right)^{\frac{\sigma\eta + \kappa\psi(1-\eta)}{\sigma(\eta + \kappa\psi(1-\eta))}} dj \right) \right)^{\frac{-\tau(\eta + \kappa\psi(1-\eta))}{\sigma\eta + \kappa\psi(1-\eta)}},$$

and as there is no simple way to implement a similar condition on subsequent panels, I hold total vacancies approximately constant by assuming that

$$\hat{C} = \left(\frac{\sum_i \beta(i) \left(\int \alpha(i, j) \left(\frac{R(i, j)R_{A,0}(i, j)}{R_0(i, j)} \right)^{\frac{\sigma\eta + \kappa\psi(1-\eta)}{\sigma(\eta + \kappa\psi(1-\eta))}} dj \right)}{\sum_i \beta(i) \left(\int \alpha_0(i, j) \left(\frac{R_{A,0}(i, j)}{R_0(i, j)} \right)^{\frac{\sigma\eta + \kappa\psi(1-\eta)}{\sigma(\eta + \kappa\psi(1-\eta))}} dj \right)} \right)^{\frac{\tau(\eta + \kappa\psi(1-\eta))}{\sigma\eta + \kappa\psi(1-\eta)}},$$

The counterfactual experiment is then to replace estimated parameter values with $\{\hat{C}, \hat{\alpha}, \hat{\beta}\}$ as defined above.

Experiment: Alternative Rent-Shares. Given a target set of firm premia FP_A , we can

implement FP_A by setting the firm's share of rent equal to

$$\hat{\eta}(i, j) = \frac{\frac{C(i, j)}{A(i, j)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}}}}{\frac{C(i, j)}{A(i, j)^{\frac{\kappa(1-\psi)(1-\eta)}{\eta}}} + \frac{\zeta^{\frac{1}{\eta}} FP_A}{\rho + \delta}}.$$

Doing so affects not only FP , but also output price $p(i, j)$ and vacancies $V(i, j)$.

B.5 Additional Results

Table B.2: Predicted and Counterfactual Wage Moments

	1993-99	1998-04	2003-10	2010-17
$\kappa = .5$				
Predicted Variance	0.0694	0.0860	0.1021	0.1106
Homogeneous FP	0.0412	0.0480	0.0527	0.0612
α, β only	0.0736	0.0862	0.0946	0.1066
Homogeneous FP	0.0436	0.0511	0.0538	0.0581
$\kappa = 1$				
Predicted Variance	0.0694	0.0860	0.1021	0.1106
Homogeneous FP	0.0399	0.0463	0.0504	0.0580
α, β only	0.0735	0.0862	0.0945	0.1065
Homogeneous FP	0.0421	0.0493	0.0515	0.0551
$\kappa = 2$				
Predicted Variance	0.0696	0.0859	0.1020	0.1106
Homogeneous FP	0.0372	0.0427	0.0457	0.0516
α, β only	0.0736	0.0861	0.0947	0.1068
Homogeneous FP	0.0392	0.0454	0.0467	0.049
$\kappa = 4$				
Predicted Variance	0.0693	0.0860	0.1021	0.1106
Homogeneous FP	0.0316	0.0357	0.0368	0.0394
α, β only	0.0735	0.0860	0.0943	0.1065
Homogeneous FP	0.0332	0.0378	0.0375	0.0376

NOTE: For between-group moments, wage and wage effects are averaged by industry-occupation cells using 3-digit codes.




Figure B.2: Predicted and Counterfactual Wage Trend, 1993-2017

Table B.3: Predicted and Counterfactual Wage Moments, By Parameter

	1993-99	1998-04	2003-10	2010-17
$\kappa = .5$				
β only	0.0859	0.0894	0.0926	0.0945
Homogeneous <i>FP</i>	0.0516	0.0524	0.0526	0.0507
α only	0.0766	0.0861	0.0912	0.1009
Homogeneous <i>FP</i>	0.0430	0.0498	0.0522	0.0584
<i>C</i> only	0.0905	0.0956	0.0995	0.0912
Homogeneous <i>FP</i>	0.0511	0.0512	0.0512	0.0513
ν only	0.0856	0.0853	0.0879	0.0930
Homogeneous <i>FP</i>	0.0482	0.0479	0.0498	0.0535
$\kappa = 1$				
β only	0.0858	0.0894	0.0926	0.0944
Homogeneous <i>FP</i>	0.0497	0.0504	0.0503	0.0482
α only	0.0766	0.0860	0.0911	0.1008
Homogeneous <i>FP</i>	0.0413	0.0478	0.0500	0.0557
<i>C</i> only	0.0904	0.0956	0.0995	0.0911
Homogeneous <i>FP</i>	0.0489	0.0489	0.0489	0.0490
ν only	0.0855	0.0852	0.0878	0.0929
Homogeneous <i>FP</i>	0.0461	0.0459	0.0476	0.0511
$\kappa = 2$				
β only	0.0858	0.0894	0.0927	0.0946
Homogeneous <i>FP</i>	0.0458	0.0461	0.0457	0.0432
α only	0.0767	0.0860	0.0913	0.1011
Homogeneous <i>FP</i>	0.0379	0.0437	0.0454	0.0503
<i>C</i> only	0.0908	0.0956	0.0996	0.0911
Homogeneous <i>FP</i>	0.0443	0.0443	0.0443	0.0444
ν only	0.0856	0.0853	0.0879	0.0930
Homogeneous <i>FP</i>	0.0419	0.0417	0.0433	0.0463
$\kappa = 4$				
β only	0.0856	0.0893	0.0925	0.0945
Homogeneous <i>FP</i>	0.0382	0.0377	0.0366	0.0336
α only	0.0766	0.0859	0.0909	0.1007
Homogeneous <i>FP</i>	0.0310	0.0358	0.0366	0.0399
<i>C</i> only	0.0904	0.0957	0.0996	0.0911
Homogeneous <i>FP</i>	0.0354	0.0354	0.0354	0.0355
ν only	0.0855	0.0852	0.0878	0.0929
Homogeneous <i>FP</i>	0.0338	0.0335	0.0348	0.0370

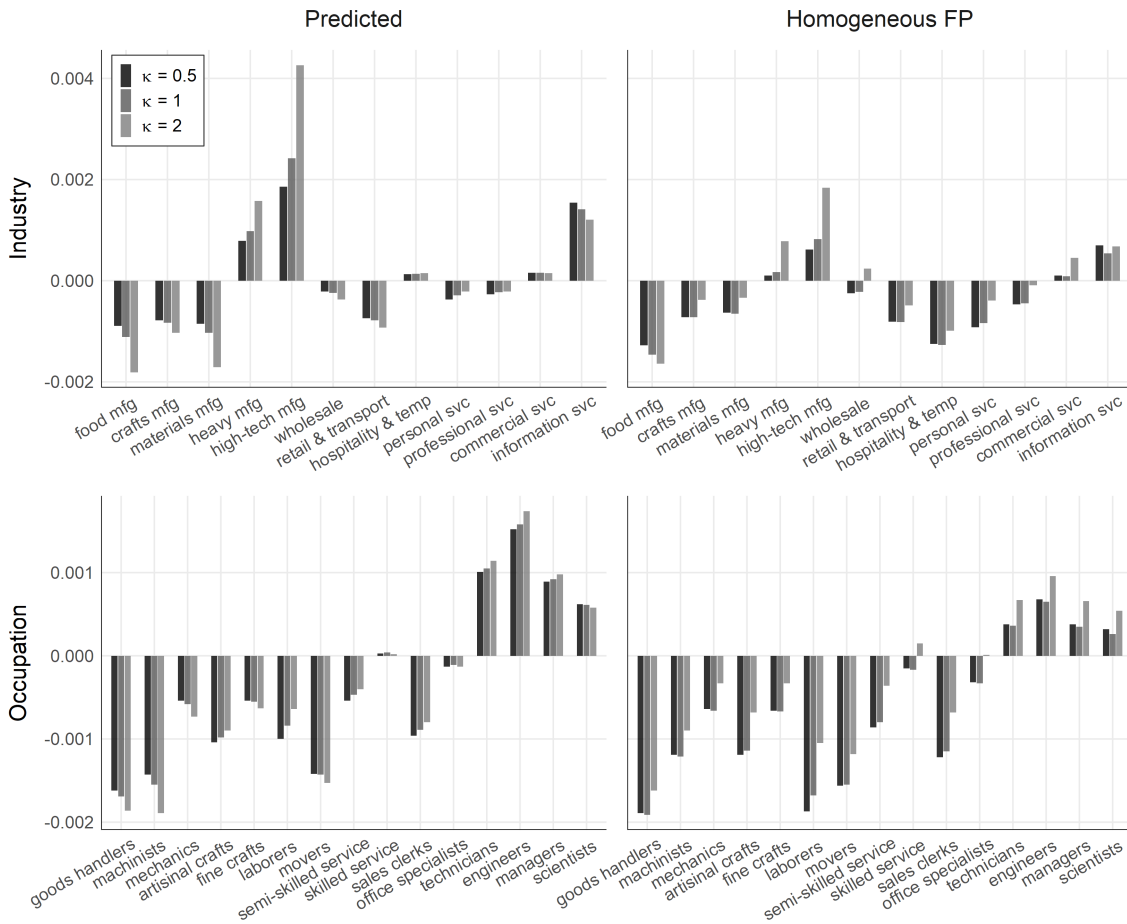


Figure B.3: Group Contribution to Change in Wage Variance, 1993-2017

NOTE: Effect on wage variance of a 1% increase in vacancy share.




Figure B.4: Estimated Type Distributions, East Germany




Figure B.5: Estimated Match Production Parameters, East Germany

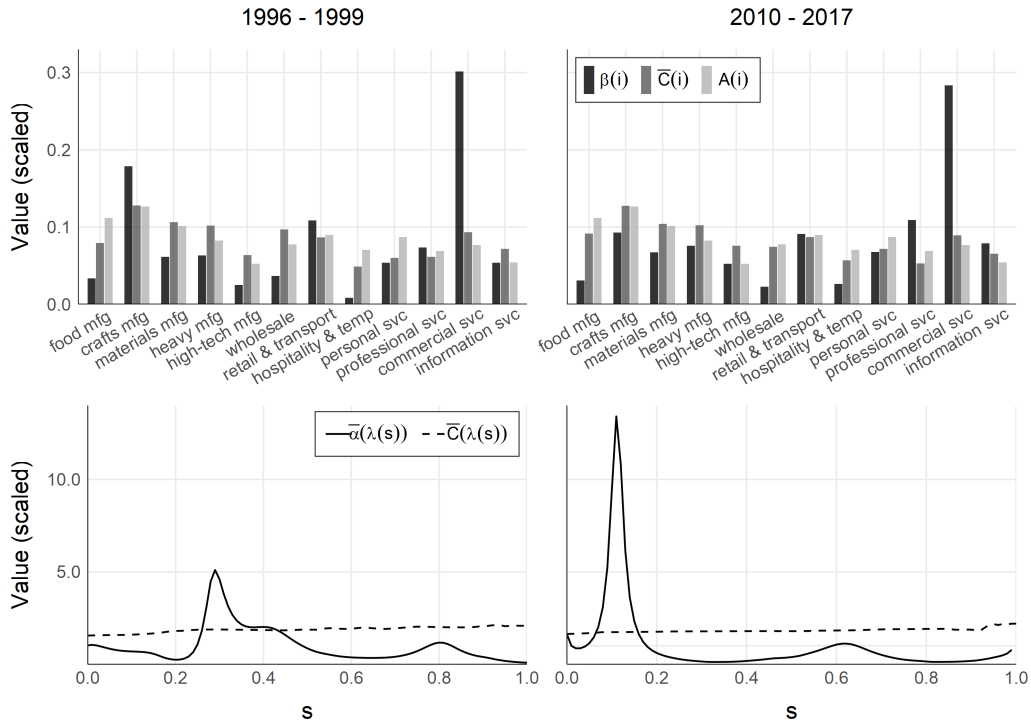


Figure B.6: Estimated Technical Parameters, East Germany

Table B.4: Aggregate Parameters, East Germany

Parameter	Definition	Value	Source
ψ	Amenity exponent	.538	Equal to $\frac{\eta}{1-\eta}$
ρ	Discount rate	.042	Discount factor of .96
δ	Separation rate	.183	$\frac{\text{Annual hires}}{\text{Total employment}}$
η	Match elasticity	.35	Kohlbrecher et al. (2016)
ζ	Match efficiency	1.29	$\frac{\text{Annual hires}}{\text{Predicted hires}}$




Figure B.7: Industry Average Firm Premia, 1993-1999



Figure B.8: Temp Agency Employment, West Germany




Figure B.9: Simulated Effect of Equal Pay for Temp Workers




Figure B.10: Simulated Effect of Entry Cost Compression

B.6 Alternative Specifications

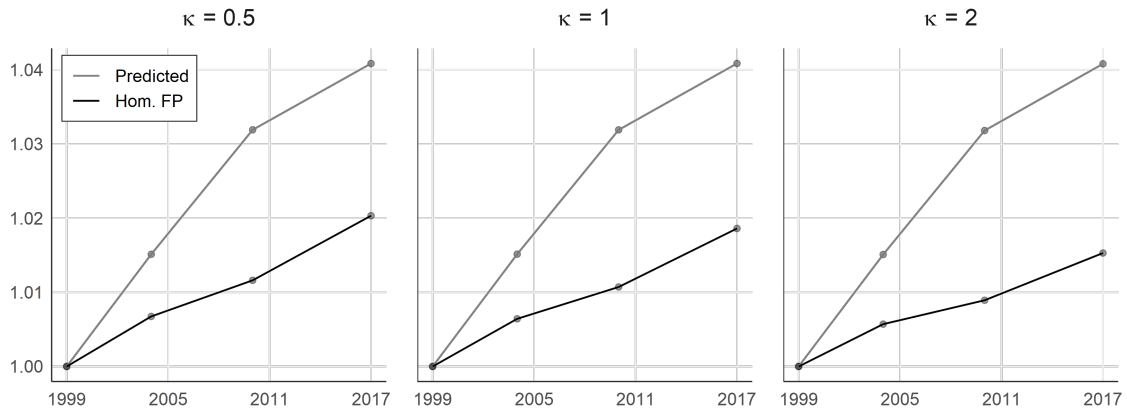


Figure B.11: Counterfactual Wage Trend, Constant Skill Distribution




Figure B.12: Counterfactual Wage Trend, Non-Unit Elasticity of Demand

Appendix C

Supplementary Materials For Chapter 3

C.1 Empirical Appendix

The first set of variables use in this analysis are those concerning job routineness, and consist of answers to the following questions:

1. **repeat tasks:** in all survey years, how often that “one and the same operation is repeated down to the last detail”
2. **follow instructions:** in all survey years, how often the execution of work is “prescribed down to the last detail”
3. **adapt to new tasks:** the 1979 survey asks respondents how often they must adapt to “new situations.” Beginning with the 1985/86 survey this changes to ”new tasks that you must first think about and become familiar with.”
4. **improve procedures:** in all survey years, how often respondents must “improve on previous procedures or try something new”
5. **solve problems:** for survey years 2006-2018, how often respondents must “react to and solve unforeseen problems”
6. **make decisions:** for survey years 2006-2018, how often respondents must “make difficult decisions independently and without guidance”

Response categories for 1979-1999 consist of five verbal frequencies: “practically always”, “often”, “every now and then”, “seldom”, and “practically never”. For 2006-2018, response categories for variables 1-4 are changed to four frequencies: “often”, “sometimes”, “seldom”,

and “never”. Variables 5-6, which are only present for 2006-2018, are coded in three frequencies: “often”, “sometimes”, and “never”. I code the responses “practically always” and “often” as 1, “every now and then” and “sometimes” as 2/3 for variables 1-4 and 1/2 for variables 5-6, “seldom” as 1/3, and “never” and “practically never” as 0. Some discrepancies are observed between the 1989-99 and 2005-06 surveys, but they are inconsistent across variables and do not suggest a systematic bias due to the change in response frequencies.

The other principal variable used in this analysis concerns personal computers. In the 1979 survey respondents are asked whether they often work with computers, EDV equipment, terminals, or screened devices on the job. In 1985/86 and 1991/92 this question is split across manufacturing and office roles and by device type, and so I assign a value of “yes” (1) if the answer is affirmative for any of these roles or devices. There is a major change in 1998/99, with workers simply being asked whether they “work with computers and data processing equipment” in their professional activity. There is again a slight change for the 2006-2018 surveys, with respondents asked simply how often they “work with computers”, the options being “often”, “sometimes”, and “never”; assign a value of “yes” to the first two responses.

With respect to other miscellaneous variables, occupation consists of 1988 KLDB 3-digit codes for the years 1979-2006, and 1992 KLDB 3-digit codes for 1999-2018. For consistency, the results in section two use 1988 KLDB codes for 1979-1992 results and 1992 KLDB codes for 1999-2018 results. Some occupations are not observed in all panels, and prior to performing difference-in-difference regressions I aggregate these codes into neighboring occupations. The same aggregated groupings are employed in the quantitative section. Occupational tasks from the 2006-2018 surveys are coded in the same frequencies as “solve problems” and “make decisions” above, and I assign them numerical values in identical fashion. The six aggregate groupings are composed as:

1. **analyze information:** “developing, researching, constructing [designing]”; “gathering information, investigating, documenting”
2. **advise others:** “organizing, planning, and preparing work processes [for others]”; “training, instructing, teaching, educating”; “providing advice and information”
3. **market goods:** “purchasing, procuring, selling”; “advertising, marketing, public re-

lations”

4. **manual labor**: “measuring, testing, quality control”; “repairing, refurbishing”; “transporting, storing, shipping”; “cleaning, removing waste, recycling”
5. **produce goods**: “manufacturing, producing goods and commodities”; “monitoring, control of machines, plants, processes”
6. **care for others**: “entertaining, accommodating, preparing food”, “nursing, caring, healing”, “protecting, guarding, patrolling, directing traffic”

Numerical values are averaged by task group, and divided by the sum across task groups in order to remove variation due to individuals who report performing all tasks more frequently.

Table C.1: Task Characteristics and Task Frequencies, 2006-2018

Sample	Analyze information	Advise others	Market goods	Manual labor	Produce goods
<i>Year fixed effects</i>					
Repetition	-.469 (.023)	-.179 (.024)	-.255 (.025)	.129 (.026)	.003 (.026)
Instructions	-.267 (.024)	-.139 (.025)	-.370 (.026)	-.002 (.003)	.155 (.027)
Adaptation	.522 (.020)	.205 (.020)	.153 (.020)	-.119 (.018)	.176 (.019)
Improvement	.281 (.022)	.039 (.023)	.085 (.023)	-.286 (.022)	.053 (.023)
Solve problems	.110 (.023)	.041 (.024)	-.048 (.023)	-.344 (.020)	-.106 (.021)
Make decisions	-.013 (.026)	.040 (.028)	.048 (.028)	-.509 (.026)	-.203 (.026)
<i>Industry, occupation, and year fixed effects</i>					
Repetition	-.269 (.030)	-.164 (.029)	-.282 (.031)	.057 (.032)	.044 (.034)
Instructions	-.179 (.031)	-.171 (.031)	-.306 (.032)	-.089 (.032)	-.058 (.035)
Adaptation	.263 (.023)	.127 (.022)	.135 (.023)	-.083 (.022)	.043 (.024)
Improvement	.208 (.025)	.074 (.025)	.135 (.026)	-.167 (.025)	.029 (.027)
Solve problems	.083 (.025)	.110 (.024)	.098 (.025)	-.163 (.023)	-.017 (.025)
Make decisions	.010 (.030)	.131 (.030)	.187 (.032)	-.332 (.030)	-.128 (.032)

NOTE. Marginal effects and robust standard errors from fractional logit regressions of task characteristics on PC use, aggregated by 3-digit occupation. All regressions include dummies for year, 3-digit occupation, 1-digit industry. Bold results indicate 95% significance.

Table C.2: Task Frequencies and PC Use, 2006-2018

Sample	Analyze information	Advise others	Market goods	Manual labor	Produce goods	Care for others
Full sample	.068 (.004)	.031 (.003)	.035 (.003)	-.053 (.003)	-.011 (.002)	-.014 (.002)
Education						
None	.048 (.009)	.022 (.008)	.030 (.007)	-.065 (.010)	-.007 (.008)	-.010 (.006)
Vocational	.057 (.004)	.027 (.004)	.023 (.003)	-.052 (.003)	-.009 (.003)	-.012 (.002)
University	.072 (.015)	.012 (.15)	.068 (.010)	-.044 (.006)	-.015 (.005)	-.016 (.004)
Wage pct.						
1-25	.058 (.006)	.017 (.005)	.020 (.005)	-.048 (.006)	.000 (.005)	-.021 (.004)
26-50	.050 (.006)	.029 (.006)	.023 (.005)	-.041 (.006)	-.013 (.005)	-.010 (.003)
51-75	.050 (.009)	.017 (.008)	.059 (.007)	-.045 (.006)	-.017 (.007)	-.002 (.005)
76-100	.044 (.019)	.042 (.019)	.097 (.016)	-.040 (.007)	-.027 (.008)	.001 (.005)

NOTE. Marginal effects and robust standard errors from fractional logit regressions of task characteristics on PC use, aggregated by 3-digit occupation. All regressions include dummies for year, 3-digit occupation, 1-digit industry. Bold results indicate 95% significance.

Table C.3: Occupation Mean Task Frequencies and PC Use (D-in-D)

Years	Analyze information	Advise others	Market goods	Manual labor	Produce goods	Care for others
2006-2012	.027 (.013)	.033 (.011)	.054 (.013)	-.049 (.008)	-.017 (.008)	.006 (.009)
2012-2018	.060 (.012)	.041 (.014)	.015 (.013)	-.042 (.011)	-.012 (.010)	-.011 (.008)

NOTE. Marginal effects and robust standard errors from fractional logit regressions of task characteristics on PC use, aggregated by 3-digit occupation. All regressions include occupation and year dummies. Bold results indicate 95% significance.

C.2 Assignment and Wages

The environment is similar to that studied by Costinot and Vogel (CV, 2010), with the difference that unit output costs for automated firms are $\frac{w(s)+r\bar{K}}{y(j,s,\bar{K})}$ and not $\frac{w(s)}{y(j,s,0)}$. Following CV, I define $\omega(s)$ to be the set of jobs j for which at least one producer hires an s -worker, and $\sigma(j)$ the worker types assigned to j . I begin with a lemma establishing that non-automated jobs are always associated with lower-skill workers than automated job:

Lemma 4. *If $s' \geq s$ and $K(j) > 0$ for some firm employing s , then $K(j') > 0$ for any firm j' employing s' .*

Proof. This result is shown by establishing that $w(s)$ and $\gamma_u(s)$ satisfy the single-crossing: there exists at most one s such that $w(s) = r\gamma_u(s)$, and that for any $s' > s$ we will have $w(s') >$

$r\gamma_u(s')$. Suppose the contrary: there exists an $s' > s$ such that $w(s') < r\gamma_u(s')$ but $w(s) > r\gamma_u(s)$. From free entry we have $p(j)y^*(j, s, K) \geq w(s)$ for any j employing worker s , and by assumption we have $y^*(j, s, K)$ increasing more quickly in s than $\gamma_u(s)$. But then $p(j)y^*(j, s', K) - w(s') > p(j)y^*(j, s, K) - w(s)$, violating the producer's profit-maximizing problem. The result then follows from (22), which states that $K > 0$ when $w > r\gamma_u$ and $K = 0$ otherwise. \square

With this result it is possible to derive the main result:

Lemma 5. *There exists a continuous and strictly increasing function $\lambda : [\underline{s}, \bar{s}] \rightarrow [0, 1]$ such that $L(j, s) > 0$ if and only if $\lambda(s) = j$, and where $\lambda(\underline{s}) = 0$ and $\lambda(\bar{s}) = 1$.*

Proof. In the proof below, I provide an abbreviated description whenever the steps follow closely those described by CV.

First, $\omega(s)$ and $\sigma(j)$ are non-empty. That $\omega(s)$ is non-empty follows from market-clearing, as labor is supplied inelastically and worker output is strictly positive. If there existed an s that was not assigned to any producer, it must be that $w(s)$ is sufficiently large that no firms wish to employ s -workers; but then market-clearing would imply that $w(s) \rightarrow 0$, a contradiction. Non-empty $\sigma(j)$ follows from the profit-maximizing condition for intermediate producers:

$$p(j) - \frac{w(s) + rK^*(j, s)}{y^*(j, s, K)} \leq 0$$

with equality for $s \in \sigma(j)$. If $\sigma(j)$ is empty for some j , then from $\omega(s)$ non-empty there must be a j' with $\sigma(j')$ non-empty and $p(j')/p(j) = \frac{A}{p(j)} = 0$ for any finite A . But then from the previous condition it must be that

$$\frac{p(j')}{p(j)} + r \frac{K^*(j', s) - K^*(j, s)}{p(j)y^*(j', s)} \geq \frac{y^*(j, s, K)}{y^*(j', s, K)} > 0$$

a contradiction since we must have the left-hand side equal to 0.

Second, $\sigma(j)$ is a non-empty interval of $[\underline{s}, \bar{s}]$, and if $j' > j$, then $s' > s$ for any $s' \in \sigma(j')$ and $s \in \sigma(j)$. Suppose instead that for some $j^+ > j^-$ and $s^+ > s^-$, we have $s^- \in \sigma(j^+)$ and $s^+ \in \sigma(j^-)$. Then from intermediate producers' first-order condition we must have

$$0 \geq p(j^-) - \frac{w(s^-) + rK^*(j^-, s^-)}{y^*(j^-, s^-)}$$

$$\begin{aligned}
&= \frac{w(s^+) + rK^*(j^+, s^+)}{y^*(j^-, s^+)} \left(1 - \frac{\frac{w(s^-) + rK^*(j^-, s^-)}{y^*(j^-, s^-)}}{\frac{w(s^+) + rK^*(j^-, s^+)}{y^*(j^-, s^+)}} \right) \\
&> \frac{w(s^+) + rK^*(j^+, s^+)}{y^*(j^-, s^+)} \left(1 - \frac{\frac{w(s^-) + rK^*(j^-, s^-)}{y^*(j^+, s^-)}}{\frac{w(s^+) + rK^*(j^-, s^+)}{y^*(j^+, s^+)}} \right) \\
&= \frac{w(s^+) + rK^*(j^+, s^+)}{y^*(j^-, s^+)} \left(1 - \frac{p(j^+)}{\frac{w(s^+) + rK^*(j^-, s^+)}{y^*(j^+, s^+)}} \right) \\
&= \frac{\frac{w(s^+) + rK^*(j^+, s^+)}{y^*(j^-, s^+)}}{\frac{w(s^+) + rK^*(j^-, s^+)}{y^*(j^+, s^+)}} \left(\frac{w(s^+) + rK^*(j^-, s^+)}{y^*(j^+, s^+)} - p(j^+) \right) \\
&\geq 0
\end{aligned}$$

a contradiction, with the strict inequality following from assumption 1 and lemma 1. The inequality holds trivially when either $K > 0$ for both s^- and s^+ or when $K = 0$ for both types. From lemma 1 the only other possibility is that $K > 0$ for s^+ but $K = 0$ for s^- , in which case $\frac{w(s^+) + rK^*(j^-, s^+)}{y^*(j^+, s^+)} = w(s^+)y(s^+, j^-, 0) - \kappa[1 - \alpha(j^-)]\gamma_u(s^+)[w(s^+) - r\gamma_u(s^+)]$. From $\alpha'(j) > 0$ we can then see that the inequality holds with even greater force than under the other two cases.

Fourth, ω and σ are single-valued almost everywhere. The proof is unchanged from CV and so I provide only the intuition: if ω (or σ) has positive measure over a domain with positive measure, then from the previous result the range of the correspondence will have measure greater than the measure of $[\underline{s}, \bar{s}]$, a contradiction.

Fifth, $\sigma(j)$ is single-valued. If this is not the case, then from step 3 there exists a non-degenerate interval $[s, s']$ for which all workers are assigned to job j . Step 4 implies that there exists another job j' that is assigned to a single worker type. But then $p(j)/p(j') = 0$, contradicting the free entry condition that $p(j)y(j, s'') \geq p(j')y(j', s'')$ for $s'' \in [s, s']$, a contradiction given that $y > 0$.

From the last step we have $\sigma(j)$ single-valued; from the third step, weakly increasing; from the first step, continuous and such that $\sigma(0) = \underline{s}$ and $\sigma(1) = \bar{s}$; and from the fourth step, σ is strictly increasing. Hence we have a continuous, strictly increasing bijection $\lambda(s) = \omega(s) = \{j \mid L(j, s) = 1\} = \sigma^{-1}(s)$. \square

Lemma 6. *There exists a single threshold skill level s^* , which may be equal to \underline{s} or \bar{s}], for which producers automate when $s > s^*$ and do not automate when $s > s^*$. The wage*

functions satisfies the differential equation

$$\frac{w'(s)}{w(s)} = \begin{cases} \frac{d}{ds} \log y(\lambda(s), s, 0) & s < s^* \\ \frac{d}{ds} \log y(\lambda(s), s, \bar{K}) & s > s^* \end{cases}$$

where $w(s)$ is continuous but not differentiable at s^* .

Proof. Market-clearing, free entry, and continuity of $y(j, s)$ imply that $w(s)$ is continuous, while from lemma 1 we have that there exists at most one s^* satisfying $w(s^*) = r\gamma_u(s^*)$. Let $\mathbb{I}[s > s^*]$ be an indicator function taking the value 1 when $s > s^*$, and 0 otherwise.

From lemma 2, for any producer employing s -labor we must have

$$p(\lambda(s)) - \frac{w(s)}{y(\lambda(s), s, K)} + \mathbb{I}[s > s^*]\kappa[1 - \alpha(\lambda(s))]r \leq 0$$

For any $s \neq s^*$, there will exist a neighborhood around s such that $\mathbb{I}[s' > s^*]$ takes the same value for any $s' \in [s - ds, s + ds]$. In the case where $s > s^*$ the following inequalities must hold:

$$\begin{aligned} & \left[p(\lambda(s)) - \kappa[1 - \alpha(\lambda(s))]r \right] - \frac{w(s)}{y(\lambda(s), s, \bar{K})} \\ & \geq \left[p(\lambda(s)) - \kappa[1 - \alpha(\lambda(s))]r \right] - \frac{w(s + ds)}{y(\lambda(s), s + ds, \bar{K})} \\ & \left[p(\lambda(s + ds)) - \kappa[1 - \alpha(\lambda(s + ds))]r \right] - \frac{w(s + ds)}{y(\lambda(s + ds), s + ds, \bar{K})} \\ & \geq \left[p(\lambda(s + ds)) - \kappa[1 - \alpha(\lambda(s + ds))]r \right] - w(s) \frac{w(s)}{y(\lambda(s + ds), s, \bar{K})} \end{aligned}$$

For $s < s^*$ we must have the following inequalities hold:

$$\begin{aligned} & \left[p(\lambda(s)) - \kappa[1 - \alpha(\lambda(s))]r \right] - \frac{w(s)}{y(\lambda(s), s, 0)} \geq p(\lambda(s)) - \frac{w(s + ds)}{y(\lambda(s), s + ds, 0)} \\ & p(\lambda(s + ds)) - \frac{w(s + ds)}{y(\lambda(s + ds), s + ds, 0)} \geq p(\lambda(s + ds)) - w(s) \frac{w(s)}{y(\lambda(s + ds), s, 0)} \end{aligned}$$

It follows that in the first case we will have

$$\begin{aligned} & \left[p(\lambda(s)) - \kappa[1 - \alpha(\lambda(s))]r \right] \left[y(\lambda(s), s + ds, \bar{K}) - y(\lambda(s), s, \bar{K}) \right] \\ & \leq w(s + ds) - w(s) \end{aligned}$$

$$\leq \left[p(\lambda(s+ds)) - \kappa[1 - \alpha(\lambda(s+ds))]r \right] \left[y(\lambda(s+ds), s+ds, \bar{K}) - y(\lambda(s+ds), s, \bar{K}) \right]$$

and in the second

$$\begin{aligned} & p(\lambda(s)) \left[y(\lambda(s), s+ds, 0) - y(\lambda(s), s, 0) \right] \\ & \leq w(s+ds) - w(s) \\ & \leq p(\lambda(s+ds)) \left[y(\lambda(s+ds), s+ds, 0) - y(\lambda(s+ds), s, 0) \right] \end{aligned}$$

We must have $p(j)$ continuous and therefore we can divide by ds and take the limit of the previous inequalities to show that

$$w'(s) = \begin{cases} p(\lambda(s))y(\lambda(s+ds), s, 0) & s < s^* \\ \left[p(\lambda(s)) - \kappa[1 - \alpha(\lambda(s))] \right] y(\lambda(s+ds), s, \bar{K}) & s > s^* \end{cases}$$

Substitution for $p(\lambda(s))$ and $p(\lambda(s)) - \kappa[1 - \alpha(\lambda(s))]$ and division by $w(s)$ yields the final result. \square

Lemma 7. *The matching function satisfies*

$$\lambda'(s) = \begin{cases} \frac{y(\lambda(s), s', 0)^{1-\rho} F'(s)}{\beta(\lambda(s))^\rho Y} w(s)^\rho & s < s^* \\ \frac{y(\lambda(s), s', \bar{K}) F'(s)}{\beta(\lambda(s))^\rho Y} \left(w(s) + \kappa[1 - \alpha(\lambda(s))] r y(\lambda(s), s', \bar{K}) \right)^\rho & s > s^* \end{cases}$$

where $\lambda(\underline{s}) = 0$ and $\lambda(\bar{s}) = 1$.

Proof. This portion of the proof follows closely Costinot and Vogel (2010). Total supply of j -labor is given by

$$L(j, s) = F'(s) \delta [j - \lambda(s)]$$

From market-clearing we have

$$Y(\lambda(s)) = \int_{s < s^*} y(\lambda(s), s', 0) L(\lambda(s), s') ds' + \int_{s > s^*} y(\lambda(s), s', \bar{K}) L(\lambda(s), s') ds'$$

and following CV we may derive the differential equation for the matching function, which is defined

piece-wise:

$$\lambda'(s) = \begin{cases} \frac{y(\lambda(s), s', 0) F'(s)}{Y(\lambda(s))} & s < s^* \\ \frac{y(\lambda(s), s', \bar{K}) F'(s)}{Y(\lambda(s))} & \mathbb{I}^s = 1 \end{cases}$$

From final good profit-maximization and free entry we have

$$Y(\lambda(s)) = \begin{cases} \beta(\lambda(s))^\rho Y \left[\frac{y(\lambda(s), s', 0)}{w(s)} \right]^\rho & s < s^* \\ \beta(\lambda(s))^\rho Y \left[\frac{y(\lambda(s), s', \bar{K})}{w(s) + \kappa [1 - \alpha(\lambda(s))] r y(\lambda(s), s', \bar{K})} \right]^\rho & s > s^* \end{cases}$$

Substitution of $Y(\lambda(s))$ and $y(\lambda(s), s)$ in the previous equation then yields the result. \square

C.3 Proofs of Theoretical Results

Theorem 1: Short-Run Polarization

Proof. To establish the first part of the result I show that $\lambda(s^*)$ is strictly smaller under automation. Suppose otherwise. It can then be shown that for all s , $\lambda(s')$ is greater under automation and hence from (24) $w'(s)/w(s)$ is everywhere larger. But then $\frac{w(s)^\rho}{Y} = \frac{w(s)^\rho}{\int [F'(s)w(s) + rK^*(\lambda(s), s)] ds}$ must be strictly smaller for $s < s^*$, implying that these workers are assigned to a smaller set of jobs and hence that $\lambda(s^*)$ is smaller, a contradiction.

For the second part of the proof, for any $j' > j$ and holding $(y(j', \lambda^{-1}(j'), 0)/y(j, \lambda^{-1}(j), 0))^{\rho-1}$ we will have $L(j')/L(j)$ greater under automation both because $\alpha'(j) > 0$ and because $\lambda(s)$ and $w(s) - r\gamma_h(s)$ are both increasing in s . Now if $\rho < 1$ and employment decreases for all automated jobs then we must have $(y(1, \bar{s}, 0)/y(\lambda(s), s, 0))^{\rho-1}$ smaller under automation for any $s < \bar{s}$ and therefore $L(j')/L(j)$ greater. On the other hand if employment increases for some $j'' > \lambda^{-1}(s^*)$, then it must increase for all $j > j''$ and the result is shown. In both cases the result is shown for $\rho < 1$. Finally, from continuity the result must also hold for $\rho > 1$ but sufficiently small.

Finally, from above we will have $w'(s)/w(s)$ smaller for $s < s^*$, whereas $\lambda(\bar{s}) = 1$ and (24) imply that $w'(s)/w(s)$ will be greater around \bar{s} . Assumption 1 and continuity of all arguments is then sufficient to establish that either $w'(s)/w(s)$ is strictly greater for $s > s^*$ and smaller for $s < s^*$, in which case I define $s' = s^*$; or that there exists at least one value of s at which $w'(s)/w(s)$ is unchanged under automation, in which case I define s' equal to the largest such value and the proof

is completed. □

Theorem 2: Long-Run Effects of Automation, $\rho = 0$

Proof. Both results may be shown by establishing that $\lambda(s)$ is strictly greater under automation for all s , which would in turn imply both that $w'(s)/w(s)$ is strictly greater and that $\int_0^j L(j, \lambda^{-1}(j))$ is strictly smaller. Now if t^* is increasing in s then holding λ fixed, $\frac{d}{ds} \log y^*(\lambda(s), s, K)$ will be smaller under automation and therefore $\frac{d}{ds} \log Y/y^*(\lambda(s), s, K)$ will be greater; and in particular, we must have $Y/y^*(\lambda(\underline{s}), \underline{s}, K)$ smaller and $Y/y^*(\lambda(\bar{s}), \bar{s}, K)$ bigger. It follows that $\lambda'(\underline{s})$ is greater under automation and $\lambda'(\bar{s})$ smaller, while labor market-clearing implies that we cannot have $\lambda(s)$ the same under automation for any s in (\underline{s}, \bar{s}) . Continuity of the matching function then establishes that $\lambda(s)$ is everywhere and the results follow. □

Theorem 2: Long-Run Effects of Automation, $\rho = 1$

Proof. I begin by noting that in the case where $\rho = 1$ we will have labor demand equal to

$$L(\lambda(s), s) = \frac{\beta(\lambda(s))Y}{(w(s) + r\kappa[1 - \alpha(\lambda(s))]y^*(\lambda(s), s, K))}$$

Making use of equation (24) it can be shown that

$$\begin{aligned} \frac{d}{ds} \log (w(s) + r\kappa[1 - \alpha(\lambda(s))]y^*(\lambda(s), s, K)) &= \frac{y_s^*(\lambda(s), s, K)}{y^*(\lambda(s), s, K)} \\ &+ \lambda'(s)r\kappa y^*(\lambda(s), s, K) \frac{[1 - \alpha(\lambda(s))] \frac{y_j^*(\lambda(s), s, K)}{y^*(\lambda(s), s, K)} - \alpha'(\lambda(s))}{w(s) + r\kappa[1 - \alpha(\lambda(s))]y^*(\lambda(s), s, K)} \end{aligned} \quad (26)$$

where the last term is strictly negative.

Suppose now that $w(\bar{s})/w(\underline{s})$ is smaller under automation. We must have $w'(\underline{s})/w(\underline{s})$ and $w'(\bar{s})/w(\bar{s})$ greater, which implies that there exists at least one s for which $w(s)/w(\underline{s})$ is unchanged under automation. For each of the smallest and largest such values of s , it must be that $w'(s)/w(s)$ is smaller, and hence that $\lambda(s)$ is strictly lower. But then between the lower of the two s values and \underline{s} we must have average wage higher under automation, while between the higher value of s and \bar{s} we must have average wage higher. Higher wages in the lower region imply a larger value of $\lambda'(s)$, and lower values of wages in the upper region imply a smaller value of $\lambda'(s)$, implying that at the points where $w(s)$ is unchanged we must have $\lambda(s)$ larger, a contradiction.

Now suppose that, for any $r \in (0, \bar{r})$, between any two points s' and s'' we have $\lambda(s)$ weakly greater relative to the case where $r = \bar{r}$. Then $w'(s)/w(s)$ must be strictly greater in this interval, and hence from (25) it must be that almost everywhere in $[s', s'']$ we have $\lambda(s)$ strictly greater. Suppose without loss of generality that $\lambda(s)$ is strictly greater in this interval. Then it must be that $\lambda'(s')$ is greater than before while $\lambda'(s'')$ is smaller. But since $w'(s)/w(s)$ is strictly larger between s' and s'' , whereas rental costs decrease proportionally in r . it must be that the sum of wages and rental costs has increased for s'' relative to s' , which from (25) implies that $\lambda'(s')$ has decreased relative to $\lambda'(s'')$, a contradiction. Hence over any interval we must have $\lambda(s)$ smaller almost everywhere, and the result follows.

Now when $r = 0$, the negative term in (26) drops out and so from the previous result we must have $\lambda'(\underline{s})$ smaller relative to $\lambda'(\bar{s})$. But then either the first of these terms is strictly smaller, or the second strictly larger, implying in either case that for some subset of $[\underline{s}, \bar{s}]$ we will have $\lambda(s)$ strictly smaller under automation, and the result follows. \square

C.4 Continuous Model

Free entry implies that

$$p(j, s) = \frac{w(s) + G(\epsilon^*(s))r\kappa[1 - \alpha(j)]y^*(j, s, \bar{K}) \int_0^{\epsilon^*(s)} G'(k)k dk}{y^*(j, s, 0) + G(\epsilon^*(s))\left[y^*(j, s, \bar{K}) - y^*(j, s, 0)\right]}.$$

while labor demand (i.e. firm entry) will be

$$\begin{aligned} L(j, s', s) &= \left(\frac{\beta(j)}{p(j, s', s)}\right)^\rho \frac{Y}{y^*(j, s, 0) + G(\epsilon^*(s))\left[y^*(j, s, \bar{K}) - y^*(j, s, 0)\right]} \\ &= \beta(j)^\rho Y \frac{\left(y^*(j, s, 0) + G(\epsilon^*(s))\left[y^*(j, s, \bar{K}) - y^*(j, s, 0)\right]\right)^{\rho-1}}{\left(w(s) + G(\epsilon^*(s))r\kappa[1 - \alpha(j)]y^*(j, s, \bar{K}) \int_0^{\epsilon^*(s)} G'(k)k dk\right)^\rho} \end{aligned}$$

Optimal choice of skill, on the other hand, implies that

$$\begin{aligned} w'(s) &= \left[G(\epsilon^*(s))p(j) - r\kappa[1 - \alpha(j)] \int_0^{\epsilon^*(s)} G'(k)k dk \right] \frac{\partial}{\partial s} y^*(j, s, \bar{K}) + [1 - G(\epsilon^*(s))]p(j) \frac{\partial}{\partial s} y^*(j, s, 0) \\ &\quad + \frac{d}{ds} \epsilon^*(s) \frac{\partial}{\partial \epsilon^*} \pi(j, s) \end{aligned}$$

where profit maximization implies that the last term $\frac{\partial}{\partial \epsilon^*} \pi(j, s)$ is equal to zero. Hence we must have

$$\frac{w'(s)}{w(s)} = \frac{A(j, s) \frac{\partial}{\partial s} y^*(j, s, \bar{K}) + B(j, s) \frac{\partial}{\partial s} y^*(j, s, 0)}{A(j, s) A(j, s) y^*(j, s, \bar{K}) + B(j, s) y^*(j, s, 0) + B(j, s) A(j, s) y^*(j, s, 0) + B(j, s) y^*(j, s, 0)}$$

$$A(j, s) = G(\epsilon^*(s)) p(j) - r\kappa [1 - \alpha(j)] \int_0^{\epsilon^*(s)} G'(k) k dk$$

$$B(j, s) = [1 - G(\epsilon^*(s))] p(j) ,$$

where since match output is log-supermodular and α is increasing, $\frac{A(j, s) \frac{\partial}{\partial s} y^*(j, s, \bar{K}) + B(j, s) \frac{\partial}{\partial s} y^*(j, s, 0)}{A(j, s) y^*(j, s, \bar{K}) + B(j, s) y^*(j, s, 0)}$ must be increasing in j and therefore $A(j, s) y^*(j, s, \bar{K}) + B(j, s) y^*(j, s, 0)$ will also be log-supermodular. Optimal assignment in the continuous model will be characterized by the differential equations

$$\frac{w'(s)}{w(s)} = \frac{A(j, s) \frac{\partial}{\partial s} y^*(j, s, \bar{K}) + B(j, s) \frac{\partial}{\partial s} y^*(j, s, 0)}{A(j, s) y^*(j, s, \bar{K}) + B(j, s) y^*(j, s, 0)} \quad (27)$$

$$\lambda'_m(s) = \frac{\left(y^*(\lambda(s), s, 0) + G(\epsilon^*(s)) \left[y^*(\lambda(s), s, \bar{K}) - y^*(\lambda(s), s, 0) \right] \right)^{1-\rho}}{\beta(\lambda(s))^\rho Y}$$

$$\times \left(w(s) + G(\epsilon^*(s)) r\kappa [1 - \alpha(\lambda(s))] y^*(\lambda(s), s, \bar{K}) \int_0^{\epsilon^*(s)} G'(k) k dk \right)^\rho \quad (28)$$

where $\lambda(\underline{s}) = 0$ and $\lambda(\bar{s}) = 1$ as before, $\epsilon^*(s) = \frac{w(s)}{r\gamma_l(s)}$, and where

$$A(j, s) = G(\epsilon^*(s)) p(j) - r\kappa [1 - \alpha(j)] \int_0^{\epsilon^*(s)} G'(k) k dk$$

$$B(j, s) = [1 - G(\epsilon^*(s))] p(j) .$$

In contrast to static model, (27) and (28) will be everywhere continuous.

C.5 Additional Qualitative and Quantitative Results

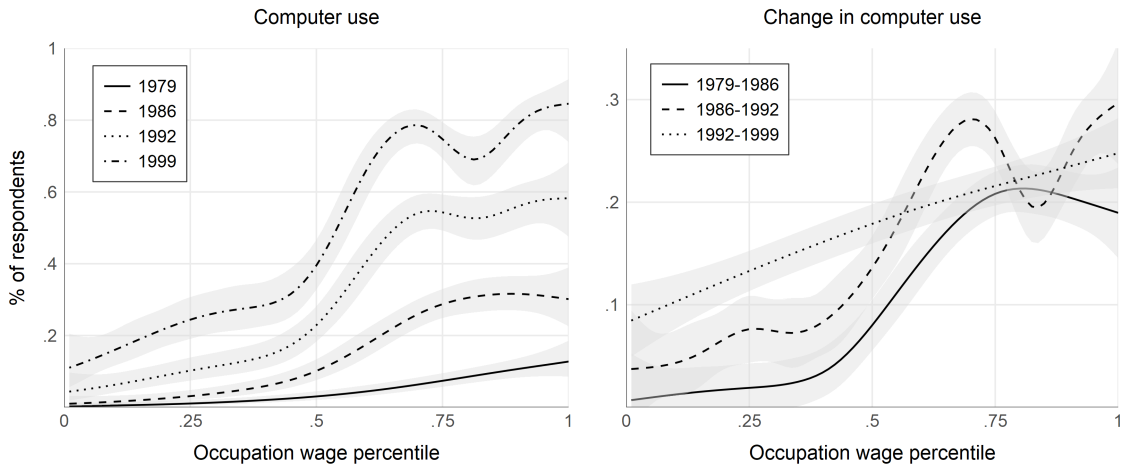


Figure C.1: Computer Use By Occupation Wage Percentile, 1979-1999

NOTE. Log wage and PC use averaged by 1988 KLDB occupation. Percentiles are time-invariant and reflect 1979 wages. Shaded regions indicate 95% confidence intervals.

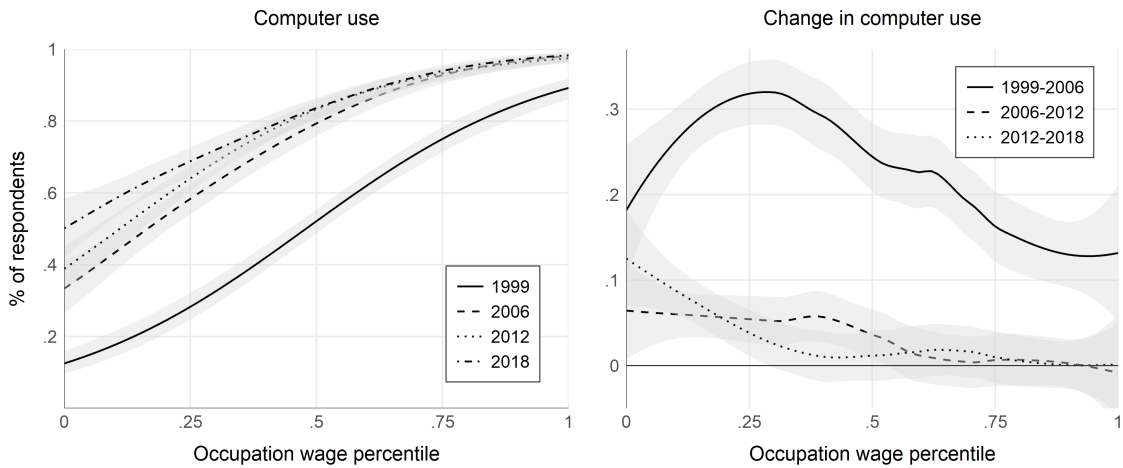


Figure C.2: Computer Use by Occupation Wage Percentile, 1999-2018

NOTE. Log wage and PC use averaged by 1992 KLDB occupation. Percentiles are time-invariant and reflect 1999 wages. Shaded regions indicate 95% confidence intervals.

Table C.4: Regression: Change in Log Occupation Employment Share, High-Skill

Independent variable	<i>Regression coefficients</i>							
	Years 1979-1999				Years 1999-2018			
PC use	.338 (.108)	.284 (.108)	.644 (.268)	.659 (.270)	.739 (.115)	.707 (.120)	.565 (.449)	.585 (.454)
Δ PC use	-.619 (.215)	-.608 (.207)	-.690 (.252)	-.719 (.257)	-.488 (.181)	-.480 (.184)	-.305 (.236)	-.291 (.233)
Log(wage)		..268 (.128)		.380 (.254)		.054 (.123)		-.174 (.260)
Occup. FE			X	X			X	X
Observations	409	409	409	409	482	482	482	482

NOTE. Difference-in-difference regression with occupational employment share as the dependent variable, occupation wage percentiles between .5 and 1. Employment shares calculated from raw survey counts. All regressions include year fixed effects.

Table C.5: Regression: Change in Log Occupation Employment Share, Low-Skill

Independent variable	<i>Regression coefficients</i>							
	Years 1979-1999				Years 1999-2018			
PC use	.102 (.165)	.086 (.166)	.557 (.266)	.585 (.262)	.209 (.122)	.067 (.141)	-1.277 (.426)	-1.402 (.438)
Δ PC use	-.042 (.296)	-.044 (.296)	-.571 (.343)	-.575 (.344)	.043 (.179)	.066 (.174)	.696 (.275)	.645 (.279)
Log(wage)		.039 (.142)		-.193 (.378)		.332 (.137)		.790 (.313)
Occup. FE			X	X			X	X
Observations	418	418	418	418	386	386	386	386

NOTE. Difference-in-difference regression with occupational employment share as the dependent variable, occupation wage percentiles between 0 and .5. Employment shares calculated from raw survey counts. All regressions include year fixed effects.

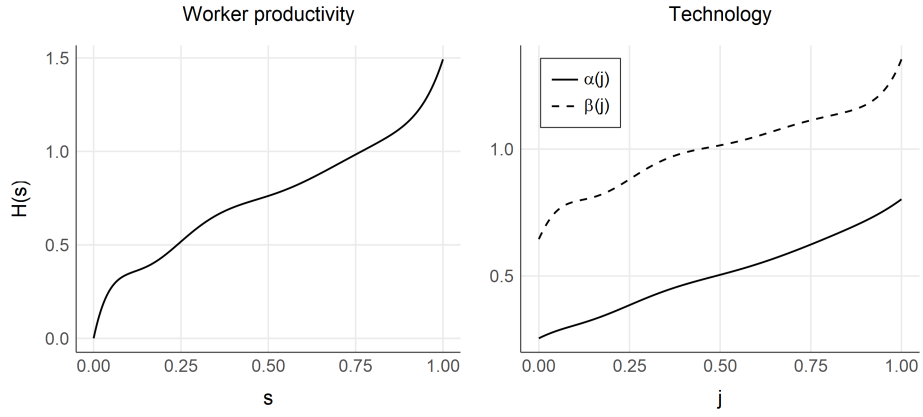


Figure C.3: Distributional Parameter Estimates

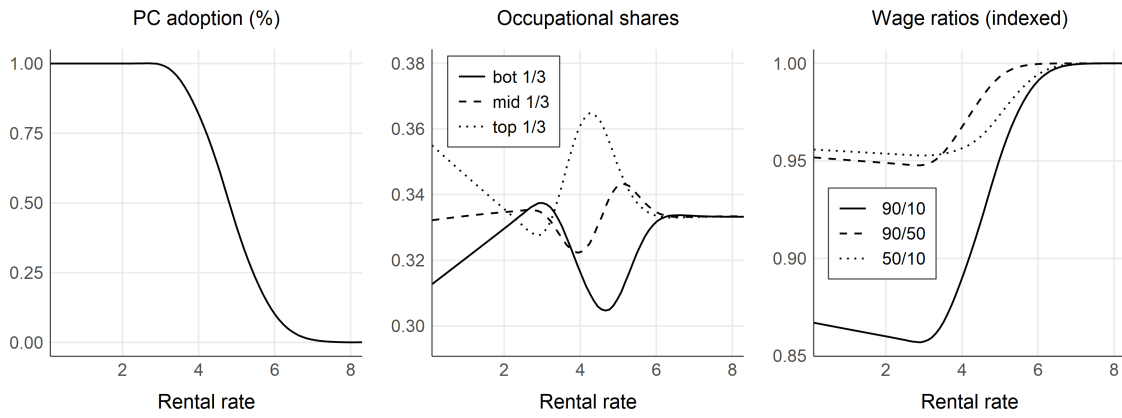


Figure C.4: Automation of Skilled Task

NOTE. Smoothed model output over a grid of rental rates r_s . Model parameters are unchanged from above and reflect the limiting case where the skill-biased technology is cost-less ($r = 0$).