Assortative Matching in Labor and Marriage Markets

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae
an der Ludwig-Maximilians-
Universität München

vorgelegt von
Bastian Schulz

Assortative Matching in Labor and Marriage Markets

Inaugural-Dissertation

zur Erlangung des Grades

Doctor oeconomiae publicae (Dr. oec. publ.)

an der Volkswirtschaftlichen Fakultät

der Ludwig-Maximilians-Universität München

2017

vorgelegt von

Bastian Schulz

Referent: Prof. Helmut Rainer, Ph.D.

Korreferent: Prof. Dr. Christian Holzner

Promotionsabschlussberatung: 31. Januar 2018

Tag der mündlichen Prüfung: 24. Januar 2018

Berichterstatter: Helmut Rainer, Christian Holzner, Philipp Kircher

Für Eva.

Acknowledgments

This doctoral thesis is the result of five years of work at the ifo Institute in Munich. As a doctoral student, I have benefited from a truly excellent research environment. I was given the opportunity to focus on the most exciting and challenging questions I could come up with, without distractions. First and foremost, I am deeply indebted to Helmut Rainer, my primary advisor, for providing this environment. Helmut taught me what good economics research is about. He was always very generous with encouragement, career advice, and feedback on my research.

I am equally indebted to Christian Holzner, my second advisor. The discussions based on our joint passion for search and matching models have influenced my thinking in countless ways and improved my work immensely. Our interactions have been a constant source of inspiration to me. Christian's support and encouragement have been a major factor in my decision to pursue an academic career.

Most of my research is inspired by the work of two great economists, Philipp Kircher and Robert Shimer. I was extremely lucky to get the chance to personally interact with both of them during my studies.

I would like to thank Philipp for being so encouraging and open to my ideas from the very beginning. Importantly, he gave me the opportunity to interact with the leading researchers in my field at his workshops. I am very grateful and excited that Philipp will join my Ph.D. committee.

To Robert I owe very special thanks. The time I spent working with him at the University of Chicago in 2016 was the highlight of my doctoral studies. Besides the enormous progress I have made on my projects under his guidance, working in the stimulating intellectual environment of the Chicago Department of Economics was a fantastic experience.

Christian Merkl has been an important teacher and mentor to me as well. Being a doctoral student himself, he led my first tutorial in macroeconomics at the University of Kiel. Later, his brilliant lectures and enthusiasm inflamed my interest in macro-labor topics. I would like to thank Christian for the constant support throughout my studies.

A number of colleagues made my time in Munich very special and memorable. The feeling of working among friends is invaluable. I would like to thank Wolfgang Auer, Natalia Danzer, Anita Fichtl, Timo Hener, Till Nikolka, Patrick Reich, Fabian Siuda, Manuel Wiegand, and especially my co-author Ben Lochner and my officemate Judith Saurer for all the interesting conversations, on economics and everything else, and loads of fun.

Finally, I would like to thank my family. My parents, Beate and Michael, have prepared me well to pursue this academic journey. Most importantly, my wife, Natalia, deserves the largest amount of gratitude. Amazingly, she challenges me to give my very best but, at the same time, keeps me grounded and helps me unwind. I would have never made it to this point without her.

Munich, September 2017

Bastian Schulz

Table of Contents

List of Figures

List of Tables

"Unlike Marxian analysis, the economic approach [...] does not assume that individuals are motivated solely by selfishness or gain. It is a method of analysis, not an assumption about particular motivations."

Gary S. Becker, Nobel Lecture
December 9, 1992

General Introduction

Whom to marry and where to work are probably the most far-reaching decisions humans make during their lifetime. Besides manifold implications on the individual level, these decisions also shape economies and societies at the macro level. As a result of vast individual heterogeneity, the choice of marital partners and the allocation of workers to jobs may have profound implications for income inequality, productivity, and welfare. The overarching theme of this thesis is the study of these allocations, their determinants and repercussions, with a particular focus on *assortative matching*.

Gary Becker has shown in his seminal work on marriage markets how and why heterogeneous men and women *sort* in a frictionless Walrasian environment (Becker, 1973, 1974). The key assumption in Becker's model is a common ordinal preference ranking of potential partners.[1] The empirical literature indeed found that individuals with high income and high educational attainment tend to marry each other.[2] Essentially, positive assortative matching is an outcome of homophily, the *love of the same*. It has the potential to shape the distribution of income in a society. Many authors have argued that marriage market sorting may increase income inequality and lead to segregation.[3]

[1] The sorting pattern then depends on the production structure of the economy. With a supermodular (submodular) production function, sorting will be positive (negative) because the agents' types act as complements (substitutes) for the joint output.

[2] The same is true along many other dimensions such as age, ethnicity, health, and physical characteristics, see Fisman et al. (2006) and Hitsch et al. (2010).

[3] Early studies on this question are Kremer (1997), Fernández and Rogerson (2001), Fernández et al. (2005), and Choo and Siow (2006). Burdett and Coles (1997) offer a theoretical analysis. More recently, Eika et al. (2014) find that marital sorting has only a limited effect on income inequality. In this debate, it is also important to distinguish within and between household inequality.

Two-sided assignment models are applicable to a wide range of problems. Today, frictional versions of the Beckerian model with sorting are an established tool for the analysis of labor markets, marriage markets, housing markets, CEO compensation, etc.[4] Frictions and decentralized trade have been introduced to the framework by Shimer and Smith (2000). Now, agents are unable to instantaneously locate their preferred partner. In models with random search, agents have to wait for the arrival of a potential partner and then weigh up the option values of forming the match or continuing search. Matching sets of acceptable and unacceptable types are formed. Unsuitable partners are rejected and, naturally, this lead to the existence of unmatched individuals in equilibrium. This is a common feature of frictional sorting models with the broader class of Diamond-Mortensen-Pissarides (DMP) search and matching models with equilibrium unemployment, which typically rely on a representative agent.

The first Chapter of this thesis attempts to work out a connection between the notion of sorting and the broader DMP literature on unemployment dynamics. It develops a sorting model of the labor market with ex-ante worker and firm heterogeneity, wage dispersion, and aggregate uncertainty. We then revisit one of the most prominent questions in the macro-labor literature in recent years, Shimer's unemployment volatility puzzle (Shimer, 2005). We shall see that sorting can improve our understanding of the cyclical dynamics of the labor market. The model is capable of explaining the full extent of unemployment fluctuations in the data. We also make a technical contribution by proposing a perturbation-based method that tackles the main computational burden of dynamic sorting models: infinite-dimensional endogenous distribution functions in the state space. This method allows us to analyze sorting models out of steady state.

Sorting is also an important topic in the broader empirical literature on wage dispersion. Abowd et al. (1999), using by now widely available matched employer-employee data, were the first to emphasize the huge importance of unobserved heterogeneity at the worker and firm level for wage dispersion. Card et al. (2013), a prominent paper building on Abowd et al. (1999), show that increasing wage dispersion in Germany is driven by ris-

[4]Chade et al. (2017) provide a comprehensive overview of the development and the relevant literature. Some recent examples of exciting new papers are Lise and Robin (2017), Lise et al. (2016), Bagger and Lentz (2016), Lindenlaub and Postel-Vinay (2016), and Eeckhout and Kircher (2017) for the labor market and Goussé et al. (2017) for the marriage market.

ing worker heterogeneity, rising firm heterogeneity, and increasing sorting.[5] The so-called AKM approach relies on strong assumptions of additive separability and monotonicity in order to identify the worker and firm components of wages from worker mobility patterns. This assumption is at odds with the basic assumption of sorting models, namely that worker and firm types interact and jointly determine output.[6]

The second Chapter of this thesis contributes to the debate about the incompatibility of the assumptions underlying AKM with sorting theory. It analyzes German matched employer-employee data through the lens of a structural sorting model in the spirit of Hagedorn et al. (2017) and studies how the allocation of workers to jobs changed during a period of profound institutional change.[7] We are able to reconcile AKM and sorting theory. The key finding is that empirical fixed effect models provide a valid approximation of observed wages and matching patterns for a large part of the data. For low-type workers, however, wages are decreasing in the type of the firm a worker is matched with. This is a prediction of sorting models which is at odds with the monotonicity assumption of AKM. After ranking both workers and firms, we show that low-type workers have become increasingly sorted into low-type firms over time, especially out of unemployment. This increase is driven by the sorting of low-type workers into their wage-maximizing matches at the bottom of the firm type distribution.

The last Chapter returns to the marriage market. Due to the similarity of assignment problems in labor and marriage markets, many empirical and theoretical tools can be used in both contexts. We study match formation and dissolution in frictional marriage markets under labor market uncertainty. The analysis makes use of a search model with transferable utility in which ex-ante heterogeneous men and women simultaneously search for partners in the marriage market and switch between employment and unemployment in the labor market. In the marriage market, individuals match assortatively on education and draw a match-specific shock component representing mutual affection. Divorces can happen due to both match-specific reasons and labor market transitions, e.g. job loss of one spouse. We structurally estimate our model using German micro

[5] Measured by the correlation of worker and firm-fixed effects.

[6] On this point see Eeckhout and Kircher (2011), Lopes de Melo (2016), Bonhomme et al. (2016), Lise et al. (2016), and Hagedorn et al. (2017) among many others.

[7] See Dustmann et al. (2014) for an overview.

data and decompose the observed flow of divorces into so-called "labor market divorces" and match-specific divorces. While more than 90% of divorces happen for reasons of changing mutual affection, the share of labor market divorces has increased significantly in Germany since the mid 2000s. Interestingly, most of this increase is driven by couples in which a previously unemployed woman starts working, more so for highly-educated women. These women also have the highest job-finding rates in our sample. As a consequence, the correlation between spouses' education has stopped increasing in Germany around 1999. It appears that the sorting trend has been mitigated by labor market effects in recent years, which is an interesting new contribution to the debate about marriage market sorting and inequality.

Each Chapter of this thesis has a separate introduction that includes a literature review. The notation in each chapter is self-contained. Selected additional results can be found in an Appendix at the end of each Chapter.

Chapter 1

Labor Market Dynamics with Sorting*

*This research is a revised version of a paper that previously circulated under the title "Wage Rigidity and Labor Market Dynamics with Sorting". Special thanks go to Robert Shimer for supporting this project and hosting me at the University of Chicago. Discussions with Matthias Doepke, Jan Eeckhout, Wouter den Haan, Leo Kaas, Philipp Kircher, Thibaut Lamadon, Rasmus Lentz, Christian Merkl, Jean-Marc Robin, Uwe Sunde, and many others have greatly improved this research. I also thank seminar and conference participants at Aarhus, Chicago, Essex, Halle, Konstanz, Munich, Northwestern, Nuremberg, SMYE 2014, ES European Winter Meeting 2014, AEA/ASSA 2015, T2M 2015, RES 2015, SaM 2015, Barcelona GSE Summer Forum 2015, ES World Conference 2015, CESifo Conference on Employment and Social Protection 2016, and EALE 2016 for many useful suggestions.

1.1 Introduction

To form a match in the labor market, both workers and firms have to expend time and resources in order to find a suitable partner. This coordination friction is the essence of the Diamond-Mortensen-Pissarides (DMP) search and matching framework and explains the coexistence of unemployed workers and vacant jobs in the labor market.[8] Intuitively, frictions can be understood as an outcome of heterogeneity: workers and firms differ in terms of skills and productivities. This makes finding a suitable match costly.

A big advantage of the standard DMP model is that it elegantly incorporates the coordination friction without making heterogeneity across workers and firms explicit. With the representative agent model, many interesting problems can be solved and studied analytically. Recently, however, the empirical literature on wage dispersion has made big advantages in identifying the importance of worker and firm heterogeneity.[9] There is positive sorting in the labor market, that is, a tendency of low (high) skill workers to work at low (high) productivity firms. We see an abundance of evidence in the literature that this is a prevalent feature of labor markets in many developed economies.[10]

This thesis chapter seeks to connect these new facts to a prominent class of search and matching models which is used to study unemployment and wage dynamics. Shimer (2005) points out that a simple dynamic DMP model fails to generate sufficient volatility in response to aggregate shocks, mainly because wages are fully flexible. Hornstein et al. (2011) add that standard search models do not generate as much wage dispersion as we observe in the data. We find that explicitly considering worker and firm heterogeneity and allowing for positive sorting enables search and matching models to match unemployment and wage dynamics.

We develop a dynamic DMP model with two-sided heterogeneity, labor market sorting, and aggregate shocks. The model features equilibrium wage dispersion, an endogenous wage rigidity, and sufficient amplification in response to shocks. At its core, the

[8]The main references for this class of models are Diamond (1982), Mortensen (1982), Pissarides (1985), and Mortensen and Pissarides (1994). Pissarides (2000) provides an excellent textbook treatment.

[9]The first paper in this spirit is Abowd et al. (1999), a seminal empirical contribution among the first using matched employer-employee data.

[10]There is evidence of PAM in Germany, Sweden, Italy, Denmark, and the United States; see Andrews et al. (2008, 2012), Card et al. (2013), Hagedorn et al. (2017), Lopes de Melo (2016), Bonhomme et al. (2016), Bartolucci et al. (2015), Lise et al. (2016), Bagger and Lentz (2016), and Chapter 2 of this thesis.

model is a Shimer and Smith (2000) economy. Sorting—generated by assuming a production complementarity between heterogeneous worker skills and firm productivities—implies that agents are not willing to match with every possible partner they meet when searching randomly. The standard coordination friction is thus enhanced because only a subset of all meetings generates new matches: endogenous matching sets contain the acceptable partner types for all worker and firm types. They are determined by the match surplus, which depends on the underlying production complementarity. In response to shocks, the surplus, the matching sets, and the distribution of worker and firm types adapt, generating rich dynamics and an amplification, in line with the data. Two new transmission channels are quantitatively important for the model's dynamics.

First, the model endogenously generates a wage rigidity of an empirically reasonable magnitude.[11] Shimer (2005) convincingly shows that the well-known lack of amplification in the standard model is due to the fully flexible wage with Nash bargaining, which reduces firms' incentive to create new jobs and limits the model's responsiveness to shocks. Adding labor market sorting to the picture leads to rigid wages, even with Nash bargaining. With sorting, wages depend on the *relative* labor market tightness, that is, the scarcity or abundance of other types in the bargaining worker-firm couple's matching sets. This property shields wages from fluctuations in aggregate labor market tightness and creates an endogenous rigidity.

Second, firms solve a type-specific dynamic job-creation problem: with free entry, they form an expectation about the future match surplus and the state-dependent distribution of unemployed worker types when deciding how many vacancies to post. We let the strength of the production complementarity be proportionate to labor productivity. This implies that the option value of being in a good match, measured by the surplus function, adjusts over-proportionately in response to shocks. This increases the firms' incentive to post vacancies and leads to amplification. Together, the endogenous wage rigidity and the dynamic job-creation problem propagate the model's response to shocks to a degree that brings it on par with empirical moments of labor market data.[12]

[11] We find a moderate elasticity of wages with respect to labor productivity of 0.75. This elasticity is close to benchmark estimates for the U.S. labor market reported by Haefke et al. (2013), who find an elasticity of 0.8 with a standard error of 0.4.

[12] Pries (2008) also analyzes how worker heterogeneity impacts the dynamics of a search and matching model. He introduces heterogeneity in a simple way—two types of workers and homogeneous firms—and

The state space of the model includes endogenous distribution functions which are infinite-dimensional. A technical contribution we make is a numerical procedure that allows keeping track of the model's complex state space. In order to compute the model's response to aggregate shocks, it is necessary to compute the adjustment path of the match surplus, the matching sets, and the endogenous type distributions. The key idea to approach this challenge is to define auxiliary state variables for the integral terms determining the state variables. This allows us to linearize the model around its steady state and use a perturbation technique to solve for the policy functions. To judge the accuracy of the computational approach, we plug the simulated data back into the model's Bellman equations and find that the errors due to discretization and approximation are on average no bigger than those made when solving standard search and matching models via perturbation.[13]

This distribution function in the state space is the main technical difference between the model developed in this chapter and Lise and Robin (2017), who show that search on the job with sequential auctions (building on Postel-Vinay and Robin (2002) and Robin (2011)) considerably simplifies the state space. They also find that a model with sorting and aggregate shocks fits many data moments well. The entry decision in their model, however, is static. Vacancy postings immediately adjust in response to shocks and, therefore, the model fails to match empirical vacancy dynamics. The dynamic entry decision proposed here overcomes this problem and generates persistent vacancy dynamics. To focus on the complexity that arises from the endogenous distributions in the state space, we abstract from search on the job. Studying the dynamic entry problem is an important complement to Lise and Robin (2017). The logical next step in this literature will be to develop a model with both search on the job and a dynamic entry, which is a fascinating avenue for future research.

The most prominent existing approaches to solve Shimer's unemployment-volatility puzzle in the literature rely on making wages less responsive to shocks, either by assuming that wages are completely rigid (Hall, 2005), by modifying the calibration in order to

finds some amplification related to a changing composition of the pool of unemployed workers over the business cycle, which does not, however, sufficiently amplify the model.

[13]Petrosky-Nadeau and Zhang (2017) show that solving the representative agent search and matching model in Hagedorn and Manovskii (2008) via log-linearization and perturbation creates a mean computational error of 3.75%. We find a mean error of 3.84% with slightly more dispersion.

increase the worker's outside option in the bargaining solution (Hagedorn and Manovskii, 2008), or by replacing Nash bargaining with an alternating offer bargaining game (Hall and Milgrom, 2008).[14,15] A counterpoint to these popular papers is Pissarides (2009), who modifies the firm entry problem, what is similar in spirit to this chapter. This chapter suggests that labor market sorting might provide an appealing alternative modification of the firm's problem in a broad class of models.

The remainder of this chapter is structured as follows: Section 1.2 introduces the setup and derives the stationary equilibrium of a hierarchical sorting model. Section 1.3 analyzes the comparative statics of the model to shed some light on the sources of amplification. Section 1.4 adds aggregate uncertainty and analyzes wage determination and firm entry in the fully dynamic model. Section 1.5 discusses the computational strategy and presents results from numerical simulations of the sorting model in comparison to a baseline search and matching model and U.S. labor market data. Section 1.6 concludes by discussing the findings in light of the related theoretical and empirical literature.

1.2 The Model

We construct a dynamic DMP model of the labor market with search frictions, sorting between heterogeneous workers and firms, and aggregate uncertainty. This Section presents the model's general setup and derives the stationary equilibrium of the model with a hierarchical production function. Section 1.4 introduces aggregate uncertainty in the form of a stochastic labor productivity process z. To simplify notation, this channel is neglected in deriving the stationary equilibrium.

[14] Hall (2005) shows that the volatility puzzle vanishes once wages are made fully inflexible, implying a counter-factual wage volatility of zero. Hagedorn and Manovskii (2008) show that the dynamics can be amplified by increasing the value of the workers' outside option of non-market activity. A higher calibrated value of the respective model parameter leads to lower wage outcomes in the Nash bargaining game, however, with the unrealistic consequence that a 15% increase in the value of non-market activity implies that the equilibrium unemployment rate doubles, see Hornstein et al. (2005) and Costain and Reiter (2008).

[15] These mechanisms are widely used to generate wage stickiness in DSGE models, for instance, to study the transmission of monetary policy shocks. A recent example is Christiano et al. (2016), who use the alternating offer bargaining game of Hall and Milgrom (2008) to induce wage inertia in a New Keynesian model.

Table 1.1: Distributions of Matched and Unmatched Worker and Firm Types

Distribution of	Relation	Aggregate Stock
Active matches	$g_m(x,y)$	$M = \iint g_m(x,y)\ \mathrm{d}x\mathrm{d}y$
Employed workers	$g_e(x) = \int g_m(x,y)\ \mathrm{d}y$	$E = \int g_e(x)\ \mathrm{d}x$
Unemployed workers	$g_u(x) = g_w(x) - g_e(x)$	$U = \int g_u(x)\ \mathrm{d}x$
Producing firms	$g_p(y) = \int g_m(x,y)\ \mathrm{d}x$	$P = \int g_p(y)\ \mathrm{d}y$
Vacant firms	$g_v(y) \rightarrow$ free entry	$V = \int g_v(y)\ \mathrm{d}y$

1.2.1 General Setup

A continuum of workers is endowed with heterogeneous skills $x \in [0,1]$ with a probability density function (pdf) $g_w(x)$. The measure of workers is exogenous and normalized to 1. A continuum of firms is heterogeneous in terms of productivity $y \in [0,1]$ with pdf $g_f(y)$. Denote by $g_m(x,y)$ the two-dimensional joint distribution of active (i.e. producing) (x,y) matches. The distributions of employed workers, $g_e(x)$, and producing firms, $g_p(y)$, can be obtained by integrating out the respective dimension of $g_m(x,y)$. The distribution of unemployed worker types, $g_u(x)$, is obtained by subtracting the distribution of employed workers from $g_w(x)$, which is exogenous and fixed. The distribution of vacant firm types $g_v(y)$ is determined by free entry. The distributions of active matches, employed/unemployed workers, and producing/vacant firms are equilibrium objects. They integrate to the stocks of employed workers, E, producing firms, P, unemployed workers, U, and vacancies, V. M is the stock of active matches, which must equal E and P. Table 1.1 summarizes the relations between the underlying density functions, the distributions of matched and unmatched workers and firms, and the aggregate variables.

The heterogeneity of workers and firms is assumed to be one-dimensional.[16] The model does not allow for imperfect information with respect to worker and firm types. All agents know their own type and the types of all potential partners they meet.

Time is discrete. Agents are infinitely lived, risk neutral, and they maximize their future discounted income streams. The common discount factor is β. A production complementarity between worker skills and firm productivities (detailed below) induces

[16] Both x and y can be viewed as a one-dimensional representation of a larger, multi-dimensional set of worker and firm characteristics. For a recent exploration of a multi-dimensional sorting model with random search, see Lindenlaub and Postel-Vinay (2016).

labor market sorting, that is, for every worker (firm) an optimal firm (worker) exists and matching with this optimal partner would maximize output.[17] Search frictions, however, prevent the formation of these optimal matches, as in Shimer and Smith (2000). In a setting with random search, heterogeneity, and sorting, not all meetings necessarily result in an employment relationship. The agents optimally choose partners by forming a matching set that comprises all acceptable types on the other side of the market. The decision as to whether a partner is acceptable or not is determined by the match surplus, $\mathcal{S}(x, y)$, which can be negative when the value of joint production falls short of the workers' and firms' outside options.

Labor is assumed to be the only production input hence capital is ignored. To focus on labor market sorting and the production complementarity between worker and firm types, we do not consider firm size and possible additional complementarities between workers within the same firm. This is equivalent to assuming that the firms' production function has constant returns to scale at the match level. Thus, firms' aggregate output is equal to the sum of what is produced by every individual worker at the firm. In other words, matches can be viewed as one-worker-one-machine relationships. The measure of active firms in the labor market is governed by free entry. Firms with a vacancy incur a per-period cost for keeping it open, representing expenses for posting the vacancy, screening applications, and the like. This cost is convex in the measure of type y vacancies posted. In the hierarchical sorting model considered below, the propensity to post vacancies will hence depend on firm type. The convex function c thus takes the measure of vacant firms $g_v(y)$ of type y as its argument. Firms enter the labor market by posting vacancies as long as the expected discounted value of production is at least as big as $c(g_v(y))$. In the hierarchical model, the convexity of $c(.)$ is critical to ensure a non-degenerate distribution of vacancies.

Only unemployed workers engage in random search.[18] We assume that meetings are governed by a standard Cobb-Douglas type matching function with constant returns to

[17] The case where all workers and firms are matched to their optimal partner corresponds to the Walrasian first-best allocation in Becker (1973).

[18] We abstract from search on the job to study a dynamic firm entry problem. In Lise and Robin (2017), search on the job with sequential auctions (Postel-Vinay and Robin, 2002) leads to a static entry problem, which has counter-factual implications for vacancy dynamics. Combining search on the job with dynamic firm entry in a sorting model would be an interesting future project.

scale, $M(U,V) = \vartheta U^{\xi} V^{1-\xi}$, where ξ $(1-\xi)$ is the elasticity of new matches with respect to unemployment (vacancies).[19] ϑ is a scaling parameter representing matching efficiency. The arrival rates are functions of aggregate labor market tightness $\theta = {}^{V}\!/_{U}$. $q_v(\theta) = {}^{M(U,V)}\!/_{V}$ is the rate at which vacant firms meet unemployed workers and, correspondingly, $q_u(\theta) = {}^{M(U,V)}\!/_{U}$ is the rate at which unemployed workers meet vacancies. $q_v(\theta)$ is decreasing and $q_u(\theta)$ is increasing in θ. Productive activity commences whenever a firm and a worker meet and find that they are jointly able to produce a non-negative surplus given their types. $\mathcal{S}(x,y)$ is then shared according to the standard Nash bargaining solution with worker bargaining power α.

Matches between firms and workers can be terminated for two reasons. They are subject to idiosyncratic separation shocks, which lead to immediate dissolution of the employment relationship. A match is subject to these shocks with an exogenous per-period probability δ. Additionally, in the presence of aggregate shocks (see Section 1.4), endogenous separations can occur at the margins of the agents' matching sets. When a negative productivity shock hits the economy, the surplus of previously marginally profitable matches may become negative. Since a negative surplus is always less than both parties' outside option, they prefer to separate.[20] In case of unemployment, workers receive flow benefits $b(x)$ every period, which represent the type-dependent value of home production or non-market activity.

The timing of the model is as follows: a period begins when the state of aggregate labor productivity z is revealed ($z = 1$ in steady state and thus muted in the following). Workers and firms form their optimal acceptance strategies based on the exogenous state z and the primitives of the model. Both endogenous and exogenous separations take place, following which new matches are formed. Workers and firms separated in the same period do not start their search until the next period. Finally, production commences and wages are paid.

[19]Note that using a linear search technology with heterogeneous workers and firms implies congestion effects between different worker and job types. Here, we stick to the Cobb-Douglas matching function for simplicity and comparability to other studies. A quadratic search technology, as used in Shimer and Smith (2000), eliminates the congestion externality. Nöldeke and Tröger (2009) extend Shimer and Smith (2000) to models with linear search technologies.

[20]This mechanism is similar to Mortensen and Pissarides (1994).

1.2.2 Production Functions

Let the output of a producing (x, y) match be denoted $F(x, y)$. The production function is non-negative and twice continuously differentiable. Labor market sorting in the model economy is induced by a complementarity between worker and firm types in production. We focus on positive assortative matching (PAM), which requires a supermodular production function, that is, positive cross-derivatives, $F_{xy} > 0$.[21]

More specifically, we rely on the equilibrium existence conditions in an optimal assignment economy with search frictions provided by Shimer and Smith (2000). Supermodularity of $F(x, y)$ alone is not sufficient in this setting. To ensure existence of a search equilibrium, $F(x, y)$, $F_x(x, y)$, and $F_{xy}(x, y)$ need to be log-supermodular.[22] These conditions also imply that the matching sets are nonempty, closed, and convex. Uniqueness cannot be established in this class of models in general. When solving the model numerically, we need to ensure that the mapping derived below is contracting in the respective parameter space by repeatedly solving the model for different initial conditions.

Depending on the functional form of the production function, labor market sorting arises from a comparative advantage (circular model) or an absolute advantage (hierarchical model) argument. In sorting models with absolute advantage, for example Shimer and Smith (2000), high type workers (firms) will always produce more than low types, no matter what type of firm (worker) they are matched with. In other words, the production function implies an unambiguous hierarchy, or ranking, of workers and firms because it is strictly increasing in both dimensions, $F_x(x, y) > 0$, $F_y(x, y) > 0$.

In the comparative advantage sorting model, the worker/firm type does not matter by itself; only the interaction of x and y determines output. Examples of sorting models with comparative advantage include the circular models in Marimon and Zilibotti (1999), Gautier et al. (2010), and Gautier and Teulings (2015). The circular production function takes as input only the distance $d(x, y)$ between x and y, measured along the circumference of a circle. Output is maximized for $d = 0$. As argued by Gautier and Teulings

[21] This implies, for any $x' > x$ and $y' > y$, $F(x', y') + F(x, y) \geq F(x', y) + F(x, y')$. See Topkis (1998) for an excellent treatment of supermodularity and complementarity in the context of lattice theory.

[22] This implies, for any $x' > x$ and $y' > y$, $F(x', y')F(x, y) \geq F(x', y)F(x, y')$, $F_x(x', y')F_x(x, y) \geq F_x(x', y)F_x(x, y')$, and $F_{xy}(x', y')F_{xy}(x, y) \geq F_{xy}(x', y)F_{xy}(x, y')$.

(2015), the circular model can be understood as a second-order Taylor approximation of a more general production technology.[23]

In this chapter, we work with the hierarchical sorting model because it is empirically more appealing: absolute advantage implies a ranking of workers and firms and most data used for constructing such rankings empirically and bring sorting models to the data, for instance, education, job tenure, firms size, or value added, are inherently hierarchical. In the sorting model with absolute advantage, high type workers (firms) always produce more, no matter with which type of firm (worker) they are matched. Following Shimer and Smith (2000), we use a simple functional form assumption that features log-supermodularity of the function itself, its first derivatives, and cross-derivatives to ensure the existence of a search equilibrium.[24]

$$F(x, y) = \exp(x \times y). \tag{1.1}$$

Recall that x and y are bounded on $[0, 1]$, so $\min(F(x, y)) = 1$ and $\max(F(x, y)) = \mathrm{e}$.[25] Let us define the surplus function

$$\mathcal{S}(x, y) = \mathcal{P}(x, y) - \mathcal{V}(y) + \mathcal{E}(x, y) - \mathcal{U}(x), \tag{1.2}$$

which depends on the four option value equations (see below) for a producing firm, a vacant job, an employed worker, and an unemployed worker for all (x, y) combinations. To capture the logic behind the matching sets in a way that is algebraically convenient, define a simple match indicator function:

$$\mu(x, y) = \begin{cases} 1 & \text{if } \mathcal{S}(x, y) > 0 \\ 0 & \text{if } \mathcal{S}(x, y) < 0 \end{cases} \tag{1.3}$$

$\mu(x, y)$ equals 1 whenever a firm of type y is willing to match with a worker of type x and vice versa. Whenever necessary, we will indicate that $\mu(x, y) = 1$ ($\mu(x, y) = 0$) by

[23] Earlier versions of this work have made use of circular production functions to induce sorting. Since it is not central to this version, we relegate the further discussion of circular models to Appendix A.2.

[24] That is, for any $x' > x$ and $y' > y$, $F(x', y')F(x, y) \geq F(x', y)F(x, y')$, $F_x(x', y')F_x(x, y) \geq F_x(x', y)F_x(x, y')$, and $F_{xy}(x', y')F_{xy}(x, y) \geq F_{xy}(x', y)F_{xy}(x, y')$.

[25] This functional form is also used in Teulings and Gautier (2004).

writing $\mu^+(x, y)$ ($\mu^-(x, y)$). The choice of with whom to match is determined solely by the surplus value function $\mathcal{S}(x, y)$. In case the surplus is positive, both parties will agree to form a match, so the decision is mutually consistent.[26] When the surplus turns out to be negative after a meeting, both parties prefer to continue their search due to their higher outside options.

$$\mathcal{E}(x,y) = W(x,y) + \underbrace{\beta\delta\mathcal{U}(x)}_{\text{separation}} + \underbrace{\beta(1-\delta)\max\{\mathcal{E}(x,y),\mathcal{U}(x)\}}_{\text{continued employment}} \tag{1.4}$$

$\mathcal{E}(x, y)$ represents the value of a type x worker employed at firm type y and consists of the flow payment of the match-specific wage, $W(x, y)$, plus the value of the two possible outcomes in the next period, discounted by β. With probability δ this match is subject to an idiosyncratic separation shock and the worker receives the value of unemployment, $\mathcal{U}(x)$. Accordingly, with probability $(1 - \delta)$, the worker continues to receive the value of employment at firm x. The max operator allows for the possibility that the value of employment falls below the value of unemployment in the next period, for instance, due to a productivity shock. The asset value of an unemployed worker of type x is defined as follows, with the limits of integration being equal to the boundaries of x and y, $[0, 1]$:

$$\begin{aligned}\mathcal{U}(x) = b(x) &+ \underbrace{\beta(1-q_u(\theta))\mathcal{U}(x)}_{\text{no meeting}} + \underbrace{\beta q_u(\theta)\int_0^1 \frac{g_v(y)}{V}\mu^+(x,y)\mathcal{E}(x,y)\mathrm{d}y}_{\text{successful match}} \\ &+ \underbrace{\beta q_u(\theta)\mathcal{U}(x)\int_0^1 \frac{g_v(y)}{V}\mu^-(x,y)\mathrm{d}y}_{\text{meet unacceptable firm}}\end{aligned} \tag{1.5}$$

In case of unemployment, all workers receive a type-specific value of home production $b(x)$ with $\frac{\partial b(x)}{\partial x} > 0$. In the following period, there is a probability $(1 - q_u(\theta))$ that the unemployed worker will not meet any firm and remains unemployed. With probability $q_u(\theta)$, the job finding rate, a meeting with a firm will occur. Whenever the type of firm y is an element of this worker's matching set (and vice versa), that is, $\mu(x, y) = 1$, a match is formed and production starts. Note that $g_v(y)/V$ under the integral sign represents the probability of meeting every specific firm type. This probability weights the surplus. In

[26] In the model with continuous distributions, the surplus is never exactly 0. Due to discretization, we allow for some smoothing of $\mu(x, y)$ when the surplus is very small.

case the firm turns out to be an unsuitable match ($\mu(x,y)=0$), both parties continue their search. The firms' asset value equations are symmetric:

$$\mathcal{P}(x,y) = F(x,y) - W(x,y) + \underbrace{\beta\delta\mathcal{V}(y)}_{\text{separation}} + \underbrace{\beta(1-\delta)\max\{\mathcal{P}(x,y),\mathcal{V}(y)\}}_{\text{continued production}} \tag{1.6}$$

The flow payment received by the firm in a productive employment relationship is the match-specific output $F(x,y)$ minus the wage, $W(x,y)$. In the next period, the match breaks up with probability δ, leading to the option value of a vacancy, or continues with probability $(1-\delta)$. Finally, the asset value of a vacancy is as follows:

$$\begin{aligned}\mathcal{V}(y) = -c(g_v(y)) + \underbrace{\beta(1-q_v(\theta))\mathcal{V}(y)}_{\text{no meeting}} + \underbrace{\beta q_v(\theta)\int_0^1 \frac{g_u(x)}{U}\mu^+(x,y)\mathcal{P}(x,y)\mathrm{d}x}_{\text{successful match}} \\ + \underbrace{\beta q_v(\theta)\mathcal{V}(y)\int_0^1 \frac{g_u(x)}{U}\mu^-(x,y)\mathrm{d}x}_{\text{meet unacceptable worker}}\end{aligned} \tag{1.7}$$

The cost of maintaining an open vacancy is determined by the convex function, $c(g_v(y))$, which must be paid every period. In the next period, there is the possibility of not meeting a worker, of meeting a suitable worker and filling the job, or of meeting an unsuitable worker and continuing to search.

The Nash bargaining solution determines how the surplus is divided in the event of a suitable match. The workers' bargaining power parameter is $\alpha \in (0,1)$. For both the worker and the firm, the respective share of surplus equals the additional value of being matched compared to the value of continued search, which serves as threat point in the bargaining game.

$$\alpha\mathcal{S}(x,y) = \mathcal{E}(x,y) - \mathcal{U}(x) \tag{1.8}$$

$$(1-\alpha)\mathcal{S}(x,y) = \mathcal{P}(x,y) - \mathcal{V}(y) \tag{1.9}$$

Note that the values of employment (production) and unemployment (a vacancy) are equalized when $\mathcal{S}(x,y)=0$, which is exactly the definition of the matching sets. It is possible to simplify the four Bellman equations (Equations (1.4) to (1.7)) by adding and subtracting the value of unemployment and a vacancy respectively. The surplus sharing rules of Equations (1.8) and (1.9) can then be plugged in and the surplus function shows up under the integral signs.[27]

$$\mathcal{E}(x,y) = W(x,y) + \beta\left(\mathcal{U}(x) + \alpha(1-\delta)\max\{\mathcal{S}(x,y),0\}\right), \tag{1.10}$$

$$\mathcal{U}(x) = b(x) + \beta\left(\mathcal{U}(x) + \alpha q_u(\theta)\int_0^1 \frac{g_v(y)}{V}\max\{\mathcal{S}(x,y),0\}\mathrm{d}y\right), \tag{1.11}$$

$$\mathcal{P}(x,y) = F(x,y) - W(x,y) + \beta\left(\mathcal{V}(y) + (1-\alpha)(1-\delta)\max\{\mathcal{S}(x,y),0\}\right), \tag{1.12}$$

$$\mathcal{V}(y) = -c(g_v(y)) + \beta\left(\mathcal{V}(y) + (1-\alpha)q_v(\theta)\int_0^1 \frac{g_u(x)}{U}\max\{\mathcal{S}(x,y),0\}\mathrm{d}x\right). \tag{1.13}$$

Stationary Equilibrium of the Hierarchical Model

Due the integral terms in the values of unemployment (Equation (1.11)) and a vacancy (Equation (1.13)), the hierarchical sorting model has no closed-form solution. Value function iteration on a discrete grid can be applied in this context to numerically approximate the model's stationary equilibrium. Technically, this procedure relies on the conjecture that in a given parameter space, the surplus function is a contraction mapping.[28]

The first step is to compute the fixed point of the surplus value function $\mathcal{S}(x,y)$, which is obtained by plugging Equations (1.10) to (1.12) into Equation (1.2). Note that $\mathcal{V}(y) = 0\ \forall\ y$ due to free entry, so the value of a vacancy drops out.

$$\begin{aligned}\mathcal{S}(x,y) = {} & F(x,y) + \beta(1-\delta)\max\{\mathcal{S}(x,y),0\} \\ & - \left(b(x) + \beta\alpha q_u(\theta)\int_0^1 \frac{g_v(y)}{V}\max\{\mathcal{S}(x,y),0\}\ \mathrm{d}y\right)\end{aligned} \tag{1.14}$$

The surplus function includes the workers' outside option in the large brackets due to the Nash bargaining assumption.[29] By solving the fixed point problem for every (x,y) combination, the equilibrium matching sets of all worker and firm types summarized in $\mu(x,y)$ are pinned down by $\mathcal{S}(x,y) \geq 0$. The agents' acceptance strategy is as follows:

[27]Note that $\int_0^1 \frac{g_v(y)}{V}\max\{\mathcal{S}(x,y),0\}\mathrm{d}y$ is equivalent to $\int_0^1 \frac{g_v(y)}{V}\mu(x,y)\mathcal{S}(x,y)\mathrm{d}y$.

[28]This approach to solve the model is similar in spirit to the procedures described in Shimer and Smith (2000), Hagedorn et al. (2017), and Lise and Robin (2017). Thanks go to Robert Shimer for sharing the code used to produce the numerical results in Shimer and Smith (2000).

[29]This is an important difference from the model of Lise and Robin (2017), who show that incorporating search on the job with sequential auctions (Postel-Vinay and Robin, 2002) removes the distributional terms from the surplus value and, hence, from the state space of the model.

for every (x, y) combination with a weakly positive surplus, $\mu(x, y)$ contains a 1 and a 0 otherwise.[30]

In the second step, knowing $\mathcal{S}(x, y)$ and $\mu(x, y)$, an equilibrium flow condition can be used to solve for the endogenous distribution of unemployed workers.

$$\delta g_m(x, y) = g_u(x)\; q_u(\theta)\; \frac{g_v(y)}{V}\; \mu(x, y). \tag{1.15}$$

The left-hand side of Equation (1.15) represents the number of dissolved active matches for every (x, y) combination in equilibrium. Without stochastic labor productivity, all match dissolutions occur exogenously with probability δ. On the right-hand side of the equation, the flow out of unemployment for the measure of all unemployed workers of type x, given by $g_u(x)$, is determined by the job-finding rate, $q_u(\theta)$, times the probability of meeting a specific firm type y, $g_v(y)/V$, times the match indicator function that takes the value 1 if the respective (x, y) combination produces a positive surplus. New matches are created solely for (x, y) combinations that produce a (weakly) positive surplus in equilibrium. Integrating the firm dimension out of Equation (1.15) and substituting in $g_w(x) - g_u(x)$ on the left-hand side yields the following expression for $g_u(x)$:

$$g_u(x) = \frac{\delta g_w(x)}{q_u(\theta) \int_0^1 \frac{g_v(y)}{V} \mu(x, y) \mathrm{d}y + \delta} \tag{1.16}$$

Note that integrating Equation (1.16) over x yields an expression comparable to the textbook equation pinning down equilibrium unemployment, that is, the Beveridge Curve:

$$U = \frac{\delta}{\delta + q_u(\theta) \iint \frac{g_v(y)}{V} \mu(x, y) \mathrm{d}y \mathrm{d}x}$$

Compared to the textbook version, the above expression has a double integral term in the denominator, which has to be smaller than 1 if $\exists (x, y). \mu(x, y) = 0$. Thus, equilibrium unemployment must be higher in the sorting model.

To determine the distribution of vacant firm types in the stationary equilibrium, we use the free entry condition for all y. Any firm type will post vacancies as long as the

[30] In practice, however, it is necessary to apply some smoothing to the acceptance strategy to ensure convergence. We allow for mixed strategy solutions close to the cutoff, following Hagedorn et al. (2017).

discounted value of the job is at least as high as the cost implied by $c(g_v(y))$.

$$c(g_v(y)) = \beta(1-\alpha)q_v(\theta)\int_0^1 \frac{g_u(x)}{U}\mu(x,y)\mathcal{S}(x,y)\mathrm{d}x. \tag{1.17}$$

Given $g_u(x)$, $\mu(x,y)$, and $\mathcal{S}(x,y)$, Equation 1.17 can be solved for the measure of vacancies posted by every firm type y in equilibrium. Integrating over $g_u(x)$ and $g_v(y)$ yields the aggregate stocks of unemployed workers and vacant firms, which in turn determine aggregate labor market tightness, arrival rates, and the flow of new meetings via the matching function.

The stationary equilibrium of the hierarchical sorting model consists of the objects $\{\mathcal{S}(x,y), \mu(x,y), g_u(x), g_v(y)\}$, which are jointly determined by the surplus value function (Equation (1.14)), the steady state flow condition (Equation (1.15)), and free entry of vacancies. The solution algorithm has two steps: it alternates between computing the fixed point of the surplus function for all (x,y) combinations and updating the distributions of unemployed workers and vacant firm types using steady-state flows and the free entry condition, respectively. This procedure is repeated until the decision rule converges.[31] The equilibrium can be computed with high precision in a relatively short amount of time. To help visualize the properties of the stationary equilibrium of the hierarchical sorting model, Figure 1.1 presents a projection of the surplus function on the type space, along with the matching cutoffs. The model is solved on a discrete grid with 1,000 worker and firm types. The hierarchy implied by the sorting model is immediately apparent, as the surplus increases quickly toward the upper-right corner. Interestingly, the surplus also increases toward the lower-left corner, that is, for matches between low type workers and low type firms. This property is a direct expression of the production complementarity in the sorting model. The surplus shrinks toward the cutoffs of the matching sets because the relatively higher ranked partner always needs to be compensated for his higher outside option. With absolute advantage and the production function of Equation (1.1), it is relatively more important to be optimally matched for low type workers and firms than for medium types. In other words, this type of worker has a higher incentive to be sorted and this is reflected in the surplus function.

[31] Convergence is achieved once the absolute difference of the surplus between two iterations is less than 10^{-12}.

Figure 1.1: Surplus and Matching Set Cutoffs in Equilibrium

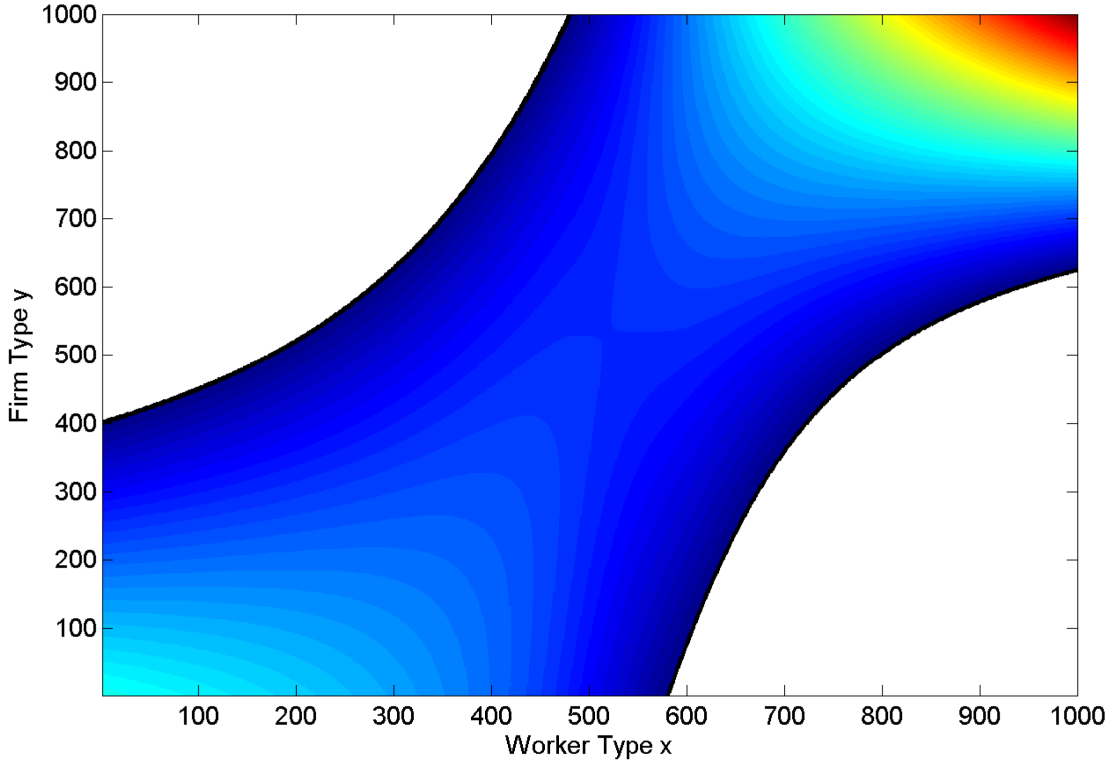

1.3 Comparative Statics

Let us now introduce aggregate labor productivity into the hierarchical sorting model. Labor productivity z can be imagined as an underlying technology that enables labor to be used productively. Thus, it influences the output of every match in equal measure:

$$F(x, y, z) = F(x, y) \times z, \tag{1.18}$$

so the output of a match between worker x and firm y with labor productivity z is simply the product of the match-specific part, $F(x, y)$, and the potentially time-varying but type-independent aggregate labor productivity, z. To better understand the properties that lead to an amplification of shocks in the dynamic model, we analyze the comparative statics of the model, shown in Figure 1.2. z changes from its steady-state value of 1 to 1.02. A key property of the production function $F(x, y, z) = z \times e^{xy}$ as introduced above is that the strength of the production complementarity is positively correlated with z.[32] This implies procyclical sorting, that is, the incentive to be optimally matched increases with labor productivity.[33]

In response to a change in z from 1.00 (blue curves) to 1.02 (red curves), the equilibrium of the hierarchical sorting model changes, as shown in Figure 1.2. First, the matching sets become wider, see Panel 1.2a. The set of acceptable matches increases for all worker and firm types. The boundaries do not shift in parallel but become (more) asymmetric in response to an increase in z. This property of the model turns out to be important for the wage adjustments in the dynamic sorting model. Panel 1.2b shows how the surplus function shifts upward for a firm of type $y = 750$. Note that even though the matching set becomes only slightly wider, the overall surplus of this firm increases more than is proportional in response to a change in z, which illustrates that being optimally sorted becomes more desirable as labor productivity increases. This property is key for it implies that the integral terms also increase more than proportionally in response to shocks in z. Panel 1.2c shows how the integral term of the firm in the job creation condi-

[32] This is a result of the (log-)supermodularity of the first derivative in the production function: $\frac{\partial F(x,y,z)}{\partial x} = F_x(x, y, z) = x\ z\ e^{xy}$.

[33] To date, there is no conclusive empirical evidence about the cyclicality of labor market sorting. Chapter 2 shows some evidence that the degree of labor market sorting in Germany seems to be roughly aligned with the business cycle.

Figure 1.2: Comparative Statics of a Change in z

(a) Matching Sets (b) Surplus of Firm Type $y = 750$

(c) Value of Integral by Firm Type y (d) Measure of Vacancies Posted by Firm Type y

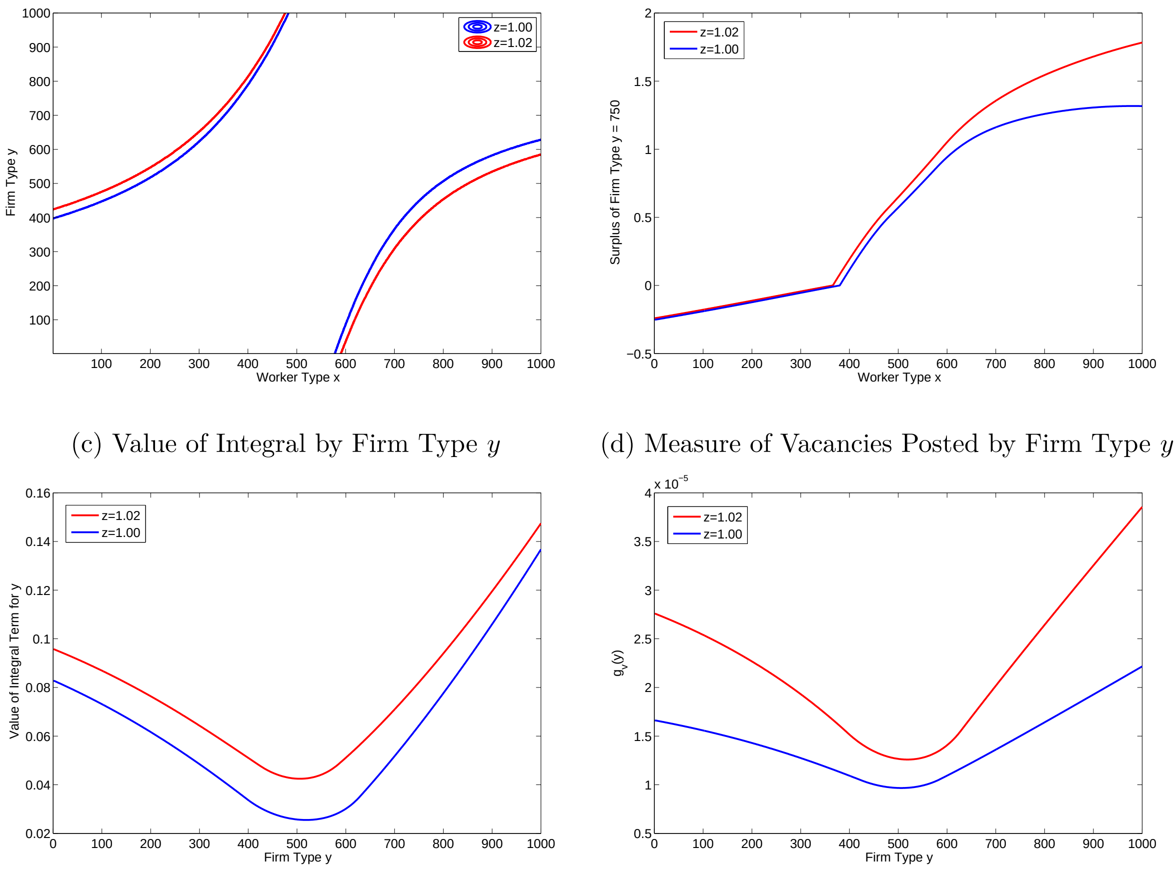

tion (Equation (1.17)) increases with z for all firm types y. Hence, all firms will have an incentive to post more vacancies as labor becomes more productive, see Panel 1.2d, which shows the measure of vacancies posted $g_v(y)$ by every firm type y. Note that vacancy posting increases particularly for low and high type firms, in line with the properties of the surplus function shown in Figure 1.1. Quantitatively, the depicted change of z from 1.00 to 1.02 leads to an overall increase of the vacancy rate by more than 50%. Unemployment falls by about 6%. These comparative statics results make a strong case for finding a significant amplification effect in the dynamic sorting model.

1.4 The Dynamic Sorting Model

In the dynamic sorting model, z is stochastic, so we study the properties of the model under aggregate uncertainty. Let z_t denote the realization of aggregate labor productivity

z in period t. I assume that z_t follows an AR(1) process (in logs):

$$z_t = \rho\ z_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \tag{1.19}$$

with the autocorrelation ρ and innovation parameter σ_ϵ calibrated below. To simplify notation, let z' be next period's realization $(t+1)$. Time subscripts are omitted in the following to avoid notational clutter.

Define the state of the system to be $\Omega(g_m(x,y,z),z)$. This state contains the exogenous state variable z and the endogenous state $g_m(x,y,z)$, which is the distribution of active matches. All endogenous objects depending on the state have Ω as an argument in the following.

With this notation, the dynamic surplus value function becomes

$$\begin{aligned} \mathcal{S}(x,y,\Omega) = {} & F(x,y,z) + \beta\mathbb{E}\left[(1-\delta)\max\{\mathcal{S}(x,y,\Omega'),0\}\right] \\ & - \beta\mathbb{E}\left[b(x) - \alpha q_u(\theta(\Omega'))\int_0^1 \frac{g_v(y,\Omega')}{V(\Omega')}\mu(x,y,\Omega')\mathcal{S}(x,y,\Omega')\ \mathrm{d}y\right]. \end{aligned} \tag{1.20}$$

$F(x,y,z)$ depends on the exogenous state only. Home production $b(x)$ does not depend on z. $\mathbb{E}$ is the expectation operator regarding the future aggregate state. It includes all information available in the period the expectation is formed.

1.4.1 Job Creation

The key for the analysis of sorting models out of steady state is the dynamic entry problem of the firm. We assume free entry for all firm types y, so $\mathcal{V}(y) = 0\ \forall\ y$. A firm of type y will post vacancies as long as the expected discounted value of production is at least as high as the cost implied by $c(.)$:

$$c(g_v(y,\Omega)) = \beta(1-\alpha)\mathbb{E}\left[q_v(\theta(\Omega'))\int_0^1 \frac{g_u(x,\Omega')}{U(\Omega')}\mu(x,y,\Omega')\mathcal{S}(x,y,\Omega')\ \mathrm{d}x\right]. \tag{1.21}$$

In the hierarchical sorting model, firm entry depends on the expectation of the integral over the firm-specific matching set, summarized in $\mu(x,y,\Omega')$, the surplus function $\mathcal{S}(x,y,\Omega')$, and the probability measure of meeting every specific worker type within the firms matching set, $\frac{g_u(x,\Omega')}{U(\Omega')}$.[34] Due to the state dependence, the job-creation condition in the hierarchical sorting model is richer than in standard search and matching models: $c = \beta(1-\alpha)\mathbb{E}[q_v(\theta(z'))\mathcal{S}(z')]$. Note that vacancy posting costs are constant in the standard problem and the expectation does not involve an integral term. The quantitative section of this chapter shows that the additional objects on the right-hand side of Equation (1.21) are critical for the model's response to aggregate shocks: all three endogenous variables of the augmented problem depend on the state and adjust in response to shocks. For a given firm, the surplus function shifts upward more than proportionally after a positive shock; recall Figure 1.2. Additionally, the cardinality of the matching sets increases since more potential matches now yield a positive surplus. There is an additional surplus to be realized both with workers who have been in the firm's matching set before *and* with new workers on the margins of the matching set. This channel is quantitatively important for a large share of the amplification result documented in Section 1.5.

1.4.2 Wage Formation

Let us now turn to the derivation of match-specific wages in the hierarchical sorting model using the Nash bargaining solution and free entry:

$$\mathcal{E}(x,y,\Omega) - \mathcal{U}(x,\Omega) = \frac{\alpha}{1-\alpha}\mathcal{P}(x,y,\Omega). \tag{1.22}$$

Plugging in the value functions, maximizing the Nash product, and some algebra yield an expression that determines the match-specific wage in the dynamic model, $W(x,y,\Omega)$:

$$\begin{aligned} W(x,y,\Omega) &= \alpha F(x,y,z) + (1-\alpha)b(x) \\ &\quad + (1-\alpha)\beta\alpha\mathbb{E}\left[q_u(\theta(\Omega'))\int_0^1 \frac{g_v(y,\Omega')}{V(\Omega')}\mu(x,y,\Omega')\mathcal{S}(x,y,\Omega')\ \mathrm{d}y\right]. \end{aligned} \tag{1.23}$$

The wage of a given worker is a convex combination of the match-specific output, $F(x,y,z)$, and the worker's outside option, his value of being unemployed. With respect to the integral term, the same logic as in the firms' job-creation problem applies: the outside

[34] The dynamic entry problem is the main difference between this research and Lise and Robin (2017). In their model with search on the job, entry is a static problem because the surplus function does not depend on the integral terms contained in Equations (1.11) and (1.13) of my model.

option depends on the expected value of the surplus with all other potential employers in the matching set, weighted by the distribution. The higher the surplus and the higher the probability of meeting every specific firm type in the matching set, the higher the bargained wage because the worker has more valuable outside options for which he needs to be compensated. After factoring out α, Equation (1.21) can be plugged into the wage equation to arrive at

$$\begin{aligned} W(x,y,\Omega) = \alpha \left(F(x,y,z) + c(g_v(y,\Omega))\mathbb{E}\left[\theta(\Omega') \frac{\int_0^1 \frac{g_v(y,\Omega')}{V(\Omega')} \mu(x,y,\Omega')\mathcal{S}(x,y,\Omega')\ \mathrm{d}y}{\int_0^1 \frac{g_u(x,\Omega')}{U(\Omega')} \mu(x,y,\Omega')\mathcal{S}(x,y,\Omega')\ \mathrm{d}x} \right] \right) \\ + (1-\alpha)b(x). \end{aligned} \tag{1.24}$$

Now, the same logic regarding the integral term over the matching set also applies for the firm in the bargaining game: the expected value of the surplus, the matching set, and the distribution of other unemployed worker types influence the negotiated wage negatively through the denominator: the more workers who are available in the firm's matching set, the better the firm's outside option of continued search, and the lower the match-specific bargained wage with worker type x because the firm needs to be compensated. Note that aggregate labor market tightness $\theta(\Omega')$ in front of the quotient cancels out with $1/V(\Omega')$ in the numerator and $1/U(\Omega')$ in the denominator:

$$W(x,y,\Omega) = \alpha \left(F(x,y,z) + c(y)\mathbb{E}\left[\frac{\int_0^1 g_v(y,\Omega')\mu(x,y,\Omega')\mathcal{S}(x,y,\Omega')\ \mathrm{d}y}{\int_0^1 g_u(x,\Omega')\mu(x,y,\Omega')\mathcal{S}(x,y,\Omega')\ \mathrm{d}x} \right] \right) + (1-\alpha)b(x). \tag{1.25}$$

The match-specific wage $W(x,y,\Omega)$ does not depend on aggregate labor market tightness, in contrast to the standard model. This is an important feature because it disconnects wages from fluctuations in aggregate tightness. Instead, the quotient, call it *relative* labor market tightness, determines wages. Relative labor market tightness is the expected ratio of the two integral terms, which include the distributions of vacancies and unemployed workers, and the surpluses with all types within the respective matching sets. (1.25) provides a natural generalization of the textbook wage equation to the framework with heterogeneous workers and firms. The key difference between the sorting model and the baseline DMP model is now obvious. Compare Equation (1.25) with its textbook counterpart in a dynamic DMP model with state z and homogeneous firms and workers:[35]

$$W(z) = \alpha(F(z) + \kappa \mathbb{E}\theta(z')) + (1-\alpha)b. \tag{1.26}$$

In the textbook model, the wage depends positively on aggregate labor market tightness. The higher $\theta = {}^{V}\!/_{U}$, the more difficult it is for firms to hire a worker. If there is fierce competition between many firms for relatively few unemployed workers, wages are higher. In this setting, the fully flexible Nash-bargained wage absorbs all the fluctuations in θ, which is the essence of the lack of amplification in the standard model (Shimer, 2005). In the sorting model, the impact of aggregate labor market tightness on wages is partially muted because θ is replaced by the relative labor market tightness:

$$\Theta(x,y,\Omega) = \frac{\int_0^1 g_v(y,\Omega)\mu(x,y,\Omega)\mathcal{S}(x,y,\Omega)\ \mathrm{d}y}{\int_0^1 g_u(x,\Omega)\mu(x,y,\Omega)\mathcal{S}(x,y,\Omega)\ \mathrm{d}x}.$$

The integral term in the numerator (denominator) represents a specific worker (firm) type's option value of search, that is, the value of the surplus function over the respective matching sets, properly weighted by the distribution of suitable types. Note that in the hierarchical sorting model, the endogenous distributions of unemployed workers and vacant firm types are typically not uniform, even with underlying uniform type densities. Relative labor market tightness can have an impact on the wage bargain that is very different from aggregate tightness in the standard model: aggregate tightness may be high (and unemployment low), but if the measure of unemployed workers within a firm's matching set is high, the firm has no incentive to pay a high wage and workers extract less, even if unemployment outside the firm's matching set is very low. Thus, the worker's bargaining position does not depend on scarcity or abundance of other unemployed workers outside the matching set of his potential employer. This mechanism de-links match-specific wages from aggregate labor market conditions.

[35]Note that this equation corresponds to Shimer's Equation (7) (Shimer (2005), p. 41), which is a slightly generalized version of Equation (1.20) in Pissarides (2000), p. 17. Equation (1.25) collapses to Equation (1.26) in the case of homogeneous workers and firms and constant output.

Wage Rigidity with Sorting

In the dynamic setting, aggregate fluctuations lead to an under-proportional wage adjustment for all (x, y) combinations where $\Theta(x, y)$ is smaller than aggregate labor market tightness θ. In the calibrated model, aggregate labor market tightness θ is normalized to 1. So whenever the denominator in $\Theta(x, y)$, the firm's integral term, is bigger than the worker's integral term in the numerator, $\Theta(x, y) < \theta$ will hold. This weakens the link between aggregate fluctuations and the wage bargain: wages no longer respond to shocks in a fully flexible manner. It is well known that mechanisms which disconnect wages and aggregate fluctuations are key to enable search and matching models to generate amplification. The finding that such a disconnect can endogenously arise in a setting with heterogeneity, sorting, and Nash bargaining is novel and a central contribution of this research. In the following, we show that the relative labor market tightness term is smaller than 1 because the surplus function exhibits an asymmetry which arises in relation to the workers' bargaining parameter.

Figure 1.3 plots the cutoffs of $\mu(x, y)$ in steady state for three values of the bargaining parameter (i.e., $\alpha = 0.5, 0.72, 0.95$). When $\alpha = 0.5$ (blue), the matching set cutoffs are slightly asymmetric. To see this, note that the lowest worker type is acceptable for all firm types below, roughly, 520, whereas the least productive firm is acceptable to all workers up to almost type 700. The asymmetry increases with α. The extreme case of $\alpha = 0.95$ (black) underlines this. The best worker type is extremely picky; he will only match with firms better than type 750. The best firm accepts workers down to rank 420 and is thus much more tolerant with respect to mismatch. Typically, calibrated search and matching models assign the worker a higher bargaining parameter. In the following, we use $\alpha = 0.72$ (red), as in Shimer (2005). This value implies a significant asymmetry, which is sufficient to ensure that $\Theta(x, y) < 1 \ \forall \ (x, y)$, implying an under-proportional wage adjustment in response to shocks. Intuitively, greater bargaining power on the worker's side translates to narrower matching sets because the worker can afford to be more picky when choosing the firm types with which to match. The worker will always extract a higher fraction of the surplus in a match, so the matching set can be smaller in order to be indifferent between the value of employment and unemployment. Firms, in turn, optimally choose wider matching sets because they command only a small share

Figure 1.3: Equilibrium Matching Sets and Worker Bargaining Power

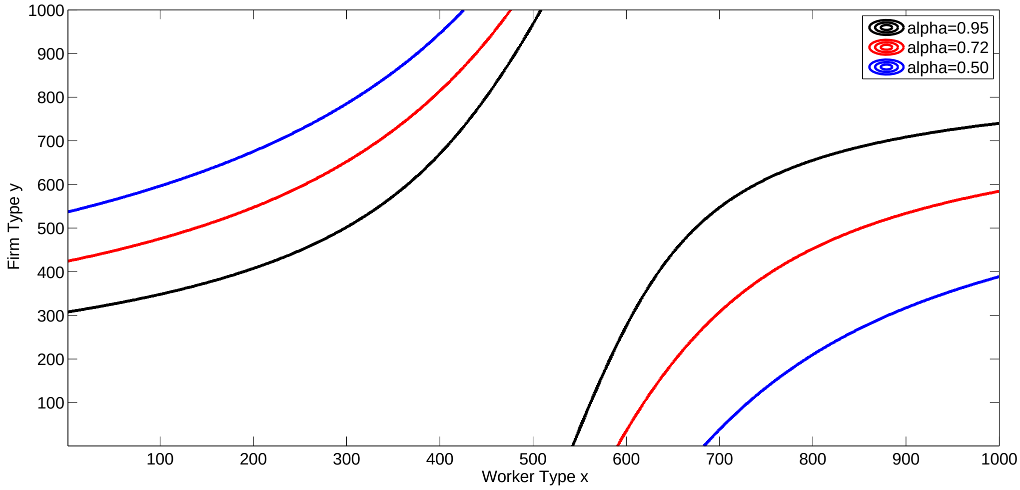

of the match-specific surplus. With small bargaining power, a wider matching set is necessary to equalize the values of production and a vacancy.

The hierarchical sorting model exhibits an asymmetry of the surplus function that translates into wider matching sets for the firms, a relative labor market tightness that is smaller than aggregate labor market tightness for all (x, y) combinations, and, accordingly, wages that do not fully adjust to shocks, and thus an endogenous wage rigidity arises. The numerical simulations in the following quantify to what degree this endogenous rigidity influences the model's dynamics in conjunction with the augmented job-creation condition.

1.5 Quantitative Analysis

We now test the quantitative importance of the augmented job creation and wage determination mechanisms in the dynamic hierarchical sorting model. It is common in this class of dynamic macro models to judge the model's empirical performance by its ability to match specific data moments. The Shimer Puzzle revolves around the search and matching model's ability (or lack thereof) to explain the volatility of the unemployment rate, the vacancy rate, aggregate labor market tightness, and the job-finding rate over the business cycle. We run numerical simulations in a stochastic environment where aggregate shocks to labor productivity drive the business cycle. The stochastic labor productivity process z is defined in Section 1.4 and calibrated below.

We find that second moments of simulated data from the sorting model are very close to moments from U.S. data, due to both the augmented wage formation mechanism and the additional margins in the job-creation condition. Both channels depend on the additional endogenous objects in the sorting model: the match surplus, matching sets, and type distributions. They all change with the state, so the computations need to handle forward-looking expectations of the match surplus, the matching sets, and the distributions of unmatched worker and firm types.

1.5.1 Computation

The state space Ω consists of the exogenous state z (labor productivity) and the endogenous state $g_m(x, y, z)$ (joint distribution of active matches). All other endogenous objects—the surplus, the matching sets, the distributions of unmatched worker and firm types, the stocks U and V, the arrival rates $q_u(\theta)$ and $q_v(\theta)$, and aggregate labor market tightness θ—follow from the state of the system Ω.

Keeping track of Ω's evolution is computationally challenging because it contains an endogenous distribution function and is hence infinite-dimensional. Discretization is inevitable, but the dimensionality of the state space is still huge for the case of 100 distinct worker and firm types, the number used in the simulations.

We approach this computational complexity by defining two auxiliary state variables for the integral terms in the job-creation condition and in the wage equation, which contain all the high-dimensional endogenous objects. This trick allows log-linearizing the model around its steady state. We can then apply standard techniques, in this case second-order perturbation, to solve and simulate the dynamic model. To keep track of the deviations of the auxiliary state variables during the simulation, we resolve the model conditional on every draw of the exogenous state z. We compute the numerical differentials of the auxiliary state variables with respect to z in a "brute-force" fashion for all (x, y) and z combinations. The key to success for this method is a fast algorithm that solves the model using the procedure outlined in Section 1.2.2. Once we know how the integral terms change in response to shocks, we can simply plug the deviations back into the log-linearized system to obtain policy functions.[36]

Since we make heavy use of discretization and numerical approximations to simulate the model, we check the reliability of the computational approach. It is well known that log-linearization around the steady state is error prone in nonlinear systems.[37] The approximation error becomes worse with large shocks that push the model far away from its steady state. Fortunately, the calibrated labor productivity process used in the context of U.S. labor market dynamics is not very volatile. The calibrated standard deviation of z is only 2%, so the model always remains in relatively close proximity to the steady state. We check the computations by plugging the simulated data back into the Bellman equations of our model to see whether the simulated data solve them. The mean computational error is quite small at 3.8%. This value is very close to the error one makes when solving simple dynamic search and matching models using log-linearization and perturbation, so the method of dealing with the additional complexity of the sorting model's state space does not appear to significantly increase computational errors.[38] Figures A1.1a and A1.1b in Appendix A.1 show the distribution of computational errors and a positive correlation of the errors with z.

1.5.2 Calibration

As in Shimer (2005), a time period is set to be one quarter. Table 1.2 shows the calibration of the model based on the U.S. labor market data used for the simulation exercise. To ensure comparability of the dynamics of the augmented model with the results in Shimer (2005), identical parameter values are used whenever possible. A value of 0.1 for the separation rate translates into an average employment spell of about 2.5 years during the United States in the relevant period (1951–2003). The quarterly discount rate is set to 0.012, representing an annual interest rate of roughly 5%. The discount factor

[36]We use Dynare (Adjemian et al., 2014) for the standard computations, e.g., checking stability of the system and computing policy functions. The actual calculations use an external routine which is called in each iteration during simulation.

[37]Using more accurate projection methods to simulate the dynamic sorting model with its complex state space is beyond the scope of this project.

[38]Petrosky-Nadeau and Zhang (2017) show that solving the representative agent search and matching model in Hagedorn and Manovskii (2008) via log-linearization and perturbation creates a mean computational error of 3.75%. We find a mean error of 3.84% with slightly more dispersion.

(as it appears in the model equations) is thus $\beta = 1/1.012 \approx 0.99$. The matching function elasticity is set to 0.72, in line with Shimer (2005), which is within the empirically supported range from matching function estimations reported by Petrongolo and Pissarides (2001). We set the bargaining parameter equal to the matching function elasticity applying the Hosios (1990) condition for socially efficient vacancy posting in decentralized equilibrium.[39]

Several parameters need to be recalibrated in the hierarchical sorting model. The value of non-market activity $b(x)$ is type dependent. We calibrate it to be $0.223 \times \underset{x}{\operatorname{argmax}}\ F(x, y)$. This implies that home production $b(x)$ has a mean of 40% of the output a worker of type x can produce in his optimal match. This assumption is a natural extension of Shimer (2005), who assumes a constant b of 0.4 when output is normalized to 1. The efficiency parameter of the aggregate matching function, ϑ, needs to be increased in the sorting model to take into account that not all meetings result in matches. A value of 2 implies, along with the other parameter values, that the *net* job finding rate, that is, the rate of matches that are formed after a meeting, is close to the value Shimer (2005) constructs from the data, which is 1.355 (quarterly).

The convex vacancy posting cost function takes the following form:

$$c(g_v(y)) = \frac{c_0}{1 + c_1} g_v(y)^{1+c_1}$$

c_0 and c_1 are set to the values shown in Table 1.2 to target a steady-state aggregate labor market tightness of 1.[40] The calibrated economy has a steady-state unemployment rate of about 7.8%.[41]

The stochastic labor productivity process z can be imagined as an underlying technology that enables labor to be used productively. It is type independent and affects all matches in equal measure. As in Shimer (2005), it is normalized to 1 in steady state and calibrated to resemble empirical labor productivity in the United States over the relevant

[39]It is unclear whether the Hosios condition holds in the sorting model. We do not claim that vacancy posting is socially efficient here.

[40]We use more curvature as compared to the values estimated in Lise and Robin (2017) to ensure quick convergence.

[41]This is slightly higher than typically targeted values of steady-state unemployment. The reason lies in the adapted Beveridge Curve (see Equation (1.2.2)), which implies higher equilibrium unemployment when the matching sets do not cover the whole type space.

Table 1.2: Quarterly Calibration of the Sorting Model for the U.S. (1951–2003)

Parameter	Symbol	Value	Source
Discount factor	β	0.99	Shimer (2005)
Separation rate	δ	0.1	Shimer (2005)
Workers' bargaining power	α	0.72	Shimer (2005)
Matching function elasticity	ξ	0.72	Shimer (2005)
Matching function constant	ϑ	2	Recalibrated
Value of nonmarket activity	$b(x)$	$0.223 \times \underset{x}{\operatorname{argmax}}\ F(x,y)$	Recalibrated
Vacancy posting costs	c_0	0.03	New Parameter
	c_1	0.4	New Parameter
First order autocorrelation	ρ	0.765	Hagedorn
Standard deviation	σ_ϵ	0.013	& Manovskii (2008)

period of time. We follow Hagedorn and Manovskii (2008) and set up stochastic labor productivity as a first-order autoregressive process:[42]

$$z_t = \rho\ z_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2). \tag{1.27}$$

$\rho \in (0,1)$ captures the degree of first-order autocorrelation of the AR(1) process. Innovations are drawn from a Gaussian distribution with mean 0 and standard deviation σ_ϵ. Both parameters are set to match quarterly U.S. labor productivity.[43] All values in Table 1.2 are based on quarterly data. Shimer's (2005) simulation results as well as our results are reported as deviations from a HP trend, which is conventional in the literature.[44]

[42]Many closely related papers use more general Markov chains to add a stochastic dimension to the model. An AR(1) process is a homogeneous Markov process iff the error terms are i.i.d. We prefer to use the AR(1) process under this assumption for computational reasons.

[43]Shimer (2005), Hornstein et al. (2005), and Hagedorn and Manovskii (2008) report the parameter values necessary to represent U.S. labor productivity "as seasonally adjusted quarterly real average output per person in the non-farm business sector constructed by the BLS" (Hagedorn and Manovskii (2008), p. 1694).

[44]The Hodrick-Prescott (HP) filter is a technique for decomposing the trend and the cyclical component of a time series (Hodrick and Prescott, 1997). Shimer (2005) sets the smoothing parameter of the filter to $\lambda = 10^5$ instead of to the more common value of $\lambda = 1600$ for quarterly data. This makes the cyclical component less volatile and more persistent. We use the same value as Shimer to generate comparable moments. Hornstein et al. (2005) point out that a more volatile trend, using the common smoothing parameter $\lambda = 1600$ for quarterly data, "has almost no effect on the relative volatilities" (p. 23).

Table 1.3: Actual and Simulated Standard Deviations of Labor Market Variables

	Standard deviations	U	V	θ	$q_u(\theta)$	z	$F(x, y, z)$
1.	U.S. data	0.190	0.202	0.382	0.118	0.02	-
2.	Results of Shimer (2005)	0.009	0.027	0.035	0.010	0.02	-
3.	No sorting, no heterogeneity	0.009	0.026	0.035	0.010	0.02	-
4.	Sorting, hierarchical model	0.102	0.277	0.380	0.168	0.02	0.06

Note: Rows 1 & 2: Based on Tables 1 and 3 in Shimer (2005), pp. 28, 39. Calculated based on quarterly U.S. data, 1951–2003. Rows 3 & 4: Standard deviations of simulated data from my model with and without sorting. All moments come from HP-filtered data with $\lambda = 10^5$.

Using the calibration in Table 1.2, the model produces realistic amounts of wage dispersion and labor market sorting. The standard deviation of log wages in equilibrium is about 0.418. Spearman's rank correlation coefficient, which is a measure for the degree of labor market sorting, is 0.095. This is only a moderate degree of positive sorting, so the strong complementarity of worker and firm types assumed via the form of the production function does not translate into a strongly sorted distribution of matches, in line with what we know from U.S. data.[45]

1.5.3 The Amplification Effect of Sorting

We find that the hierarchical sorting model produces a large amount of amplification in response to shocks. Second moments of simulated time series data are of the same order of magnitude as the volatility observed in U.S. labor market data for the relevant period of time. In particular, the simulated standard deviations of unemployment, vacancies, labor market tightness, and the job-finding rate are much closer to empirical second moments than simulated data from standard search and matching models. Table 1.3 compares my results to those of Shimer (2005) and the empirical data moments.

The first two rows of Table 1.3 show the well-known unemployment volatility problem emphasized by Shimer (2005). The standard deviations of unemployment, U, vacancies, V, labor market tightness, θ, and the job-finding rate, $q_u(\theta)$, in simulated time series data

[45] Some evidence for the degree of labor market sorting in the United States is presented by Lise et al. (2016), who make a parametric assumption about the production function (CES) and directly estimate the elasticity of substitution. They find evidence for a relatively small degree of positive sorting, but the magnitude of the estimated substitution elasticity is not readily comparable to a rank correlation coefficient.

miss the standard deviations in the data by a factor of about 10 to 20. As a first exercise, we run simulations of a model from which worker and firm heterogeneity is removed with output set to 1 and the exact same calibration as in Shimer (2005). The results are in the third row of Table 1.3. They are nearly exactly the same as Shimer's when sorting and heterogeneity are switched off.

The main results are reported in the fourth row of Table 1.3. Note that the volatility of z, the calibrated underlying labor productivity process, is the same in all models. In the hierarchical sorting model, however, overall match-specific output ($F(x, y, z)$) fluctuates somewhat more than z. This is because, as explained in the comparative statics exercise in Section 1.3, the chosen production function implies that sorting itself is procyclical. It becomes relatively more valuable to be optimally matched as z increases, so $F(x, y, z)$ is more volatile in response to shocks than z alone. This is included in the amplification effect we find.

The second moments of simulated time series data from the hierarchical sorting model are much closer to the data than are those from the standard model without sorting. The standard deviation of the HP-filtered time series of labor market tightness (0.380) is very close to the data (0.382). The sorting model generates realistic dynamics of labor market tightness via two transmission channels: the additional margins of adjustment in the firms' job-creation problem (Section 1.4.1) and the endogenous wage rigidity (Section 1.4.2). The simulated standard deviations of vacancies (V) and the job-finding rate (q_t^u) are also much higher than in the baseline model; they even overshoot their empirical counterparts to some extent. The standard deviation of unemployment (U) is amplified as well but remains somewhat lower than the empirical value.

To illustrate the dynamics of the hierarchical sorting model, Figure 1.4 shows impulse response functions of six key variables: unemployment, vacancies, aggregate labor market tightness, and the job-finding rate, as well as the autoregressive labor productivity process and wages. In response to a positive shock to labor productivity, unemployment falls by about 6 percentage points initially and shows a hump-shaped return to steady state. Vacancy posting, job finding, and aggregate tightness of the labor market show strong positive reactions directly after the shock. All impulse responses show a high degree of persistence as well as realistic correlations and cyclical properties. For instance,

Figure 1.4: Impulse Response Functions of Key Variables in the Search and Matching Model with Sorting

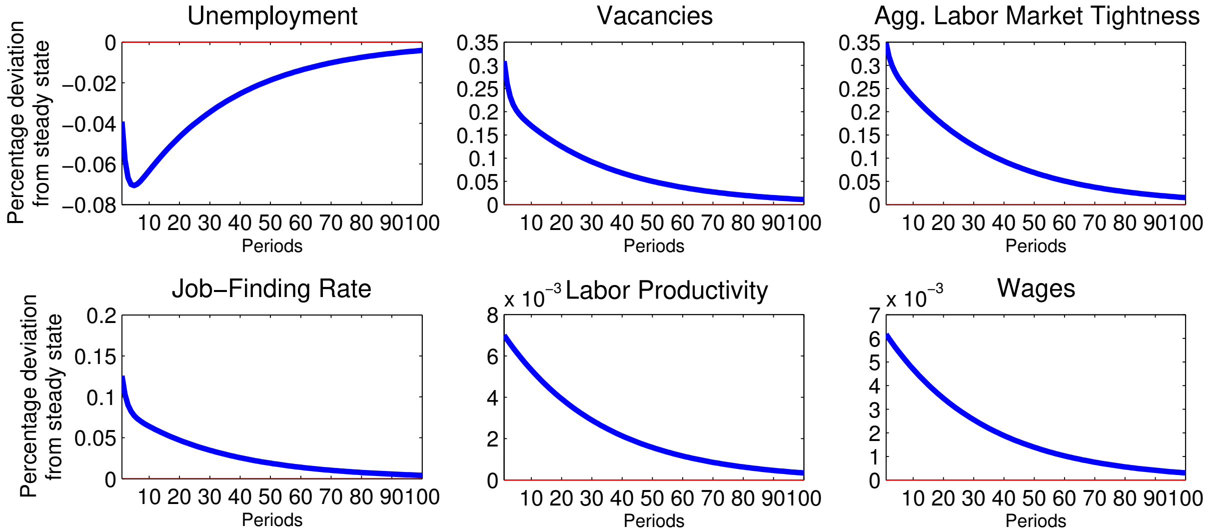

unemployment and vacancies move in opposite directions in response to the shock and are thus highly negatively correlated (Beveridge Curve). Note also the under-proportional adjustment of wages to the shock in labor productivity: the initial adjustment of wages is roughly 85% of the jump in labor productivity in the depicted example, so the endogenously generated wage rigidity becomes apparent.

Let us now take a closer look at the simulation results. The sorting model includes two channels that lead to amplification: job creation and the endogenous wage rigidity. How does each channel contribute quantitatively to the improved empirical performance and how large is the implied degree of wage rigidity?

1.5.4 The Degree of Wage Rigidity

That rigid wages can amplify search and matching models' response to shocks is well known in the literature and typically based on altering the assumptions underlying wage determination. (Hall, 2005; Hagedorn and Manovskii, 2008; Hall and Milgrom, 2008). The argument in favor of rigid wages in the context of the sorting model is novel and different. It was shown in Section 1.4.2 that rigid wages originate from an equilibrium property of the model: with sorting, the values of workers' and firms' outside options are asymmetric, leading to a skewed surplus function. Workers optimally choose narrower matching sets because they have greater bargaining power. This is based on the fact

Figure 1.5: Simulated Time Series of Wages and Labor Productivity

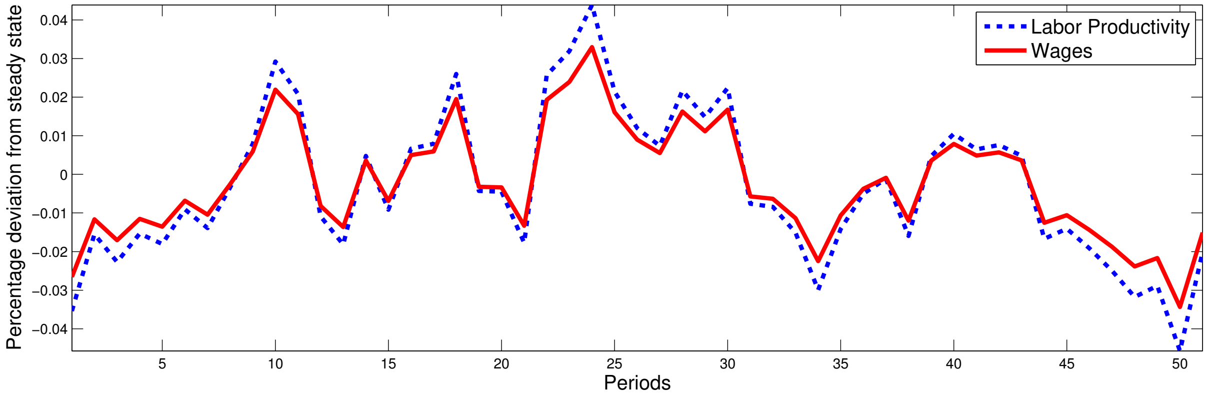

Note: Fluctuations around trend of simulated data for labor productivity and wages, HP-filtered with $\lambda = 10^5$. Hierarchical model. The 50-period time frame was randomly chosen.

that workers receive a positive value of home production in the event of unemployment, whereas a vacant firm earns nothing.

Perfectly flexible wages—as in the textbook model—always fully respond to changes of labor productivity, leading to a one-to-one co-movement. The elasticity of wages with respect to labor productivity is 1. Totally rigid wages, on the other hand, would not react at all to changes in labor productivity, implying an elasticity of 0 (as in Hall (2005)). Figure 1.5 shows how simulated wages (red, solid) and labor productivity (blue, dashed) deviate from trend in the model with sorting. The data stem from a simulation of the hierarchical model and the 50-period timeframe is randomly chosen. In the standard search and matching model, these two time series are congruent. In the model with sorting, however, wages turn out to be less volatile and do not fully adjust, as can be seen in Figure 1.5. Wages follow labor productivity closely but adjust imperfectly. The elasticity of wages with respect to labor productivity, a measure for the rigidity of wages, hence must be somewhat smaller than 1. This is exactly the result one would expect from comparing the sorting model to the baseline model. In the baseline model, wages are too responsive. After a favorable shock, workers soak up most of the extra productivity via immediately adjusting wages. This leads to an insufficient responsiveness of other model variables, particularly vacancies, due to a lack of incentives for the firm. In the model with sorting, however, the one-to-one link between wages and labor productivity is weakened. The endogenous wage rigidity, which is visible in Figure 1.5, limits the extent to which wages adjust in response to shocks. This increases firms' incentives to create new jobs and leads to amplification.

We go to the empirical literature to investigate whether the model-generated rigidity is of a reasonable magnitude. Haefke et al. (2013) focus primarily on the different degrees of wage rigidity for newly hired workers as compared to established employment relationships. I can compare the reported wage elasticity with respect to labor productivity for new hires to my results.[46] Haefke et al. (2013) find an elasticity of wages with respect to labor productivity of 0.8 with a standard error of 0.4. Using our simulated data, this elasticity can be computed as the estimated coefficient η_1 from a simple linear regression of wages on labor productivity in logs and first differences:

$$\Delta \log W_t(x, y, z) = \eta_0 + \eta_1 \Delta \log z_t + \varepsilon_t \tag{1.28}$$

Running this regression yields a wage elasticity of $\eta_1 = 0.751$ for the hierarchical model, which is well within the empirically supported range. Hagedorn and Manovskii (2008) also compute an elasticity from U.S. wage and productivity data and report a coefficient of 0.449, which is at the lower end of the empirically supported range. The elasticity/derivative implied by the alternating-offer bargaining game proposed by Hall and Milgrom (2008) is about 0.7, close to what we find. It is reassuring that our rigidity lies close to these benchmarks.

1.5.5 Rigid Wages vs. Job Creation

It is instructive to decompose the sorting effect into she shares explained by the endogenous wage rigidity and by the modified job creation condition. Sorting influences firms'

[46] Haefke et al. (2013) note that this value "is an appropriate and informative calibration target for search and matching models." (p. 898). Since the baseline search and matching model with Nash bargaining is essentially a model of new hires, the elasticity of wages with respect to labor productivity for this group is a reasonable target to match. Note that wages do not play an allocational role in a random search model. The Nash bargaining solution simply determines how the surplus is shared in every time period given the state of the model in the same period. Thus, the length of an employment spell does not influence wages and there is no meaningful distinction between new hires and existing employment relationships.

forward-looking vacancy posting decisions. Recall the firms' entry condition:

$$c(g_f(y)) = \beta(1-\alpha)\mathbb{E}\left[q_v(\theta(\Omega'))\int_0^1 \frac{g_u(x,\Omega')}{U(\Omega')}\max\{\mathcal{S}(x,y,\Omega'),0\}\ \mathrm{d}x\right]. \tag{1.29}$$

After a positive productivity shock, the inversely U-shaped surplus function shifts upward as shown in the comparative statics exercise. In our parametrization, firms choose wider matching sets in response to positive shocks because the surplus from an increased number of potential matches is larger than zero. This leads to a higher expectation of the value of matches with all workers in the firms' matching sets. An opposing force is the changing distribution of unemployed worker types. Since unemployment falls in response to a positive shock, it will be harder for firms to meet workers within their matching sets.[47]

The combined effect of the additional channels in the sorting model is that firms post more vacancies in response to shocks by increasing $g_v(y)$ subject to the convex cost function. Relative to the standard model, this creates an amplification of shocks because the value of being optimally sorted increases and the matching set and surplus function adjust accordingly. For a productivity shock of equal magnitude, the right-hand side of Equation (1.29) increases more than the expected future value of a job in the standard model without sorting. Additionally, this implies that wages do not need to be extremely rigid in the sorting model to generate sufficient volatility.

To quantify the contributions of the endogenous wage rigidity and the larger number of vacancy postings to overall amplification, let us take the elasticity of wages with respect to labor productivity calculated above as given and impose it on a model without sorting and heterogeneity. The gap in volatility, which cannot be accounted for by the effect of rigid wages alone, must then be the effect of sorting on job creation. The last row of Table 1.4 shows the results of this exercise.

The volatility of labor market variables with an imposed wage rigidity of $\eta_1 = 0.751$ is too small compared to the data. As expected, the rigidity amplifies the model's response, but only by a factor of about 2–3 (depending on the variable). Thus, the relatively moderate model-generated wage rigidity does not suffice to generate sufficient amplification.

[47]The described adjustments are not unambiguous and depend on the model's parameterization. For instance, as the value of being optimally matched increases after a positive shock, the matching sets could also become smaller in principle.

Table 1.4: Amplification Effect of an Imposed Wage Rigidity

	Standard deviations	U_t	V_t	θ_t	q_t^u
1.	U.S. data	0.190	0.202	0.382	0.118
2.	Results of Shimer (2005)	0.009	0.027	0.035	0.010
3.	Sorting, hierarchical model	0.102	0.277	0.380	0.168
4.	Rigidity imposed, $\eta_1 = 0.751$	0.031	0.084	0.114	0.051

Note: Rows 1 & 2: Based on Tables 1 and 3 in Shimer (2005), pp. 28/39. Calculated based on quarterly U.S. data, 1951–2003. Rows 3 & 4: Standard deviations of simulated data from my model. All moments come from HP-filtered data with $\lambda = 10^5$.

This is not surprising. Additionally, the effect of endogenous separations, which is present in the dynamic sorting model, is quantitatively small. This is in line with the comparative statics result that the matching sets do not change much in response to changes in labor productivity. We conclude that the large amplification effect of sorting must be primarily driven by additional job creation.

1.6 Conclusions

This research proposes a search and matching model with two-sided heterogeneity, sorting, and aggregate shocks. The model's relationship to previous literature is best understood from two reference points. The first reference point is the optimal assignment model of heterogeneous agents in a frictionless market following Becker (1973), the classical sorting model. In this frictionless environment, a production complementarity makes it optimal for all types to match with only their unique optimal counterpart. Shimer and Smith (2000) take this model out of its Walrasian equilibrium by adding frictions and making search costly, thus introducing the concept of matching sets that do not cover the whole type space.

The second reference point is the baseline DMP search and matching model of the labor market and its dynamic version.[48] The DMP model has been very successful in explaining equilibrium unemployment and a number of important stylized facts of labor market data (e.g., the Beveridge Curve). In the DMP model, every agent is willing to

[48]The main references for this class of models are Diamond (1982), Mortensen (1982), Pissarides (1985), and Mortensen and Pissarides (1994). Pissarides (2000) provides an excellent textbook treatment.

match with every other agent, but frictions limit the number of encounters in the labor market and, therefore, equilibrium unemployment exists.

We construct a model in between these two reference points: in a labor market with two-sided heterogeneity, search frictions, sorting, and aggregate shocks, suboptimal matches between heterogeneous jobs and workers are formed and persist through time. This setup has the potential to match the data well. It generates sufficient volatility in response to shocks due to an endogenous wage rigidity and the firms' dynamic entry problem. Recent advances in empirically identifying the extent of sorting in labor markets[49] suggest that it provides an empirically supported complement to existing approaches to better align search and matching models with the data.

Beyond the search and matching literature, wage and price rigidities play a key role in modern macroeconomics. Blanchard and Galí (2010) survey a number of papers that explore the implications of real and nominal rigidities in different variants of real business cycle (RBC) or New-Keynesian macro (NKM) models. The models developed by Merz (1995), Andolfatto (1996), Christoffel and Linzert (2005), Krause and Lubik (2007), Faia (2008), and Gertler and Trigari (2009) are all examples of DSGE models that integrate variants of price and wage rigidities to create persistence and sufficient amplification for the analysis of different kinds of shocks. Gertler and Trigari (2009), for example, adapt the well-known Calvo (1983) pricing structure for wage formation in the labor market. Rigid wages are generated by allowing only a fraction of matches to renegotiate wages in every period. Another recent example is Christiano et al. (2016), who use the alternating-offer bargaining game of Hall and Milgrom (2008) to induce wage inertia in a New Keynesian model. In the light of this literature, a key contribution of this chapter is to show that a search and matching model with two-sided heterogeneity and sorting can generate rigid wages endogenously and match the large unemployment fluctuations in the data.

[49]See Chapter 2 as well as Andrews et al. (2008, 2012); Card et al. (2013); Hagedorn et al. (2017); Lopes de Melo (2016); Bonhomme et al. (2016); Bartolucci et al. (2015); Lise et al. (2016); Bagger and Lentz (2016).

Appendix to Chapter 1

A.1 Computational Appendix

To check the accuracy of the computational method described in Section 1.5.1, we plug simulated data from the dynamic sorting model back into the Bellman equations of the model. Ideally, the simulated data would solve these equations exactly. However, we make heavy use of discretization and approximation techniques, so it is reasonable to expect some imprecision. For convenience, we use the wage equation for this test because it contains both the firms' and the workers' integral terms:

$$W(x,y,\Omega) - \alpha\left(F(x,y,z) + c(g_v(y,\Omega))\mathbb{E}\left[\frac{\int_0^1 g_v(y,\Omega')\max\{\mathcal{S}(x,y,\Omega'),0\}\ \mathrm{d}y}{\int_0^1 g_u(x,\Omega')\max\{\mathcal{S}(x,y,\Omega'),0\}\ \mathrm{d}x}\right]\right) \\ - (1-\alpha)b(x) \stackrel{?}{=} 0.$$

Solving the dynamic sorting model by log-linearization and perturbation results in a mean computational error of 3.84%. The 2.5th, 50th, and 97.5th percentiles of the distribution are −15.0%, 4.67%, and 17.6%, respectively. This distribution is slightly left skewed due to the fact that the model's response to shocks is not symmetric around the steady state, for example, because of endogenous separations, which only happen after negative shocks. Figures A1.1a and A1.1b show a histogram of the computational errors and a scatter plot that shows the positive correlation of the errors with z.

A recent paper by Petrosky-Nadeau and Zhang (2017) can serve as a benchmark for the size and distributions of the errors. The authors show that solving a representative agent search and matching model via log-linearization and perturbation—they use Hagedorn and Manovskii (2008) as an example—creates a mean computational error of 3.75% with the 2.5th, 50th, and 97.5th percentiles of the distribution being −11.1%, −3.66%, and 8.76%, respectively. We conclude from this that the errors resulting from the computational approach we propose lead to errors of an expectable magnitude, even though the errors we find are slightly more dispersed than what Petrosky-Nadeau and Zhang (2017) find for a representative agent search and matching model.

Figure A1.1: Computational Errors

(a) Histogram of Computational Errors (b) Scatter Plot of Errors and Labor Productivity z

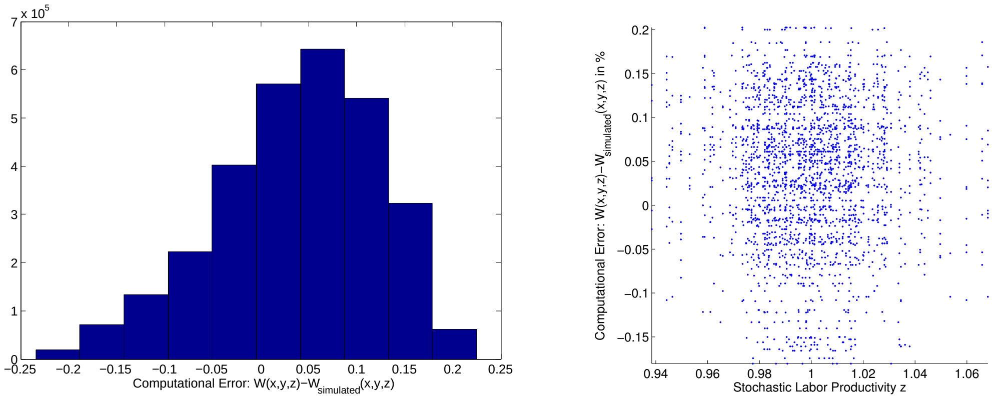

A.2 The Circular Sorting Model

Define the distance between a worker type x and a firm type y match along the circle as

$$d(x, y) = \min\{x - y + 1, y - x\}.$$

Without loss of generality, we consider only the case $x > y$ because $y > x$ is completely symmetric. The maximum distance is $1/2$ since it starts decreasing again halfway around the circle. The circular production function $F(d)$ maps d onto output with an interior maximum at $F(d = 0)$:

$$F(d) = \bar{F} - \frac{1}{2}\gamma d. \tag{1.30}$$

The functional form we consider is comparable to that of Marimon and Zilibotti (1999), Gautier et al. (2010), and Gautier and Teulings (2015).[50] $\bar{F} > 0$ is the output of an optimal match with $d = 0$. The distance $d \in (0, \frac{1}{2})$ is a measure of mismatch between workers and firms. γ governs how quickly output is decreasing in distance. It represents the strength of the production complementarity and the cost of mismatch.[51] The smaller γ, the more substitutable are different types in production.[52] We consider only interior

[50] Marimon and Zilibotti (1999) have a non-negative lower bound, which I do not need for my purposes. Gautier et al. (2010) and Gautier and Teulings (2015) let the output decrease quadratically in distance. This difference only affects the calibration of the model.

[51] γ is also known as the "complexity dispersion parameter" (Teulings and Gautier, 2004; Teulings, 2005).

[52] The limiting case $\gamma = 0$ would be a labor market without worker and firm heterogeneity. With $\gamma = \infty$, agents would match only with their optimal partner and nobody else.

solutions in which matches are unacceptable beyond a certain distance, so γ has to be sufficiently high to ensure that, in equilibrium, workers do not accept jobs at every distance.

The circular sorting model has two properties that vastly simplify it compared to the hierarchical version. First, it follows from Lemma 1 in Marimon and Zilibotti (1999) that if unemployed workers x are distributed uniformly along the circle initially, the distribution of vacancies y will also be uniform. This implies, second, that the values of unemployment, $\mathcal{U}$, and vacancies, $\mathcal{V}$, are the same for all worker and firm types. The value of non-market activity, b, and the vacancy posting costs, c, are constants and not type dependent. The values of employment and production, $\mathcal{E}(d)$ and $\mathcal{P}(d)$, depend on d only and not on the underlying worker and firm types. The same is true for the Nash wage bargain $W(d)$. These properties highlight why the circular version is a comparative advantage sorting model: the underlying worker and firm types do not matter for production by themselves; there is no hierarchy. We solve the circular model to obtain a closed-form expression for the interior cutoff distance, which we call d^*. In an interior solution ($d^* < \frac{1}{2}$), workers and firms will accept matches only up to a distance of d^* and reject matches beyond that point. This conforms with the matching set logic of the more general hierarchical model.

The following circular sorting model is a discrete time version of Marimon and Zilibotti (1999). The steady-state value functions are standard expect for the integral term in the value of unemployment and a vacant firm. Unmatched agents will accept matches up to d^* both to the left and to the right of their own position on the circle. For this reason, the integral terms are multiplied by 2.[53]

$$\mathcal{E}(d) = W(d) + \beta(\delta\mathcal{U} + (1-\delta)\mathcal{E}(d)) \tag{1.31}$$

$$\mathcal{U} = b + 2\beta q_u(\theta) \int_0^{d^*} (\mathcal{E}(c) - \mathcal{U})\mathrm{d}c \tag{1.32}$$

$$\mathcal{P}(d) = F(d) - W(d) + \beta(\delta\mathcal{V} + (1-\delta)\mathcal{P}(d)) \tag{1.33}$$

$$\mathcal{V} = -c + 2\beta q_v(\theta) \int_0^{d^*} (\mathcal{P}(c) - \mathcal{V})\mathrm{d}c \tag{1.34}$$

[53]We use c as the variable of integration to avoid notational confusion.

Due to the free entry assumption, $\mathcal{V}$ is 0 in equilibrium, so $c = 2\beta q_v(\theta) \int_0^{d^*} (\mathcal{P}(c))\mathrm{d}c$. The standard Nash bargaining solution then implies

$$\mathcal{E}(d) - \mathcal{U} = \frac{\alpha}{1-\alpha}\mathcal{P}(d), \tag{1.35}$$

where α is the workers' bargaining power. By plugging in (1.31), (1.32), and (1.33) and using the fact that $\frac{c}{q_v(\theta)}$ must be equal to $\mathcal{P}(d)$ with free entry, We obtain the following wage equation

$$W(d) = \alpha(F(d) + c\theta) + (1-\alpha)b, \tag{1.36}$$

which is just the Pissarides (2000) textbook wage equation with the additional twist that wage and output depend on the distance d in the circular sorting model.

Stationary Equilibrium of the Circular Model

The main advantage of the circular model from a computational perspective is that it delivers a closed-form solution for the matching cutoff d^*. We find the cutoff by using the fact that the value of a producing firm must be 0 at the cutoff in an interior solution.

$$\mathcal{P}(d^*) = (1-\alpha)(F(d^*) - b) - \alpha c\theta = 0 \tag{1.37}$$

Equation (1.37) can easily be solved for d^*. To find the value of aggregate labor market tightness θ compatible with d^* and free entry, plug the value functions into Equation 1.35, substitute out the wage using $\mathcal{P}(d)$, integrate both sides, and substitute the integral term out of the entry condition. The equilibrium θ value is unique because we consider interior solutions only.

The Dynamics of the Circular Model

As a final computational test, we run simulations of the simple circular sorting model. The model has a closed-form solution for the matching cutoff $d^* \in (0, \frac{1}{2})$, the maximum distance along the circle workers and firms are willing to accept when forming a match. Due to this closed-form solution, solving and simulating the circular sorting model is much simpler than the hierarchical version. The matching cutoff is procyclical in the circular model, so the acceptable distance from the optimal match increases in response to a positive shock. This channel amplifies job creation in a manner similar to that found with the hierarchical model. The question is whether the simpler circular model also produces sufficient amplification or if the additional complexity of the hierarchical sorting model is necessary to match empirical labor market dynamics

We simulate the model using the Shimer (2005) calibration, the value of home production b and vacancy posting costs c are constants. We calibrate the parameters of the circular production function (Equation (1.30)) in order to match the output dispersion of the hierarchical production function. We find that the circular model produces a small amount of amplification. The standard deviation of simulated labor market tightness is about 0.052. This is a small increase in comparison to Shimer (2005) (0.035) but still far from the empirical moment to be matched (0.382). We conclude from this exercise that the additional features of the more complex hierarchical sorting model, particularly the endogenously adjusting distributions in the state space, are important for explaining the observed cyclical dynamics of the labor market.

Chapter 2

Labor Market Sorting in Germany*

*This research is joint work with Benjamin Lochner, University of Erlangen-Nuremberg and Institute for Employment Research (IAB). We thank Michele Battisti, Anja Bauer, Stéphane Bonhomme, Bernd Fitzenberger, Timo Hener, Christian Holzner, Philipp Kircher, Patrick Kline, Thibaut Lamadon, Rasmus Lentz, Jonas Maibom, Christian Merkl, Helmut Rainer, Jean-Marc Robin, Johannes Schmieder, Heiko Stüber, Coen Teulings, and Rune Vejlin for their comments and suggestions. Feedback from seminar and conference participants at the Universities of Aarhus, Cambridge, Nuremberg, and Regensburg, at the Ifo Institute and the Institute for Employment Research (IAB), and at SMYE 2016 (Lisbon), SaM 2016 (Amsterdam), ES-NASM 2016 (Philadelphia), SED 2016 (Toulouse), and the conference of the DFG priority programme 1764 has greatly improved this work.

2.1 Introduction

Increasing wage inequality is a topic of high interest for both policymakers and academics. For Germany, Dustmann et al. (2009) show that between 1990 and 2000 real wage growth for full-time working men was negative below the 18th percentile of the wage distribution and increasingly positive above.[54] Using by now widely available matched employer-employee data, Abowd et al. (1999) (henceforth AKM) first demonstrate how to quantify the respective contributions of unobserved worker and firm heterogeneity to wage dispersion. AKM models provide a very good fit and typically find that most of the observed wage dispersion is explained by unobserved heterogeneity. Card et al. (2013) (henceforth CHK) apply the AKM methodology to the universe of German social security records.[55] They decompose the increase of wage dispersion into three elements: rising worker heterogeneity, more dispersion of wage premiums paid by employers, and increased sorting based on unobservables, measured by a rising correlation of worker and firm-fixed effects.

This Chapter is motivated by the discussion about the potential misspecification of the AKM two-way fixed effect model. As pointed out by, among others, Gautier and Teulings (2006), Eeckhout and Kircher (2011), and Lopes de Melo (2016), the assumption that log wages are additively separable into worker and firm-fixed effects leaves, by construction, no role for match-specific effects.[56] In theoretical sorting models, however, output and wages are match-specific; they are determined by a production complementarity between workers' skill and firms' productivity types. If this match-specific component of wages was quantitatively important, the AKM model would be inapplicable to quantify wage dispersion. CHK argue that the abstraction from match-specific effects is defensible

[54]Remarkably, the yearly increase of wage dispersion in Germany has an order of magnitude comparable to the United States; this trend can be traced back well into the 1980s/70s. Dustmann et al. (2009) find that the gap between the 85th and 50th percentiles of the German wage distribution has increased by about 0.6 log points per year between 1975 and 2004. This is comparable to Autor et al. (2006), who report that the gap between the 90th and 50th percentile in the United States has increased at a rate of roughly 1 log point per year during the same period.

[55]Having access to the universe of data is crucial for the AKM methodology because the observation of all workers per firm reduces the potential impact of the so-called "limited mobility bias", a problem emphasized by Andrews et al. (2008, 2012).

[56]In the remainder of this Chapter, we will sometimes refer to this assumption as the "AKM assumption".

because deviations in terms of wage residuals appear to be small for most, but not all, combinations of worker and firm types.

We build a bridge between these two opposing views on the sources of wage dispersion. We analyze German matched employer-employee data[57] through the lens of a structural sorting model featuring heterogeneity on both sides of the market, search frictions, and search on the job. Our main analytical tool is the identification technique proposed by Hagedorn et al. (2017) (henceforth HLM). To identify the sign and strength of sorting, they estimate economy-wide rankings of both workers and firms using data on wages and labor market transitions only. In their model, the match-specific output is non-parametrically identified. HLM report a rank correlation coefficient, a natural measure for sorting, of 0.76. The respective correlation of worker and firm-fixed effects in CHK is much lower, about 0.21 for a comparable period.[58] First, we make a methodological contribution by using the German labor market as a laboratory to reconcile these seemingly incompatible approaches. Using the worker ranking procedure proposed by HLM and a firm ranking which does not depend on wages, we confirm that matching patterns are positive assortative in Germany. We find an overall rank correlation coefficient of 0.24, which is closer to the CHK-AKM benchmark than to HLM.

Our second contribution is a thorough quantitative investigation of labor market sorting in Germany. We find that sorting increased significantly over time. It rose particularly for matches formed by workers out of unemployment. Low-skill workers are increasingly sorted into low-productivity firms. This development can be linked to increasing domestic outsourcing of firms.[59] Sorting of high-skill workers into high-productivity firms has, if anything, slightly decreased during our period of observation. Interestingly, the increased sorting of workers at the bottom of the type distribution is related to non-monotonic wage patterns. These workers maximize their wages in matches with low-type firms and the increased sorting leads to wage gains. We provide direct empirical evidence that for some

[57] German social security register data are provided by the Institute for Employment Research (Institut für Arbeitsmarkt- und Berufsforschung, IAB) of the German Federal Employment Agency (Bundesagentur für Arbeit).

[58] For simplicity, this is the arithmetic mean of the correlations reported by CHK for the two subperiods which span the same period as HLM.

[59] Goldschmidt and Schmieder (2015) analyze domestic outsourcing in Germany. They use an AKM model and find that displaced workers suffer wage losses. Using our more flexible framework, we observe that the wage of low-type workers increases on average when they move to firms that offer services which are typically outsourced (cleaning, security, temporary work agencies).

worker types wages are not monotonically increasing in the productivity-type of the firm they are matched with. This non-monotonicity of wages is a key feature of theoretical sorting models like Shimer and Smith (2000), Atakan (2006), and Eeckhout and Kircher (2011), but it is at odds with the AKM two-way fixed effect approach, because in the reduced-form model wages mechanically increase in the firm effect. Accordingly, the residuals of the CHK-AKM model are substantial for some combinations of worker and firm types, specifically those including low-type workers.[60] We find that increased sorting and the non-monotonicity of wages are most pronounced for these high-residual workers at the bottom of the type distribution. For medium and high-type workers, however, wages increase monotonically in the firm type, consistent with CHK and AKM. This explains why two-way fixed effect models generate a good fit overall. Nevertheless, it is important to emphasize that the AKM model cannot explain wage determination for low-type workers. The outcomes for this group of employees are the ones with which labor market policy is typically most concerned.

In the related literature, a number of recent papers explores ways to identify the sign and strength of sorting and its contribution to wage dispersion without relying on the AKM assumption. Bagger and Lentz (2016), Gautier and Teulings (2015), and Lise et al. (2016) develop structural search models and take them to the data. Bonhomme et al. (2016) propose a flexible empirical framework with a discrete number of types (finite mixture model) which also allows for unrestricted interactions of worker and firm heterogeneity. Our analysis is inspired by HLM, which also belongs to this group of papers. HLM show how match-specific output, wages, and the sign and strength of sorting are non-parametrically identified with standard matched employer-employee data available for many countries. The key challenge for the empirical study of sorting is to construct credible global rankings of workers and firms. While studies in the AKM tradition implicitly rank workers and firms by their fixed effects, HLM solve a Kemeny-Young rank aggregation problem to merge intra-firm wage rankings into a global ranking of workers. This procedure uses the largest connected set of workers. A computational algorithm running on this graph can effectively maximize the likelihood of the correct

[60]In the AKM context, worker and firm types are equivalent to their estimated fixed effects. Regarding residuals in CHK, see Figure VI on p. 996 and the discussion on pp. 989-991 in Card et al. (2013).

global worker ranking, as proven by Kenyon-Mathieu and Schudy (2007). Once workers are globally ranked, HLM show how to use the model structure to rank firms based on the value of a job vacancy. This approach delivers a rank correlation which lies considerably above the correlation for Germany estimated by CHK and above related results for other countries in the studies mentioned before. The German case is well-suited to take the HLM approach to a test. We have access to very detailed information for a large number of firms from the IAB Establishment Panel. If all information necessary to rank both workers and firms globally was indeed contained in wages alone, additional firm data would be redundant and not alter the results in any meaningful way. We construct an efficiency-based firm ranking using the distance to an estimated production frontier. Our ranking is independent of wages and produces lower rank correlations, which are more in line with the empirical sorting literature and with CHK for Germany in particular. We also construct a profit-based firm ranking inspired by Bartolucci et al. (2015), which yields similar results.

A closely related paper is Kantenga and Law (2016). They use the full HLM model to structurally decompose increasing overall wage dispersion into the contributions of worker heterogeneity, firm heterogeneity, search frictions, and sorting. They reject the assumption of additive separability statistically by following pairs of workers that jointly change employers. We do not have to impose the full structure of the HLM model on the data because we rank firms independently of wages. This prevents us from performing a structural decomposition of wage dispersion. We go beyond Kantenga and Law (2016), however, by providing an explanation for why additive separability is rejected in some parts of the type space. Thus, we see our results as largely complementary.

This Chapter is organized as follows: Section 2.2 describes our data. Section 2.3 briefly discusses the theory of sorting that guides our thinking. Section 2.4 explains our approach to identifying the sign and strength of sorting. Section 2.5 presents our results: correlations of our estimated worker and firm rankings, the dynamics of sorting over time, and how this relates to wages, wage inequality, and increased domestic outsourcing. Section 2.6 concludes and selected additional results and robustness tests can be found in the Appendix.

2.2 Data

We use matched employer-employee data for Germany provided by the Institute for Employment Research (Institut für Arbeitsmarkt- und Berufsforschung, IAB) of the German Federal Employment Agency (Bundesagentur für Arbeit). We use the "LIAB Mover Model"[61] and restrict our analysis to the years 1998-2008.[62] This data set is ideal for our purposes as it provides information about a large number of workers moving between firms and it can be linked directly to firm-level survey data from the IAB Establishment Panel. We provide a short overview of our data preparation procedures in the following; more details on sample selection and imputation procedures can be found in Appendix A.1.

2.2.1 Data Preparation

One major advantage of German matched employer-employee data is the high quality of the wage data due to plausibility checks carried out by the social security institutions followed by sanctions for non- and misreporting. In our raw data, we observe nominal gross daily wages, which we deflate using the consumer price index from German national accounts with 2005 as the base year. Every wage observation corresponds to one employment spell, which can last from one day up to one year due to the reporting rules of the German social security system.[63] On average, the workers in our sample have 6.4 job spells between 1998 and 2008 (with a standard deviation of 5.0) and the average spell lasts 289 days (with a standard deviation of 114). We drop workers with more than 150 job spells (less than 0.01%). Observed spells are sometimes identical except for the ending date and the corresponding wage, which could be the result of multiple reports in case of a changing contract end date. In such cases the longer spell is typically associated with a higher wage, possibly due to Christmas bonuses or other salary supplements. We always keep the spell with the higher wage.

[61] File: `LIAB_MM_9308`. See Heining et al. (2012) for a detailed description of the data set.

[62] We chose this period of time because it is roughly split in half by the German labor market reforms passed and implemented between 2002 and 2005 (Hartz I-IV).

[63] The reporting rules require employers to file a report whenever an employee joins or leaves the establishment or, in the event of no change in an ongoing employment relationship, on December 31 each year.

The education variable in German social security data suffers from missing values and inconsistencies, essentially because misreporting has no negative consequences. We impute missing and inconsistent observations using the methodology proposed in Fitzenberger et al. (2006). Missing values cannot be imputed for about 2% of the data and we drop these spells. A second limitation of the wage data is that the German social security system tracks earnings only up to a certain threshold, the contribution assessment ceiling ("Beitragsbemessungsgrenze").[64] We follow the procedure suggested by Dustmann et al. (2009) and impute the upper tail of the wage distribution by running a series of Tobit regressions, allowing for a maximum degree of heterogeneity by fitting the model separately for years, education levels, and eight five-year age groups.[65]

We restrict our data set to West German male employees between 20 and 60 years of age who were liable for social security contributions. Self-employed workers, civil servants, and students are not included in the sample. Although it would be very interesting to analyze the sign and strength of sorting for the excluded subgroups, we put a higher priority on using a dataset which is, on the one hand, homogeneous and, on the other, comparable to data used in previous studies of wage dispersion in Germany, CHK in particular. To minimize working-time effects in our wage data, we further exclude part-time and marginal employment.[66]

After the initial data preparation, our sample consists of 16,361,068 employment spells, 1,824,580 workers, and 472,869 establishments for the years 1998-2008. Note that the employers we observe are not firms in the legal sense, but establishments or local production units. These do not necessarily coincide with the legal entity to which they belong. We use the terms "firm", "establishment", and "employer" interchangeably. The standard deviation of log wages in this sample is 0.455. This value is close to the measures

[64]The average yearly censoring rate is 13.6% of wage observations. We define a wage observation as censored whenever the reported wage is higher than 99% of the censoring threshold.

[65]An alternative to imputing the censored part of the wage distribution would be to simply drop top-coded wages. Although we would lose the ability to analyze sorting patterns of workers earning very high wages, our main findings would not be affected as they pertain to workers with low and hence non-censored wages. Moreover, the calculation of residual wages, the main input for our analysis, is largely unaffected by the imputation: Table A2.2a shows that a wage variance decomposition delivers very similar results with and without the imputed part of the wage distribution.

[66]To reliably identify spells of marginal or part-time employment we use both indicator variables in the data and, additionally, drop spells with wages below the time-varying marginal employment threshold, which is on average 12.2 Euro per day across years in our sample ("Geringfügigkeitsgrenze").

of wage dispersion reported by CHK, indicating that our data preparation procedure indeed generates a comparable sample.[67]

2.2.2 Residual Wages

We analyze how unobserved worker and firm heterogeneity as well as their potential interaction influence wages and the selection of workers into jobs, i.e sorting. Therefore, we use residual wages—wages net of the effects of observable worker characteristics—as the main input to rank workers. To compute them, we use a simple empirical model of wage dispersion of the AKM-type and regress, following CHK, log wages on a person-fixed effect, an unrestricted set of year dummies, and quadratic and cubic terms in age fully interacted with educational attainment:

$$\ln w_{it} = x'_{it}\gamma + \alpha_i + r_{it}. \tag{2.1}$$

$\ln w_{it}$ denotes the log real daily wage of a worker i in year t, x'_{it} includes the time-varying observable characteristics, α_i is a worker-fixed effect, and r_{it} is the error term. Typical of an AKM-type wage regression, the explanatory power of this model is very high. The adjusted R^2 is 81%, slightly below the around 90% reported by CHK. This difference is easily explained by the lack of an establishment-fixed effect in our specification, which is, however, inconsequential for our ranking purposes.[68] Table 2.1 shows the decomposition of the variance of log wages. The unobservable components of wages explain the vast majority of wage variation in our data, namely, 91%. The person-fixed effect alone explains 74% of log wage variance; the residual absorbs another 17%. To compute the residual wages, we simply subtract the share of the wage predicted by observable characteristics

[67] CHK report a standard deviation of log wages of 0.432 for 1996-2002 and 0.499 for 2002-2009.

[68] In contrast to CHK, who have access to the universe of German social security records, we cannot reliably estimate establishment-fixed effects because we do not observe the full workforce of the establishments in our sample. This is not a problem for our analysis. The ranking procedure proposed by HLM relies on pairwise wage comparisons of workers who are employed at the same establishment (see Section 2.4.2). The unobserved firm effect affects the wages of two workers at the same establishment by exactly the same amount. The ranking of a pair of workers employed by the same firm is thus not affected by the firm effect.

Table 2.1: Decomposition of the Variance of Log Wages

	$\ln w_{it}$	$x'_{it}\hat{\gamma}$	$\hat{\alpha}_i$	$\hat{r}_{it}$
$\ln w_{it}$	0.207			
$x'_{it}\hat{\gamma}$	0.014	0.008		
$\hat{\alpha}_i$	0.158	0.006	0.152	
$\hat{r}_{it}$	0.035	0.000	0.000	0.035

Notes: Variance-Covariance matrix of regression model 2.1. The variance of log wages ($\ln w_{it}$) is decomposed into the variance of observable characteristics ($x'_{it}\hat{\gamma}$), the person-fixed effect ($\hat{\alpha}_i$), and the residual ($\hat{r}_{it}$). Rounded to three decimal places.

from the observed wage for every individual:

$$\ln \tilde{w}_{it} = \ln w_{it} - x'_{it}\hat{\gamma}. \tag{2.2}$$

The log residual wage has a standard deviation of 0.433 (variance 0.187). It is thus only slightly less dispersed than observed wages in our sample, underlining the small role that observable characteristics play for wage dispersion. The correlation between observed and residual wages is very high (0.98).

To address concerns about the influence of occupations on wages which we do no account for in the specification above, we augment Equation (2.1) by controlling for 32 occupational categories.[69] We interact them with education and year dummies. Reassuringly, controlling for occupations does not significantly increase the explanatory power of the wage regression. The adjusted R^2 stays virtually the same (81%) and the portion of variance explained by observable characteristics rises only slightly from 3.9% in the baseline regression to 6.7% when including occupational controls.[70] The correlation between baseline residual wages and residual wages net of occupational effects is very high, above 0.99. Taking into account the workers' occupation to calculate residual wages would not

[69]In accordance with Bundesagentur für Arbeit (1988) the classification in our data consists of a about 330 occupational codes on the 3-digit level. We use 32 2-digit values ("Berufsabschnitte") in order to being able to estimate a large set of interaction terms.

[70]Table A2.2b shows the decomposition of wage variance including occupational controls.

lead to a significantly different worker ranking and hence not affect our results. We are confident that the residual wages $\tilde{w}_{it}$ are the appropriate input for our analysis.[71]

2.2.3 Firm Data

We use the IAB Establishment Panel to calculate measures of firm performance which we develop and use in Section 2.4.3. The IAB Establishment Panel draws a stratified random sample of establishments from the register data.[72] We restrict our attention to establishments that employ at least 10 workers on average during our period of interest.[73] The data contains no direct information about the capital stock on which an establishment operates. To circumvent this shortcoming, we use the perpetual inventory method proposed by Müller (2008) to approximate the capital stock of the establishments in our sample. This method uses information about the average economic lives of different capital goods (buildings, IT, production machinery, transport equipment), which is available from national accounts, and the firms' net investments in the different categories of capital goods. We exclude the public sector of the economy and firms which do not report revenues as their primary measure of output.[74] Finally, we merge our firm sample with the data on employment spells. This leaves us with 4,901 establishments for which we observe 1,714,450 employment spells of 234,800 workers.

2.2.4 Matches

The main unit of our empirical analysis in the following is a match between a worker and a firm. For our study of the allocation of workers to jobs, we focus on the first person-year observation of a new employment spell. Hence, we abstract from the spell length. This ensures that certain match types are not over-represented (under-represented) in the distribution of matches simply because they last longer (shorter) on average, leading

[71]CHK provide a comprehensive analysis of the explanatory content of additional controls in their setting. In line with our finding, they report that occupational (and industry) controls do not significantly increase the explanatory power of AKM-type models of wage dispersion.

[72]For details on the IAB Establishment Panel see Kölling (2000) and Fischer et al. (2009).

[73]We lose only 0.2% of employment spells due to this restriction.

[74]The public sector includes health care, education, non-profit organizations. We exclude these firms because they typically do not seek to maximize profits or minimize costs, what is at odds with the rationale of our firm ranking procedure. The most important group of establishments which do not report revenues are banks.

to more (less) person-firm-year observations. This is the most conservative way to study sorting. Technically speaking, we define a new match as the first worker-id establishment-id combination we observe in the data. We further distinguish between two types of matches: matches out of unemployment and matches resulting from job-to-job switches of an employee. A match out of unemployment occurs whenever a particular worker is employed after a period of registered unemployment or an uncompensated time gap between two consecutive jobs which is longer than 1 month. Job-to-job matches are defined as job switches with no time gap or a time gap smaller than 1 month. Whenever a worker and an establishment match twice, we only consider the first match and exclude job recalls.

The worker ranking procedure we apply relies on the HLM model with search on the job.[75] In this framework, identification of the worker ranking is based on wage information from matches formed by workers out of unemployment only. These wages can be shown to be monotonically increasing in the unobserved worker type in the context of the model. As a result, we lose all workers in our sample who are never unemployed and only switch jobs. Workers unemployed at some point in time, however, can be ranked and we follow them as they switch jobs.[76] Our final sample consists of 183,156 matches, 75,831 arise out of unemployment and 107,325 from workers switching jobs.

2.3 Sorting Theory

To support our analysis of labor market sorting in Germany, this section briefly summarizes the theory of sorting that guides our thinking about the allocation of workers to jobs and the determination of wages. Thinking about sorting conceptually requires heterogeneity of agents in a two-sided market. Since this is research in the context of the German labor market, we denote our heterogeneous types *workers* and *firms.* Firms are heterogeneous in terms of their productivity and workers have different skill levels.

[75] Details are discussed in Section 2.4.2.

[76] All our conclusions are unaffected by this choice. Ranking workers based on all their wages, not just out of unemployment, leads to a very similar worker ranking with a correlation of above 0.99. In this case, the final sample contains about three times as many matches. Even though the difference between the two rankings is very small, we use the ranking out of unemployment for the sake of methodological cleanness.

Workers and firms carry identifiers i and k, respectively. We assume for now that unambiguous rankings of both worker and firm types exist and are observable to the researcher. In practice, rankings are unobservable. Establishing these rankings empirically is thus the key challenge in measuring the sign and strength of labor market sorting.

Let the true rank of worker i be denoted $x(i)$ while the true rank of firm k is $y(k)$. For simplicity, heterogeneity across workers and firms is assumed to be one-dimensional.[77] The starting point for the theory of labor market sorting is the neoclassical optimal assignment model proposed by Becker (1973). In the absence of search frictions, every worker finds his optimal firm assignment instantaneously and matches are formed. This leads to a Walrasian first-best allocation of workers to firms. The specific nature of the optimal assignment depends on the production structure of the economy. One possible assignment pattern is positive assortative matching (PAM). In Becker's theory, a complementarity between worker and firm types gives rise to PAM. Technically, the production function, which takes worker and firm types as arguments, must be supermodular, that is, the cross-derivatives must be strictly positive. To maximize output, the complementarity then requires the most productive firm to employ the most skilled worker, the second most productive firm to employ the second most skilled worker, and so forth. The worker with the lowest skill level is optimally matched with the least productive firm. This allocation implies a perfect positive correlation of worker and firm ranks: $\rho = \frac{\text{Cov}[x(i),y(k)]}{\sigma_{x(i)}\sigma_{y(k)}} = 1$, where ρ is Spearman's rank correlation coefficient, computed by dividing the covariance of the rank variables by the product of their standard deviations.[78] Spearman's ρ is a natural measure for the sign and strength of sorting. A negative rank correlation, $\rho < 0$, implies negative assortative matching (NAM) with a submodular production function. NAM features high-productivity firms employing low-skilled workers (and vice versa) to maximize output. In turn, $\rho = 0$ indicates the absence of sorting patterns in allocating workers to jobs, corresponding to a modular production technology.

A major step forward for the theory of labor market sorting has been the incorporation of search frictions into the Beckerian model. Shimer and Smith (2000) show existence and

[77]Of course, this is a simplifying assumption. However, it is not obvious ex-ante that having more than one dimension of heterogeneity would meaningfully increase the explanatory power of the model. For some recent explorations of the empirical content of multi-dimensional sorting models see Lise and Postel-Vinay (2016) and Lindenlaub and Postel-Vinay (2016).

[78]Spearman's ρ is equal to a standard Pearson correlation coefficient applied to the rank variables.

Figure 2.1: A Simple Sorting Model

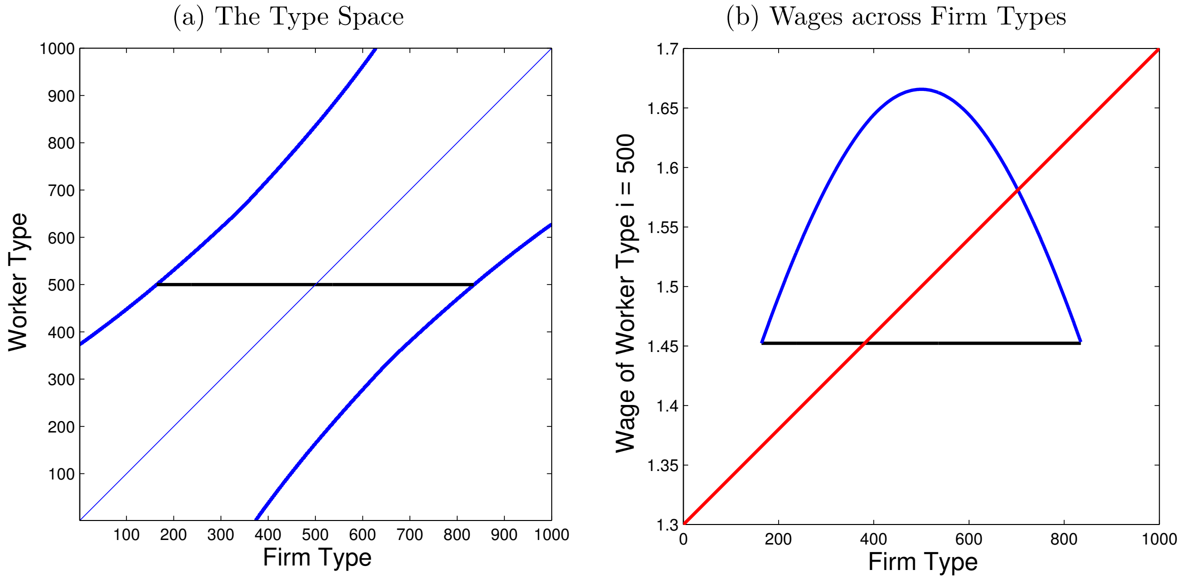

Note: The blue curves depict the equilibrium of a simple Shimer and Smith (2000) economy. The black line represents the equilibrium matching set of a worker of type 500. The model is a simplified version of the structural search model presented in HLM. It is solved numerically using standard parameter values. For the derivation and solution procedure see also Chapter 1. The red line in Panel (b) is a simple stylized representation of a monotonic relation between worker's wages and the type of the firm as it is assumed in two-way fixed effect models of wage dispersion.

the properties of an equilibrium in a two-sided assignment model with frictions. With a continuum of types and random search, frictions imply that the first-best allocation cannot be realized because the probability of meeting any specific partner type on the other side of the market is zero. Therefore, a subset of firm and worker types in the vicinity of the optimal allocation must be acceptable. These subsets of types, the matching sets, are determined by the option values of a all potential matches.[79] Figure 2.1 depicts the equilibrium of a simple Shimer and Smith (2000) economy. This illustrative economy is populated by a population of 1,000 discrete and heterogeneous worker and firm types, each of them having an equal mass normalized to one. Since we abstract from the ranking problem for now, the indices $(i, k) \in \{1, ..., 1000\}$ are directly interpretable as the workers' and firms' ranks. The production function exhibits a complementarity which induces PAM as explained above. Panel (a) of Figure 2.1 shows the space of all worker and firm type combinations, the type space. The 45-degree line represents the Beckerian first best allocation. The blue curves depict the cutoffs of the matching sets of all worker and firm types. The black line is the matching set of one specific worker type, $i = 500$. This worker type is willing to match with, roughly, all firm types above 180 and

[79]Matches are acceptable if their option value is higher than the value of continued search.

below 820. Within these boundaries, the surplus is (weakly) positive and matches are formed. Matching with firm types outside of the boundaries, in turn, is precluded due to a negative surplus. According to the model, no matches are formed in the upper-left and lower-right corners of the type space, indicating the absence of matches between high-type workers and a low-type firms as well as low-type workers and high-type firms.

Panel (b) shows how the model-generated wage of the exemplary worker type, $i = 500$, varies across firm types. In the model, wages are determined by splitting the match surplus according to the Nash bargaining solution. Note that due to the production complementarity in this economy, the wage of a given worker type is non-monotonically related to the firm type. For worker type $i = 500$, it is maximized in a match with firm $k = 500$ (first best allocation). Apart from the wage-maximizing allocation, the worker has to accept lower wages because the production complementarity is not fully exploited. The worker needs to compensate the firm for the foregone option value of waiting for a better hire. This non-monotonicity lies at the heart of the critique of the AKM two-way fixed effect approach (Gautier and Teulings, 2006; Eeckhout and Kircher, 2011; Lopes de Melo, 2016). In the log-linear AKM model, the wage must monotonically increase in the estimated firm-fixed effect, which is simply being equated with the unobservable firm type. Such a monotonic relation between wages and firm type might look like the stylized red line in Panel (b) of Figure 2.1. The higher the firm type, the higher is the wage of every worker the firm employs, a simple wage premium. The theoretically predicted non-monotonicity is clearly at odds with the AKM model, which assumes monotonicity. This potential misspecification of the AKM model led to a number of alternative strategies to quantify the contributions of worker and firm heterogeneity as well as their potential interaction to overall wage dispersion. The HLM identification strategy, which we rely on to a significant extent, makes direct use of the model structure described in this section. We discuss it in relation to alternative identification strategies in the next Section.

2.4 Identifying Sorting

We identify the sign and strength of sorting in the labor market by combining the wage-based worker ranking procedure (global rank aggregation) proposed by HLM with firm rankings which are constructed independently of wages. The next subsections explain our motivation for this approach and our ranking procedures in detail before we turn to the results in Section 2.5.

2.4.1 Motivation

The main contribution of HLM is to show how to construct global worker and firm rankings and identify the sign and strength of sorting using data on wages and labor market transitions only. Their identification strategy is viable in the class of structural models discussed in Section 2.3, building upon Shimer and Smith (2000). The quality of wage and transition data is typically high in commonly available matched employer-employee data sets. Detailed information about the firm, however, is not always available. It is thus an asset that the firm data requirements of the HLM approach are small. For Germany, we have access to very detailed information for a large number of firms from the IAB Establishment Panel.[80] The German case is therefore well-suited to take the HLM approach to a test: if all necessary information to rank both workers and firms globally were contained in wages, additional firm data would be redundant and not alter the results in any meaningful way. We show below that using additional firm data leads to a different firm ranking and different results.

Once workers are ranked and binned using the global rank aggregation algorithm (see details below), the value of unemployment for a specific worker type can be estimated by interpreting the lowest observed wage of all workers of the same type as this type's reservation wage. Subsequently, HLM use the model structure, primarily Nash bargaining, to show that firms can be ranked based on the value of a vacant job.[81] It can be constructed as a statistic of wages and transition rates. The value of a vacant job is shown to be monotonically increasing in the firm type and it is identified from the observed differences of wages paid to different workers of the same type and their estimated reservation wage. This wage premium is proportional to the expected surplus of a vacancy due to the

[80] HLM use data from the IAB Establishment Panel as well. Vacancy data at the firm level greatly simplify the computation but is not necessary for their approach to work in principle.

[81] Technically, Nash bargaining is not specifically required. It is sufficient that both parties benefit from an increase of the surplus, the proportions do not have to be fixed.

assumption of Nash bargaining.[82] Thus, the bargaining assumption is crucial to identify the value of a vacancy using wage data. Once the worker-specific value of unemployment and the firm-specific value of a vacancy are known, the model structure allows HLM to simply invert the wage equation for match-specific output. Thus, the production function is identified non-parametrically at the match level in this setting. As a consequence, it is possible to analyze the sign and strength of sorting without making functional form assumptions about the production function.

The assumption that wages are determined by contemporaneous bargaining made by HLM is critical for identification as it enables them to rank firms based on wages and invert the wage equation for match-specific output. Empirically, the assumption that wages are determined by splitting the surplus of a match period-by-period is backed up by evidence presented in Hagedorn and Manovskii (2013). For the U.S., they show that wages are robustly not history dependent but driven by current aggregate labor market conditions and idiosyncratic match-specific productivities, consistent with period-by-period bargaining.[83] For Germany, new evidence suggests that history dependence is, in contrast to the U.S., a prevalent feature of wage determination. Bauer and Lochner (2016) use the on-the-job-search model proposed by Hagedorn and Manovskii (2013) and estimate it on German social security data (SIAB-7514). Following Hagedorn and Manovskii (2013), they explicitly control for unobserved match-quality and even allow for heterogeneity across occupations. Unlike in the U.S., wages turn out to depend on the initial labor market conditions or on the best outside offer a worker receives while being in a specific job.[84] The influence of current labor market conditions is thus smaller than in the U.S. This suggests that the bargaining assumption necessary to rank firms and invert the wage equation might be less appropriate for Germany than it would be in the U.S. context.

[82] See Hagedorn et al. (2017) p. 15 ff. for more details and proofs.

[83] In contrast, earlier studies for the U.S., most notably Beaudry and DiNardo (1991), find history dependence in wages and interpret it as evidence for infrequent bargaining or implicit contracts. This result disappears after controlling for unobserved match quality as Hagedorn and Manovskii (2013) show.

[84] The labor market conditions at a specific point in time are measured by the unemployment rate. Empirically, these papers use the unemployment rate at the begin of an employment spell to control for initial conditions. The likelihood of receiving outside offers, which the firm may match to retain the worker, is captured by the lowest measured unemployment rate during an employment spell.

The baseline results in HLM are derived using an extended version of the model with search on the job. As in Postel-Vinay and Robin (2002) and Cahuc et al. (2006), in case an employed worker meets another firm, the firms engage in Bertrand competition and depending on the potential surplus at the new firm the worker might switch jobs. HLM show that their constructive proof of identification extends to this richer framework under the assumption that unemployed workers extract the full surplus of the match when they are hired. In this case, the wage is monotonically increasing in the unobserved worker type, the key requirement to rank workers. A disadvantage of identifying the model from out of unemployment wages only is that workers who are never observed in an unemployment spell cannot be ranked.[85]

Using German data, HLM find a very high correlation of estimated worker and firm ranks, it is 0.76 for the years 1993-2007. This value suggests a high degree of PAM and severe misspecification of AKM models. It is much higher than the correlation of estimated worker and firm-fixed effects reported by CHK using the AKM methodology on German data. CHK report correlations of 0.17 (1996-2002) and 0.25 (2002-2009).[86] The HLM result also lies considerably above rank correlations reported in related studies using alternative identification strategies and data from other countries. Bagger and Lentz (2016) use Danish data and estimate a structural model with endogenous search intensity and on the job search. Job-switchers are used to rank firms by their poaching rank. They report a correlation of only 0.11. Bonhomme et al. (2016) propose a clustering technique (finite mixture model) to identify a discrete number of firm types based on the similarity of their within-firm wage distributions. This very flexible empirical model allows for unrestricted interactions between worker and firm heterogeneity. Using Swedish data, they find correlations around 0.44. Bartolucci et al. (2015) rely on high-quality Italian firm data (balance sheets). They develop a refined reduced-form approach to identify the sign and strength of sorting and find a correlation of worker and firm types of 0.52. Finally, evidence for the U.S. is presented by Lise et al. (2016). They make a parametric assumption

[85] We find that ranking workers based on all their wages, not just out of unemployment, leads to a very similar worker ranking with a correlation of above 0.99 for the workers in both samples. However, we stick to the ranking out of unemployment for the sake of methodological cleanness.

[86] Recall that CHK have access to the universe of German social security records. Using the AKM model on those data does not suffer from the limited mobility bias emphasized by Andrews et al. (2008, 2012).

about the production function (CES) and directly estimate the elasticity of substitution. They also find evidence for PAM, but the magnitude of estimated substitution elasticity is not readily comparable to a rank correlation coefficient.

The sign of sorting found in all these studies is unambiguously positive, suggesting that some degree of PAM is indeed a prevalent feature of labor markets in developed economies. The variation of estimated correlation coefficients, however, is huge.[87] Using our method, we shed light on the large spread of estimated correlation coefficients for Germany. As mentioned before, the benchmark studies are CHK using the AKM model (0.17-0.25) and HLM using a structural framework (0.76). Both studies use German social security records and comparable time periods. Different correlation coefficients must therefore stem from methodological differences. Using additional firm information from the IAB Establishment Panel is instructive because we can construct firm rankings independently of wages and thus test whether the HLM approach of ranking both workers and firms on wage data alone is a reliable way of capturing both worker and firm heterogeneity.

We find confirmatory evidence for PAM in the German labor market. Interestingly, our results are much closer to the CHK-AKM benchmark than to HLM. Using the HLM worker ranking and our alternative firm ranking, we find an overall rank correlation of 0.24 (1998-2008). Thus, the specific structural assumptions made to rank firms based on wages, particularly wage bargaining, appear to severely increase the measured rank correlation. Overall, we find that the AKM model applied by CHK seems to deliver a satisfactory approximation of the data. However, we do not conclude from this observation that match-specific effects play no role for wage determination whatsoever. Low-type workers have large estimated residuals in CHK.[88] For these workers, the empirical wage profiles we observe are at odds with the monotonicity assumption. The sorting of low-type workers indeed appears to be driven by match-specific interaction effects, possibly as a result of production complementarities as suggested by theory. Applying the AKM model shrouds these effects: the monotonicity assumption may be met for large fraction

[87] To some extent the large variation of estimated correlation coefficients is probably driven by cross-country differences rather than methodological issues. Also, possible cyclical fluctuations of the degree of labor market sorting are not taken into account in the aforementioned studies. Both topics are exciting avenues for future research.

[88] CHK show residuals for type combinations, measured in worker and establishment effect deciles, in Figure VI on p. 996. In fixed effect models, the estimated fixed effects of both workers and firms are implicitly interpreted as their cardinally rankable type.

of the data, leading to the well-known high explanatory power of the model. Deviations for smaller groups of workers, however, remain undetected. Whether AKM or a more flexible model is applicable therefore depends on the specific research question: the AKM model can be a valid approximation to study the extent of wage dispersion in the labor market as a whole, jointly analyzing all worker and firm types. In a study that focuses on low-type workers, the group that labor market policy typically is most concerned with, the AKM model would fail because wage patterns are highly non-monotonic in this group. Below, we show in detail for which worker types wages across firms are indeed non-monotonic (in line with the theory depicted in Figure 2.1) or monotonically increasing in the firm type (in line with the AKM assumption of additively separable log wages). Before we turn to the results, we describe our ranking procedures in more detail.

2.4.2 Ranking Workers

HLM present a computational algorithm that merges intra-firm wage rankings into a global ranking of workers by solving a Kemeny-Young rank aggregation problem.[89] The procedure makes use of the fact that matched employer-employee data allows wage comparisons across firms because workers switch jobs: co-workers at one firm move from firm to firm over time and form a graph (or connected set) of workers with comparable wage observations. A computational algorithm running on this graph can effectively maximize the likelihood of the correct global ranking, as proven by Kenyon-Mathieu and Schudy (2007). The input of the algorithm are the workers' residual wages, $\tilde{w}_{it}$, net of observable effects as presented in Section 2.2.2. The algorithm is initialized by ranking workers according to a simple wage statistic which needs to be monotonically increasing in the unobserved worker type.[90] Using a Bayesian approach with a normal prior, HLM show how to compute the probability of worker i being ranked higher than worker j given wage

[89] Rank aggregation is an ancient problem that originated in social choice theory. Kemeny-Young rank aggregation solves this problem by minimizing the number of disagreements between potentially inconsistent rankings of voting alternatives by different voters, see Kemeny and Snell (1962). In the HLM application to the labor market, the firms are the voters and the workers the voting alternatives.

[90] HLM prove that, in the context of their model, the reservation wage, the maximum wage, and the adjusted average wage of a worker are monotonically increasing in the unobserved type. Importantly, average wages, sometimes used to rank workers in empirical applications, are not monotonically increasing in the type because they do not factor in the values of workers' interjacent unemployment spells.

histories at firm k in the presence of measurement error:

$$c(i,j) = P(\tilde{w}_{i,k} > \tilde{w}_{j,k}) = \Phi\left(\frac{\bar{\tilde{w}}_{i,k} - \bar{\tilde{w}}_{j,k}}{\frac{\sigma^2}{n_{i,k}} + \frac{\sigma^2}{n_{j,k}}}\right). \tag{2.3}$$

Φ is the standard Normal CDF. Observed (residual) wages are assumed to follow a noisy process: $\tilde{w}_{i,k,t} = \tilde{w}_{i,k} + \epsilon_t$, with σ^2 being the variance of ϵ. Intuitively, the difference of the average residual wages $\bar{\tilde{w}}_{i,k} - \bar{\tilde{w}}_{j,k}$ at firm k is weighted by the wage variance σ^2 in proportion to the number of wage observations for workers i and j at firm k, $n_{i,k}$ and $n_{j,k}$. The more available observations, the smaller is the potential impact of measurement error on the average wage of worker i at firm k and the more plausible is the ranking implied by the wage observations at this firm, resulting in a higher value of $c(i,j)$.[91] Note that σ^2 is the overall wage variance and not firm-specific because HLM make the assumption that all variation in wages for a specific job stems from measurement error only.[92] The probability $c(i,j)$ is defined for worker pairs employed at the same firm. In case a pair of workers is observed at more than one firm, the wage observations are considered to be independent and the probabilities are simply multiplied. By comparing the initial ranking with the ranking implied by the posterior probabilities $c(i,j)$, the algorithm iteratively increases the value of the following objective function and, hence, maximizes the likelihood of the global ranking:

$$\sum_{i>j} \left[c(i,j)\ \Pi(i,j) + c(j,i)\ \Pi(j,i)\right]. \tag{2.4}$$

[91]For details of the derivation of $c(i,j)$, see Appendix III.1 in Hagedorn et al. (2017).

[92]While this assumption is consistent with period-by-period wage bargaining in their model, from an empirical perspective it could be desirable to allow for heterogeneity of the within-firm wage distributions beyond the mean. Imagine a firm using different contracts to discriminate between worker types: two workers could have different slopes in their wage profile over time because tenure is remunerated differently. Such patterns could be due to history dependence, as evidenced by Bauer and Lochner (2016), or due to the coexistence of wage bargaining and wage posting, as evidenced by Gartner and Holzner (2015) (both for Germany). Ranking the two workers based on their mean wage in this setting might not yield the correct ranking. In contrast, the k-means clustering technique proposed by Bonhomme et al. (2016) allows for heterogeneity of within-firm wage distributions even beyond the second moment. However, the computational complexity of this method increases quickly with the number of moments to be estimated, hence the number of clusters/types is limited. The HLM method, in turn, allows for (almost) unique worker and firm ranks. The researcher faces a trade-off: to allow for more heterogeneity of the within-firm wage distributions, the number of types to be identified must be smaller. Our results indicate that it can be insightful to have a large number of types: it allows us to detect non-monotonic wage patterns for small worker groups.

Table 2.2: Properties of Worker Ranking

	$\bar{w}_i$	$\hat{\alpha}_i$	age	education
Correlation with $\hat{x}(i)$	0.75	0.87	0.19	0.48

Note: The table shows correlations of our estimated worker ranking ($\hat{x}(i)$) with other statistics used to rank workers in the literature: individual mean wages ($\bar{w}_i$) and estimated person-fixed effects (extracted from running the wage regression 2.2, $\hat{\alpha}_i$), as well as workers' observable characteristics, the individual means of age and education.

$\Pi(i,j)$ ($\Pi(j,i)$) is an indicator function that takes on the value 1 in case i (j) is ranked higher than j (i) and 0 otherwise. Whenever $c(i,j) > c(j,i)$ but $\Pi(i,j) = 0$ and $\Pi(j,i) = 1$, the values of the indicator functions are swapped and the value of the objective rises. The procedure continues until no further swap of workers increases the value of the objective. It runs on the set of worker pairs who are employed by the same firm at some point in time. The employment spells do not have to overlap.[93] We choose the "LIAB Mover Model" version of German matched employer-employee data because the sampling procedure maximizes the numbers of observed coworker pairs in our data, an ideal environment for the outlined computational procedure to run.[94] Importantly, we do not need to observe all workers of a given establishment to compute $c(i,j)$. The pairwise comparison of residual wages of two workers at the same firm is not affected by a potential wage premium (or firm-fixed effect) because both workers receive it.[95] We arrive at a final ranking which gives an estimate of the unobserved type, $\hat{x}(i)$, for every individual worker in our data. We group workers into 100 bins of equal size.[96] In the following, workers within one bin should be thought of as workers with the same estimated type $\hat{x}(i)$.

To understand the properties of our worker ranking $\hat{x}(i)$, we compare it to other wage statistics commonly used to rank workers in the empirical literature: mean wages, person-fixed effects (AKM), and rankings based on the observable characteristics age

[93]Recall that residual wages are deflated and net of time effects.

[94]See Appendix A.1 for sampling details.

[95]This is a big advantage over the AKM model with respect to data requirements. In a reduced-form fixed effect model, the full workforce of a firm needs to be observed in order to reliably estimate the firm-fixed effect. CHK meet this data requirement by using the universe of German employment records.

[96]The number of individual workers in every bin must be sufficiently large. We find that 100 bins/types is a good compromise between observations per bin and fineness of the type space which allows us to detect wage non-monotonicity even for small groups of workers.

and education. Table 2.2 shows correlations of those statistics with our final ranking. Notably, the rank aggregation procedure produces a ranking which is markedly different from rankings based on the alternative statistics. The correlations of our final ranking with individual mean wages ($\bar{w}_i$) and estimated person-fixed effects ($\hat{\alpha}_i$) are, naturally, positive but markedly different from 1. The correlation with the estimated worker-fixed effect (0.87) can be interpreted as a measure for how different the worker rankings in CHK-AKM and HLM are. Moreover, our final ranking is only weakly correlated with age and education. This is reasonable since the ranking procedure uses residual wages, which are net of the effect of observables.

To understand how the binning of workers modifies the ranking, Table 2.3 shows a decomposition of the respective variances of workers' observed wages, residual wages, age, and education into the shares explained within and between the bins. A relatively homogeneous distribution of a variable within the bins indicates that the binning provides a meaningful summary of the underlying heterogeneity in the respective dimension. Since our worker ranking is based on wages, the share of variance explained between the bins is relatively high, roughly two-thirds both for log wages ($\ln w_{it}$) and log residual wages ($\ln \tilde{w}_{it}$). Hence, our bins are internally homogeneous in terms of wage variation. Since our ranking is based on the residual wage (net of observables), the bins are much less homogeneous for the covariates age and education with, respectively, 95% and 66% share of the overall variance within the bins. We find workers of almost all ages (20-60 in our sample) in every bin. A high-type worker is not necessarily old and low-type workers are not simply young workers without experience. For our six education categories, the same is true, albeit to a lesser extent. Plots showing the distribution of age and education across worker bins are shown in Appendix A.2, Figure A2.1.

2.4.3 Ranking Firms

As reasoned before, we seek to rank heterogeneous firms independently of wages. This approach sets a counterpoint to both CHK-AKM and HLM, who use wage data and worker mobility to capture firm heterogeneity. The downside of our approach is that the measures we use to rank firms are by construction firm-specific and not match-specific (like wages). Ideally, the researcher would like to directly observe the firms' share of

Table 2.3: Properties of Worker Bins

	$\ln w_{it}$	$\ln \tilde{w}_{it}$	age	education
Overall Variance	0.182	0.163	96.927	1.910
Between bins	0.119 (65%)	0.107 (65%)	4.725 (5%)	0.666 (35%)
Within bins	0.064 (35%)	0.057 (35%)	92.202 (95%)	1.244 (65%)

Note: The table decomposes the overall variance of log wages ($\ln w_{it}$), residual wages ($\ln \tilde{w}_{it}$), age, and education in our sample into the respective shares explained within and between the worker bins. The age of individual workers in our sample ranges from 20 to 60. There are 6 education categories: 1 = "no degree", 2 = "vocational training", 3 = "high school", 4 = "high school and vocational training", 5 = "technical college", 6 = "university".

rents obtained from matches with each individual worker.[97] In the absence of these data, we have to be content with the argument that match-specific influences are most likely largely integrated out of firm-level statistics, particularly in large firms.

We interpret firm heterogeneity in terms of efficiency. Conceptually, this idea builds upon stochastic frontier analysis (SFA).[98] An individual firm's distance to the production frontier is a natural measure of relative performance. Given its inputs, capital and labor, the most productive firm in the economy is located on the frontier. Less efficient firms generate less output relative to their inputs. They have a positive distance to the frontier in output space. In this setting, each individual firm's distance to the production frontier is inversely related to the firm-fixed effect, which is equivalent (in absolute terms) to the technical efficiency residual in SFA terminology. A firm on the frontier has a distance of zero and thus the highest estimated fixed effect.

To rank firms indexed by k based on the distance to the production frontier, we estimate a rich empirical model of firm performance with a large number of controls from the IAB Establishment Panel. We use a flexible translog specification with firm and time effects. Log value added is our measure of output, $\ln v_{kt}$. x'_{kt} contains time-varying explanatory variables. In the translog setting, these include capital, labor, capital and labor squared as well as their interaction.[99] z'_{kt} includes dummies for 32 sectors of the

[97] This would be equivalent to directly observing match-specific output. In this hypothetical case, it would not be necessary to make a structural assumption on bargaining (like HLM do) to identify the value of a firm's vacancy and rank firms according to it.

[98] SFA has been developed in the context of cross-sectional data, see Aigner et al. (1977). It was then extended to the panel data context by, among others, Schmidt and Sickles (1984) and Cornwell et al. (1990).

[99] We approximate the capital stock using a perpetual inventory method, see Section 2.2.3. The labor input is measured by the size of the workforce.

Table 2.4: Properties of Firm Ranking

	$\ln \bar{v}_k$	$\ln \frac{\bar{v}_k}{\bar{N}_k}$	$\hat{\phi}_k$	$\ln \bar{N}_k$	$\bar{\pi}_k$
Correlation with $\hat{y}(k)$	0.47	0.75	0.94	0.17	0.60

Note: The table shows correlations of our firm ranking ($\hat{y}(k)$) with other statistics that could be used to rank firms: log mean value added ($\bar{v}_k$), log mean value added per worker ($\ln \bar{v}_k / \bar{N}_k$), and estimated firm-fixed effects (extracted from running regression 2.5, $\hat{\phi}_k$), the log of the average size of a firm's workforce, $\bar{N}_k$, as well as average profits per worker, $\bar{\pi}_k$.

economy which we include as additional controls.[100] Time effects are captured by ω_t, ϕ_k is the firm-fixed effect, and r_{kt} is the residual.

$$\ln v_{kt} = \phi_k + \omega_t + x'_{kt}\beta + z'_{kt}\gamma + r_{kt}. \tag{2.5}$$

After running this regression, we rank firms based on their estimated fixed effect, $\hat{\phi}_k$. This is our estimate of each individual firm's type. We group the firms into 30 bins of equal size, denoting the estimated rank of all firms within one group $\hat{y}(k)$.

Table 2.4 shows the correlations of our firm ranking with other firm-level statistics: log value added, log value added per worker, the log of the firms' workforce size, and profits per worker (all in firm-level means), as well as the estimated firm-fixed effect. It is not surprising that the production frontier based ranking is highly correlated with the estimated firm-fixed effects due to the close link between the two. The moderately high correlations with log value added and log value added per worker, however, are not mechanic and show that the various controls included in the regression lead to a very different firm ranking as compared to, for instance, ranking on value added alone. Moreover, controlling for the size of the workforce, its square and the interaction with the capital stock in the translog setting leads to a low but positive correlation with the log of the mean workforce size, $\bar{N}_k$. The correlation of the production frontier based ranking with our measure of profits per worker (discussed below) is of the same order of magnitude as the ranking's correlation with value added and value added per worker.

[100] We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel. The WZ classification of the German Federal Statistical Office is compatible to the common international classifications of industries, NACE and ISIC.

Table 2.5: Properties of Firm Bins

	$\bar{\pi}_k$	$\ln \bar{v}_k$	$\ln \frac{\bar{v}_k}{\bar{N}_k}$	Sector
Overall Variance	4.734E+09	3.584	0.733	66.825
Between bins	2.715E+09 (57%)	1.084 (30%)	0.475 (65%)	6.321 (9%)
Within bins	2.019E+09 (43%)	2.500 (70%)	0.258 (35%)	60.504 (91%)

Note: The table decomposes the overall variance of profits ($\bar{\pi}_k$), log value added ($\ln \bar{v}_k$), log value added per worker ($\ln \bar{v}_k/\bar{N}_k$), and of the sectors the firms operate in into the respective shares explained within and between the firm bins in our sample. We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel, which is compatible to the common international classifications of industries, NACE and ISIC. We use 32 industries, roughly classified as follows: 1-2 = "Agriculture & Mining", 3-18 = "Manufacturing", 19-20 = "Construction", 21-23 = "Retail Trade", 24-32 = "Service Sector".

As we did with the binned worker ranking, we decompose the variance of some key firm variables in our data into the shares explained between and within our firm bins to show in which dimension the bins are internally homogeneous and in which they are not. Table 2.5 shows the decomposition. The bins are internally homogeneous in terms of average profits per worker and log value added per worker with the majority of variance between the bins. This is not true for log value added alone, underlining the importance of controlling for size effects in the ranking exercise. Moreover, the variable indicating in which sector a firm operates has a very high share of within variance. Thus, the sector, which we also control for in the production function estimation, is not a strong determinant of a firm's position in the final ranking. In every bin, we find firms from almost all sectors. Additionally, we show in Figure A2.2 the plotted distributions of firm size, the sector variable, as well as collective bargaining and employee representation dummies. There is no clear relation of our firm bins to collective bargaining schemes or employee representation, the dispersion of these attributes within the firm bins is huge.

The results we present do not hinge on a specific way of ranking firms. To illustrate this, we use an alternative measure of firm performance, the average profit per worker, and check the robustness of our findings. The average profit per worker is a simple and transparent statistic to rank firms. We build on Bartolucci et al. (2015), who use very detailed firm data (balance sheets) to study labor market sorting in the Italian region of Veneto.[101] Their main argument in favor of profits is that all firms share a similar objective: maximizing profits. In addition, as argued before, firms tend to be matched

[101] They test a variety of potential measures of firm performance based on the firms' balance sheets and find that economic profits per worker are well suited to rank firms.

to a large number of workers and match-specific noise should thus be integrated out of firm-level profits.[102] On the other hand, in the context of the model presented in Section 2.3, average profits are not necessarily increasing in the unobserved type of the firm. Similar to the argument presented for the non-monotonicity of workers' wages, profits could be non-monotonically related to the unobserved firm type. This would invalidate a firm ranking based on profit data. Since the production frontier ranking is unaffected by this theoretical obstacle, we chose it to generate our baseline results.[103] The correlation between the profit-based ranking and the production frontier ranking is very high, 0.79, indicating that both rankings do an almost similarly good job in summarizing firm heterogeneity. Hence, all our main conclusions are unaffected by the ranking choice. The next Section presents our baseline results, the robustness check using the profit ranking is relegated to Appendix A.3.

2.5 Labor Market Sorting in Germany

Having ranked both workers and firms, we are now in a position to construct and analyze the empirical bivariate density of matches across all possible combinations of worker and firm types. We use 100 worker and 30 firm bins based on the types estimated in the preceding section. To compute the density of matches, we use our definition of a match described in Section 2.2.4. A match is always the first new person-year observation of an employment spell between a worker and an establishment in our data, so we count every individual worker-establishment combination only once to be as conservative as possible when estimating rank correlations. This also excludes recall. Our match definition is well-suited for the analysis because our primary interest is the allocation of workers to jobs. The abstraction from the spell length ensures that certain match types are not over-represented (under-represented) in the distribution of matches simply because they last longer (shorter) on average, leading to more (less) person-firm-year observations. Figure 2.2 plots the bivariate density of matches to illustrate the sorting patterns in our

[102] Conversely, workers are typically matched only to a small number of employers throughout their career, creating noisy wage histories.

[103] An earlier version of this research used the profit-based ranking as baseline and checked robustness with the production frontier ranking. We exchanged the two rankings solely for the purpose of methodological cleanness, none of our main results changed.

Figure 2.2: Empirical Bivariate Match Density in Germany (1998-2008, $\rho = 0.24$)

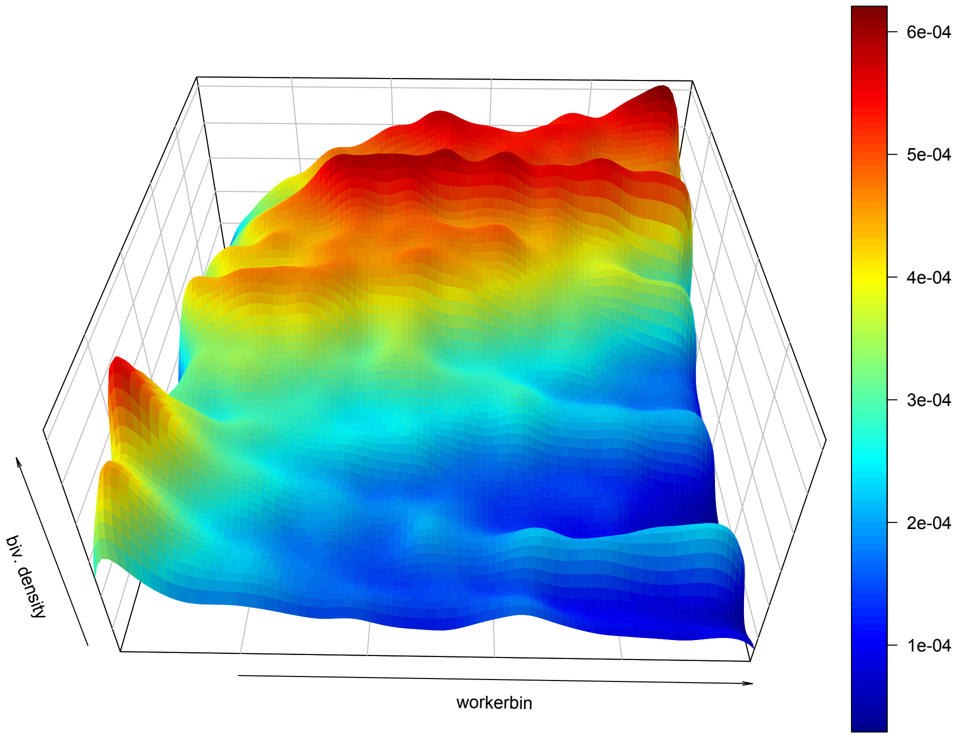

Note: Two-dimensional kernel density estimation with an axis-aligned bivariate normal kernel, evaluated on a grid with dimensions 100×30 (#worker types $\times$ #firm types).

data. There is a distinct tendency of low-type workers to be matched with low-type firms as indicated by the two spikes in the lower-left corner of the plot. Most of the matches in our data are located in the upper half of the plot, representing matches with firms in the upper half of the firm ranking. Dispersion is higher in this part of the distribution, many low-type workers are employed by medium and high firm types. The maximum match density, however, is observed for matches between the highest worker and firm types. Correlating our binned worker and firm rankings for the full sample, we find that the rank correlation coefficient (Spearman's ρ) is significantly positive with a value of 0.24. This indicates that the matching process in the German labor market features positive sorting of workers to jobs, albeit not to a very high degree. Sorting is more pronounced for matches out of unemployment with a rank correlation of 0.26. In the sample of job-to-job switchers, the rank correlation is 0.20. The degree of sorting we find for all matches is somewhat higher than the values reported by CHK for the first half of our sample. CHK report a correlation of the estimated person and establishment effects of 0.17 for the years 1996-2002. We find a rank correlation of 0.21 for 1998-2002. For the second half of our sample, the correlation found by CHK is 0.25 (2002-2009); ours is virtually similar at 0.25 (2003-2008).[104] HLM do not report rank correlations for sub samples, only an overall correlation of 0.76 (1993-2007).

Figure 2.3 presents a direct comparison of the theoretical model and the empirical density of matches. Panel (a) shows the theoretically optimal allocation along the 45-degree diagonal and the matching sets around it; Panel (b) is a contour plot of our bivariate match density in Figure 2.2. Recall that theory predicts that no matches are formed by type combinations outside the matching sets, that is, no matches of high-type firms and low-type workers or vice versa. Empirically, we clearly see in Panel (b) that a large number of matches is located in the vicinity of the diagonal, in line with theory. This explains why we find a positive rank correlation. If one would connect the points with the highest match density for every worker bin, however, the result would not be a straight line from the lower-left to the upper-right corner. Rather, we would get a concave curve above the diagonal, suggesting that the empirical allocation of workers to

[104]See Table III in Card et al. (2013), "correlation of person/establ. Effects", p. 994. Note that CHK compute the correlation for all person-firm-year observations in their sample, not just for new matches. All rank correlations reported here are rounded to two decimal points.

Figure 2.3: Sorting: Theory and Empirical Evidence

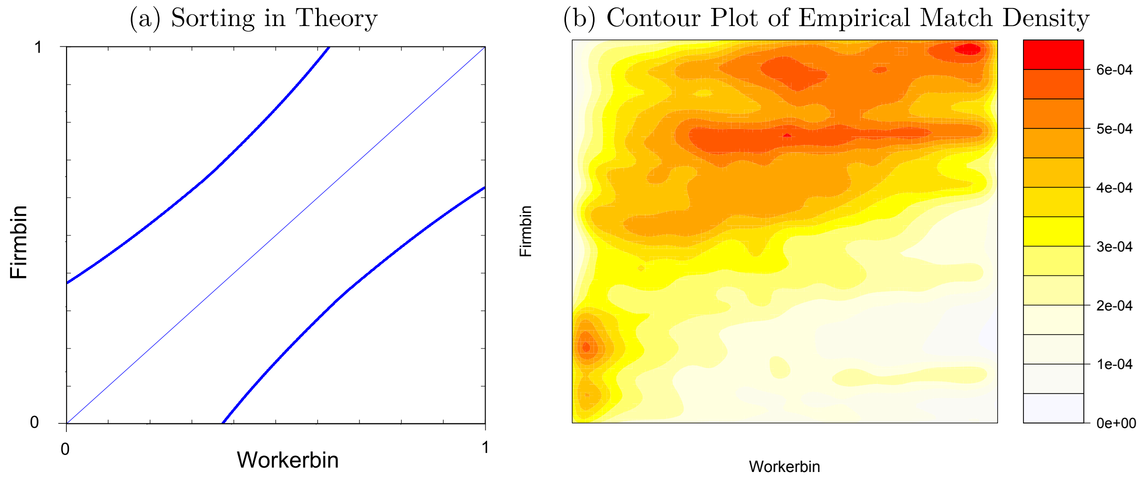

Note: In both Panels, worker and firm types are normalized into the unit interval to facilitate comparison. Panel (a) depicts the equilibrium of a simple Shimer and Smith (2000) economy. The model is a simplified version of the structural search model presented in Hagedorn et al. (2017). It is solved numerically using standard parameter values. For the derivation and solution procedure see Chapter 1. Panel (b) shows a contour plot of the empirical match density depicted in Figure 2.2.

jobs is different from the optimal allocation in the simple model. The density of matches in the lower-right corner is very low, indicating that high-skilled workers almost never work at low-productivity firms, in line with the theoretical matching sets. The prediction regarding the opposite corner, however, is not met by the data. The density is stretched out towards the upper-left corner, indicating that matches between low-type workers and high-type firms are common. The symmetry of the simple theoretical plot is thus rejected by the data, along with the diagonal optimal allocation.[105]

2.5.1 Sorting over Time

Interestingly, the depicted sorting patterns in Germany changed considerably over time. We show this in Panels (a) and (b) of Figure 2.4. Splitting our sample into two distinct time periods reveals that sorting between high-type workers and high-type firms was very pronounced in the first half of our sample (Panel (a), 1998-2002).[106] In the second half

[105]In the model, symmetry hinges on the assumption of equal bargaining powers of workers and firms. By relaxing this assumption, it is easily possible to align the theoretical model with the apparent asymmetry of matching patterns in the data. HLM show how the bargaining power can be measured in the data.

[106]The time periods are 1998-2002 and 2003-2008. We interpret our worker and firm ranks as time-constant and do not rerun the ranking procedures on the subsamples. We simply condition on the year when computing the match density.

Figure 2.4: Empirical Bivariate Match Density by Sub-Period

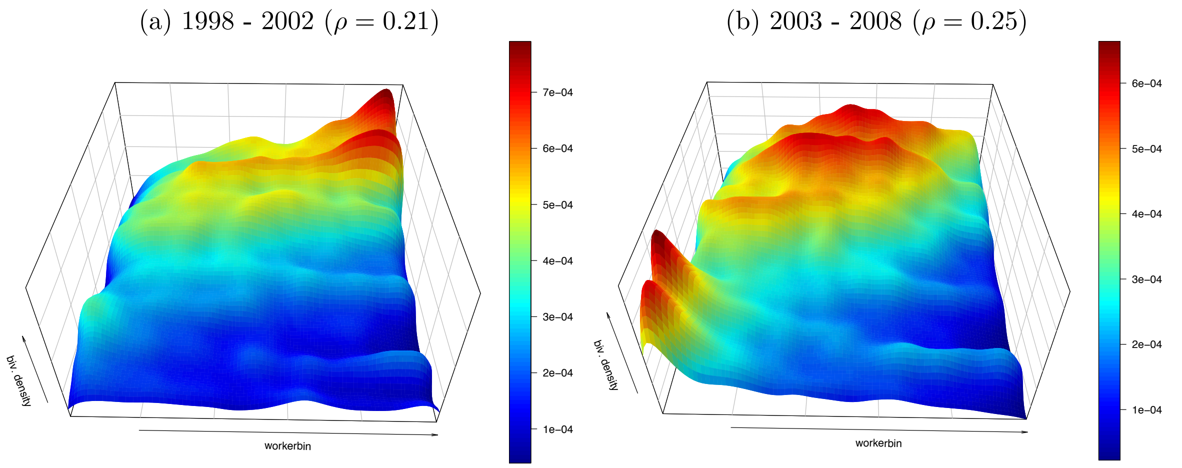

Note: Two-dimensional kernel density estimation with an axis-aligned bivariate normal kernel, evaluated on a grid with dimensions 100×30 (#worker types $\times$ #firm types). Ranks of workers and firms are time invariant, the plots are simply conditioned on the year of the match.

(Panel (b), 2003-2008), the allocation in the upper half of the plot became much more dispersed, the spike disappeared. The opposite is true at the bottom of the distribution. There was almost no match density for low-type workers and low-type firms before 2003[107], but in the second sub period we observe a huge increase of new matches between low worker and firm types. These contrary trends in the lower and upper half of the match density indicate that the overall increase of the rank correlation from 0.21 to 0.25 across the two time periods is the combined effect of two opposed developments: sorting strongly increased for low-type workers, their allocation moved closer to the theoretically predicted optimal allocation on the diagonal, increasing the overall rank correlation. The allocation of high-type workers, however, became more dispersed and moved away from the diagonal, reducing the overall rank correlation. Note that our sample split corresponds to the announcement and implementation of a big labor market reform program in Germany, the Hartz reforms.[108] Simply conditioning the match density on years suggests that the reforms potentially had a large effect on the allocation of workers to jobs in the labor market. We will investigate this hypothesis in more detail later

[107] The small hump is not present in the conditional densities before the year 2000.

[108] The so-called Hartz reforms consist of four "Acts for Modern Services on the Labor Market" ('Gesetze für moderne Dienstleistungen am Arbeitsmarkt'), which came into effect on January 1, 2003 (Hartz I and II), on January 1, 2004 (Hartz III), and on January 1, 2005 (Hartz IV). Hartz reforms I – III were mainly concerned with active labor market policies (ALMPs). Hartz IV significantly reformed the unemployment benefit system.

Figure 2.5: Rank Correlations by Match Type over Time (1998-2008)

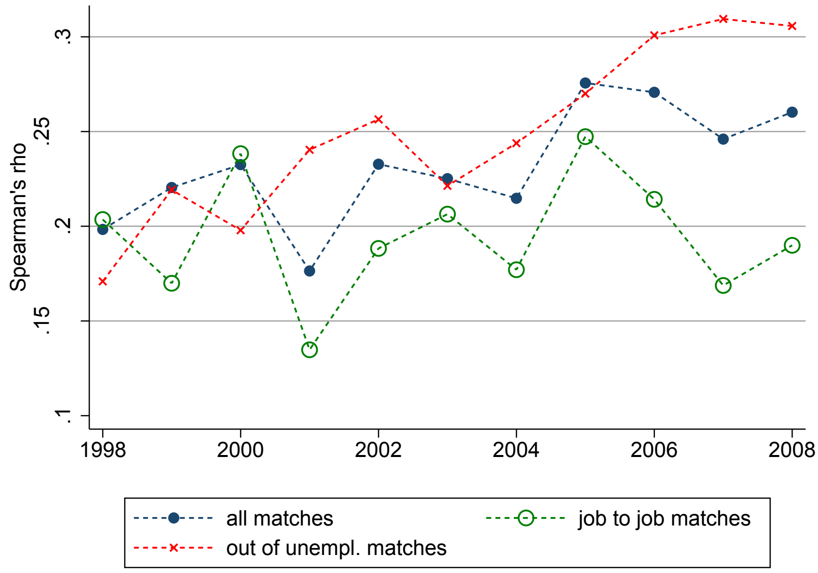

Figure 2.5 shows a different decomposition of sorting patterns over time. The blue line plots rank correlation coefficients for all matches. We then differentiate between the two types of matches: new matches of workers out of unemployment and matches from workers who switch jobs, meaning they changed their employer without an intervening unemployment spell. About 41% of the matches in our sample are formed out of unemployment, the remaining 59% are matches of job switchers. This relation is stable over time. Overall, the extent of positive sorting in the German labor market increased during our period of observation. For all matches, the rank correlation rose from 0.20 in 1998 to 0.26 in 2008 (blue line). We see no uniform trend and considerable year-to-year variation of the correlation coefficient. Specifically, the correlation fell by about 24% from 2000 to 2001. This might be related to the recession in the early 2000s, which is suggestive of a cyclical component in sorting patterns. Conditioning the rank correlation on the different match types sheds some more light on the time variation. For matches out of unemployment, a surprisingly clear trend emerges (red line): the rank correlation increased almost uniformly from 0.17 in 1998 to 0.31 in 2008. In turn, most of the year-to-year variation we observe for the overall rank correlation appears to be driven by job-to-job switchers. The green line exhibits considerable fluctuation and there is no clear time trend for this

group of matches. The observed degree of sorting for new matches resulting from search on the job did not change significantly during our period of observation. If anything, it has slightly decreased from a correlation of 0.20 in 1998 to 0.19 in 2008.[109]

Our analysis of the changing rank correlation over time suggests that the relative importance of the two match types for the overall degree sorting has reversed. Specifically, in the beginning of our sample, the rank correlation was slightly higher for job-to-job switchers, but starting in 2001, the rank correlation for new matches out of unemployment dominates. This is remarkable since theoretical models of on the job search typically suggest that sorting patterns should be more pronounced for new matches resulting from search on the job.[110] Recall that we use only wage information from employment spells formed out of unemployment to rank workers using the HLM procedure. This implies that we cannot rank workers who switch jobs throughout our sample but are never observed as being unemployed. We therefore suspect that the observed rank correlation for job-to-job switchers might be somewhat depressed in our results. This is, however, inconsequential for the remaining analysis because we largely focus on low-type workers and match formation out of unemployment.

2.5.2 Sorting by Education

To better understand empirical sorting patterns, it is also instructive to condition on education. For the U.S., Lise et al. (2016) similarly find that matches are not uniformly distributed across the type space. By splitting their sample into college graduates and workers with a high school degree or less, they show that sorting patterns are very pronounced for highly educated workers but less so for the low education group. For these workers, their results are only suggestive of a mild production complementarity, which cannot be distinguished from linearity. The German picture looks different. Panel (c) of Figure 2.6 shows, consistent with the findings for the U.S., a high concentration of

[109] A table with all rank correlations and numbers of observations by year and match type can be found in Table A2.3 in Appendix A.2.

[110] In a sequential bargaining model à la Postel-Vinay and Robin (2002) and Cahuc et al. (2006), workers gradually move toward their optimal employer (in terms of wages) and observed sorting is expected to increase. Unemployed workers, however, have an incentive to accept all job offers as long as the option value of the job lies above their reservation value, including unemployment benefits. The observation of a lower degree of sorting from search on the job as compared to sorting out of unemployment is somewhat puzzling from a theoretical perspective.

Figure 2.6: Sorting by Education

(a) No Degree (b) Secondary Degree (c) Tertiary degree

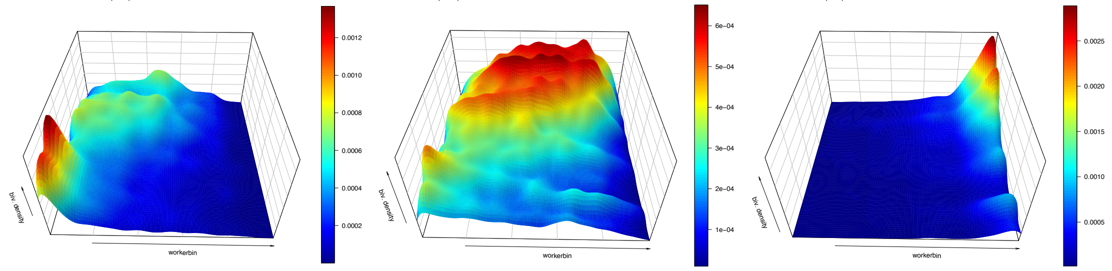

Note: Two-dimensional kernel density estimation with an axis-aligned bivariate normal kernel, evaluated on a grid with dimensions 100 × 30 (#worker types × #firm types). Panel (a) contains workers in category 1 of our education variable (no degree), Panel (b) contains workers in categories 2, 3, and 4 (vocational training only, high school only, high school and vocational training), Panel (c) contains workers in categories 5 and 6 (technical college, university).

workers with tertiary degrees at the highest firm types. These are the workers at the top of our ranking and they are strongly sorted. The allocation of workers with secondary degrees in Panel (b) is much more dispersed. We find these workers in all bins besides the highest ones. They contribute to most of the match density in the upper half of the firm ranking. However, we also observe a spike for matches of the lowest worker types in this education category with low-type firms, consistent with our finding that sorting of low-type workers into low-type firms is very pronounced in Germany. Panel (a) underlines this picture. We find workers with no educational degree exclusively in the bottom half of our worker ranking and they are concentrated in matches with low-type firms to a high degree. In contrast to what Lise et al. (2016) find for the U.S., the sorting of low-type workers is very pronounced in Germany and it has increased over time (recall Figure 2.4).

2.5.3 Distributional Dynamics

We have established that labor market sorting in Germany has increased between 1998 and 2008 in terms of aggregate rank correlations. This increase is primarily associated with a higher sorting propensity out of unemployment, while the degree of sorting of matches resulting from job switches is roughly constant over time in our data. To understand which worker-firm combinations contributed to the respective trends, we track the changing univariate distributions of worker types across firms over time, both for matches out of unemployment and for job-to-job switches. This allows us to show precisely for which worker types the distribution across firms shifted towards the theoretically predicted optimal allocation, where it shifted away from it, and where no significant change occurred. Figuratively, we slice through the empirical bivariate density of matches depicted in Figure 2.2 and compare the "density slices" for different time intervals to see where the distribution across firms changed and where it remained constant. Figure 2.7 shows the estimated univariate density functions for the worker bins 1, 10, 50, 90, and 100.[111] We compare the first half of our sample, 1998-2002 (red line), to the second half, 2003-2008 (black line). Notably, the distribution changed significantly primarily for low-type workers, see bin 1 in Panel 2.7a. For workers of medium and high-types, the density functions largely lie on top of each other and cannot be distinguished statistically.

For the lowest type of workers in bin 1, we observe an increase of the density of matches with firms of the lowest types. The density of matches with firms below, roughly, bin 8 has increased, significantly so below bin 6.[112] At the same time, the density of matches with firms in bins 12 to 22 has decreased significantly. Apparently, the distribution of workers in bin 1 across firms shifts to the left during our period of observation, see Panel 2.7a. Low-type workers become more concentrated in low-type firms, leading to increased sorting because workers of the lowest type moved towards their theoretically predicted optimal allocation. Having in mind the bivariate density in Figure 2.2, the spike of the bivariate density in the lower-left corner moved closer to the 45-degree line and became more pronounced. Distinguishing between matches out of unemployment and job-to-job switches reveals that the described shift is significant only for matches formed after an unemployment spell of the worker. This is in line with our observation that the aggregate rank correlation increased strongly for matches out of unemployment, recall Figure 2.5. Due to space constraints, we cannot plot estimated densities for all 100 worker bins. A significant distributional shift is present for all workers up to bin 8. The shifts are consistently more pronounced for matches out of unemployment than for job-to-job switches. In worker bin 10, we start observing distributions which do not change

[111] Histograms of the raw match data can be found in Appendix A.2, Figure A2.3.

[112] We deliberately choose to be conservative by determining statistical significance based on the overlap of confidence intervals. It is always true that with non-overlapping confidence intervals two statistics are significantly different. The converse, however, does not necessarily hold. As long as the difference of two statistics is significantly different from zero, a possible overlap of the confidence intervals does not imply insignificance.

Figure 2.7: Estimated Density Functions: 1998-2002 (red) vs. 2003-2008 (black)

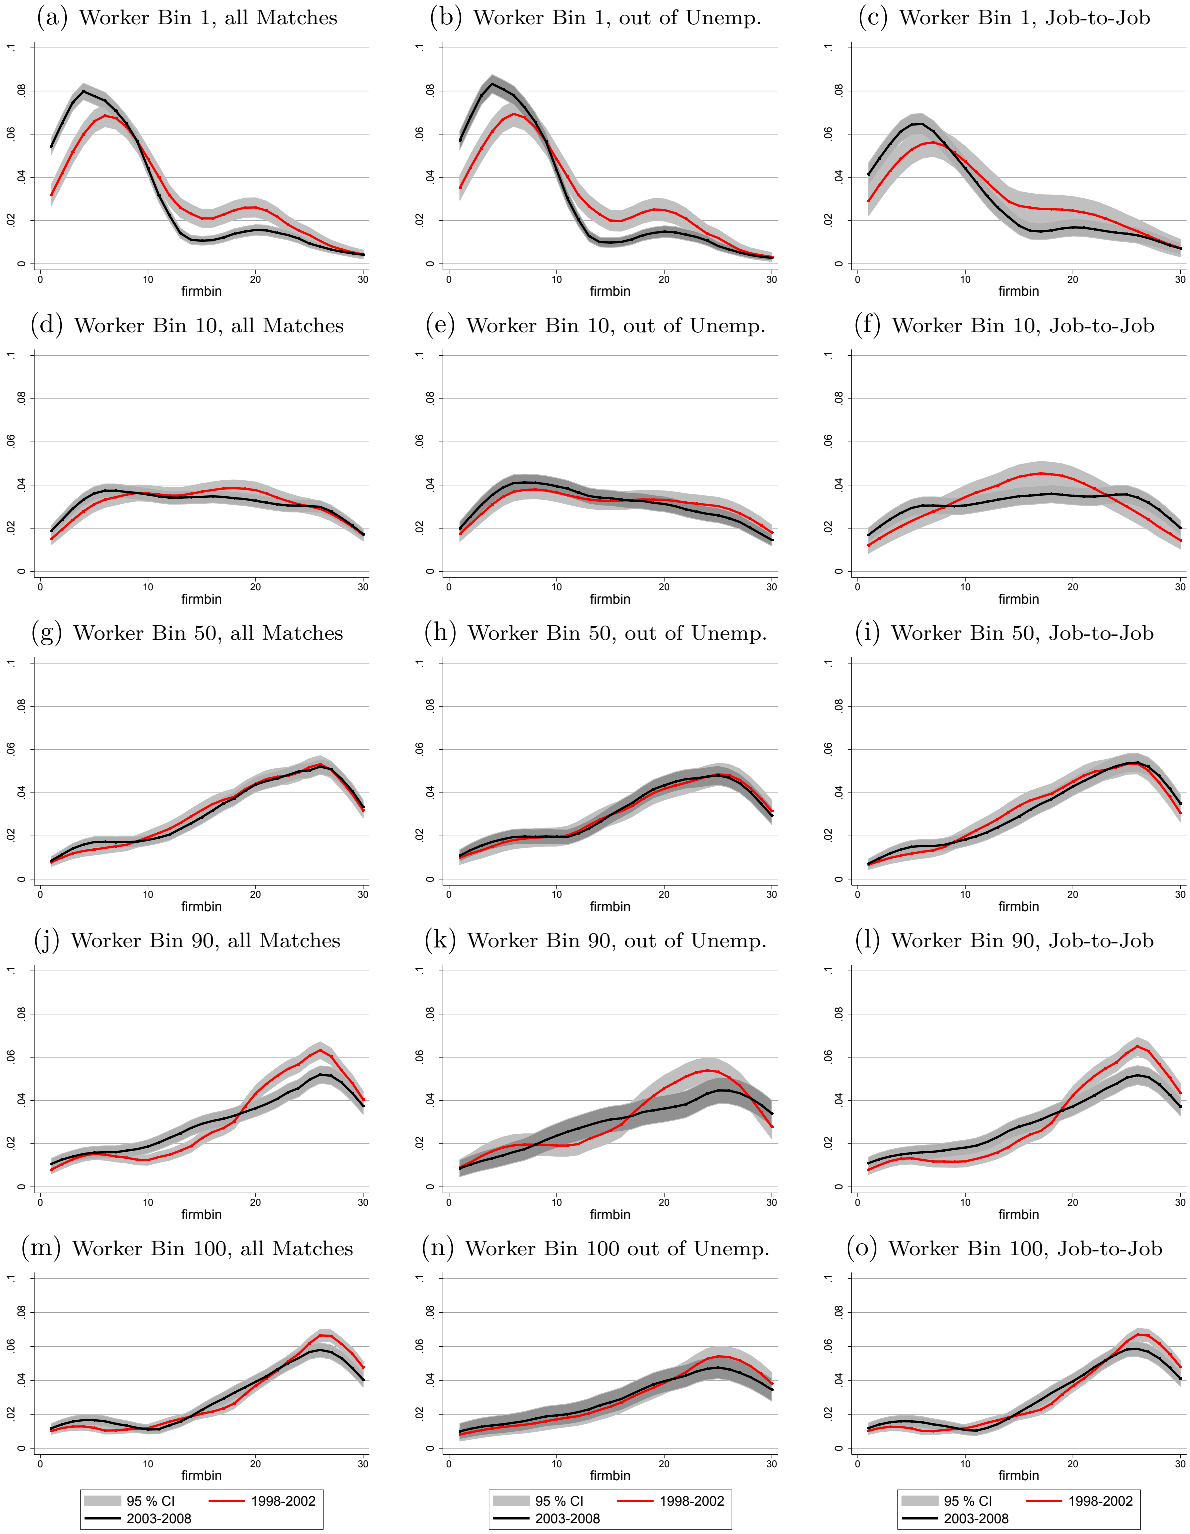

(a) Worker Bin 1, all Matches
(b) Worker Bin 1, out of Unemp.
(c) Worker Bin 1, Job-to-Job
(d) Worker Bin 10, all Matches
(e) Worker Bin 10, out of Unemp.
(f) Worker Bin 10, Job-to-Job
(g) Worker Bin 50, all Matches
(h) Worker Bin 50, out of Unemp.
(i) Worker Bin 50, Job-to-Job
(j) Worker Bin 90, all Matches
(k) Worker Bin 90, out of Unemp.
(l) Worker Bin 90, Job-to-Job
(m) Worker Bin 100, all Matches
(n) Worker Bin 100 out of Unemp.
(o) Worker Bin 100, Job-to-Job

Note: Estimated univariate kernel densities of matches conditional on worker bins, time, and match type. The kernel is estimated using an Epanechnikov kernel function. The bandwidth is calculated by Silverman's rule of thumb. Pointwise confidence intervals are calculated using a quantile of the standard normal distribution.

significantly over time, see Panels 2.7d through 2.7i. Bin 10 workers are distributed almost uniformly across firms. As we move up the worker ranking, the estimated density becomes more concentrated at the top of the firm ranking. We observe a small but significant decrease of the match density at high-type firms starting around worker bin 90, which is mainly driven by job-to-job switches, consistent with the slightly shrinking rank correlation that we observe for this match type. For the highest worker type in bin 100, this distributional shift is still visible but only marginally significant. Overall, the distribution of high-type workers shifts away from high-type firms. Dispersion increases, in line with the development depicted for the whole type space in Figure 2.4, and sorting becomes less pronounced. The mode of the distribution of high-type workers, however, is well above firm bin 25 in both halves of our sample. This is already true in bin 50, see Panels 2.7g through 2.7o.

We use estimated kernel density functions in Figure 2.7 to analyze the changing distributions of workers across firms with confidence intervals. Our findings do not hinge on the density estimation procedure, they are present in raw data as well. To show this in the context of our two sub periods, 1998-2002 and 2003-2008, we plot the difference of the two empirical densities conditional on worker bin, time interval, and match type in Figure A2.4. In case of a positive (negative) difference for a certain firm bin, the empirical density of matches has shifted upwards (downwards) for this worker-firm combination. The density differences confirm the observed distributional shifts towards more sorting at the bottom of the worker ranking, small changes in the middle, and slightly less sorting at the top.

2.5.4 (Non-)Monotonicity of Wages in Firm Types

Wages are the main determinant of how workers select themselves into jobs.[113] Hence, they are the key to understanding the distributional shifts and the increase of labor market sorting we observe in Germany. Our analysis is motivated by the theoretical prediction

[113]Sorkin (2015) uses revealed preference information contained in worker mobility patterns to construct a firm ranking and divide wage dispersion into the respective contributions of rents (or wages) and compensating differentials. Sorkin finds that compensating differentials explain up to 15% of the wage variance in the U.S. While this is a novel and interesting finding, the rents/wage-based explanation of worker selection into jobs that we and the papers we build upon focus on clearly dominates quantitatively.

Figure 2.8: Mean Wages for all Worker-Firm Type Combinations (1998-2008)

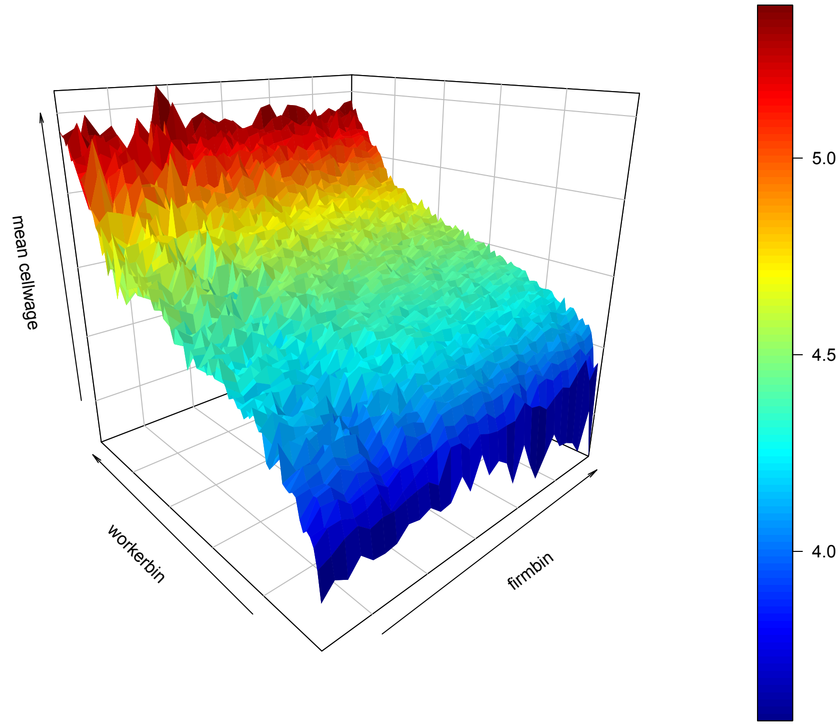

Note: The Figure shows the mean of the log real daily wage for all combinations of worker and firm types on a grid with dimensions 100 × 30 (#worker types × #firm types).

of non-monotonic wage patterns of worker types across firm types. With sorting, wages are maximized in the type-specific optimal allocation, arguably due to a production complementarity. Apart from this point, wages fall in both directions, what is at odds with the monotonicity assumption of the AKM model. Moving to a better firm does not necessarily lead to a higher wage because the worker needs to compensate the firm for not waiting for a better match. Our semi-structural identification method does not restrict the interaction of worker and firm heterogeneity. Nevertheless, we estimate rank correlations of the same order of magnitude (around 0.24) as the benchmark estimates provided by CHK, using the more restrictive AKM method. Figure 2.8 illustrates why this is the case. Simply plotting the mean of observed log wages for all combinations of worker and firm types according to our rankings reveals that linearity in logs, overall, is not a bad assumption to empirically analyze wage dispersion. The pattern of wages across the type space is not very different from a simple linear plane. This explains why the AKM framework produces a good fit. It does a good job in decomposing wage dispersion in the respective contributions of worker and firm heterogeneity, especially when the researcher has access to a data set as big and detailed as the universe of German social security records used by CHK.[114] Consistent with CHK, we find that worker heterogeneity is the dominant source of wage dispersion. Wages increase strongly in the direction of the worker type in Figure 2.8. In the firm dimension, however, they hardly visibly increase. It almost appears as if, given the worker type, it does not matter much for the wage at what type of firm a worker is employed. As we will show in the following, this conclusion, which would call the importance of labor market sorting into question, is premature. We find that the large overall wage dispersion in a joint analysis of all workers and firms masks non-monotonic wage patterns for specific worker types, particularly at the bottom of the ranking. This is consistent with the fact that deviations from monotonicity are most pronounced for worker types with large residuals in an estimated AKM model. CHK themselves state that those residuals might indicate systematic departures from the monotonicity assumption and demand an in-depth analysis.[115] We attempt to provide such an analysis in the remainder of this Chapter.

Figure 2.9 plots wage patterns across firms for single worker types. The observed higher tendency of low-type workers to sort into low-type firms indeed appears to be associated with a non-monotonic wage pattern across firms, in line with theoretical sorting models. The mean log wage of our lowest worker type (worker bin 1) increases at first, reaches its maximum in matches with firms around bin 10 and decreases thereafter in the type of the firm. This wage pattern suggests selection of low-type workers into low-type firms, simply because these are the wage-maximizing matches for them. This explains the increase of sorting at the bottom of the bivariate match distribution. To the best of our knowledge, we are the first to provide direct empirical evidence for a negative relation between a worker's wages and a performance measure of the firms he is matched with. Note, however, that this negative relation is visible in our data only for the lowest worker types. Consistent with the good overall fit of the AKM model, we find that wages are monotonically increasing in the firm type for the majority of workers in our sample. In

[114]Recall that it is imperative for applying the AKM model to observe all workers at every establishment in order to consistently estimate firm-fixed effects. Most samples of matched employer-employee data do not meet this requirement and suffer from the limited mobility bias emphasized by Andrews et al. (2008, 2012). Our semi-structural method is well-suited also for smaller samples because it is not necessary to jointly estimate worker and firm-fixed effects.

[115]A plot of the residuals and the related discussion can be found on p. 996 in CHK.

Figure 2.9: Mean Wages across Worker Types

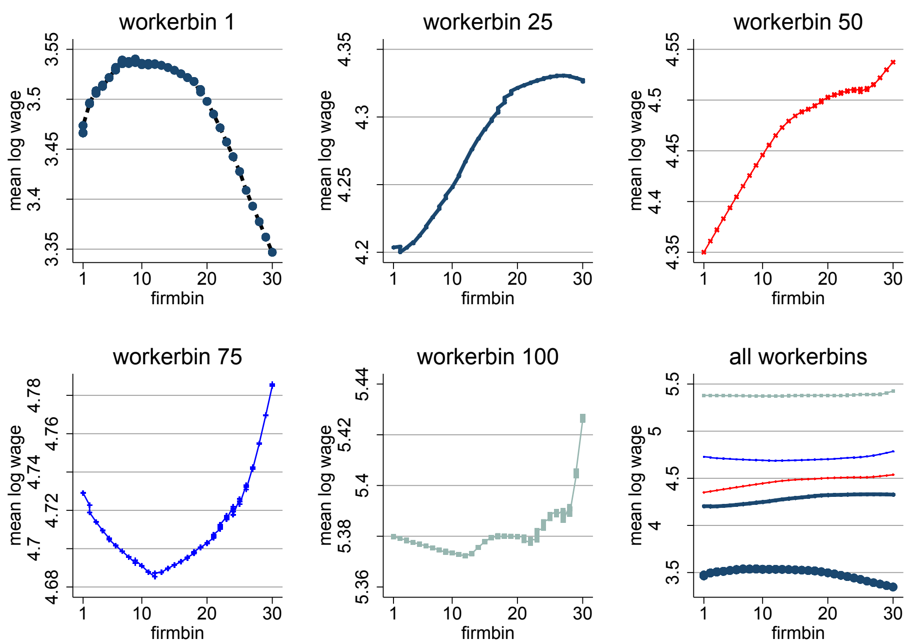

Note: Wages are means of real log cell-wages for the given worker bins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function.

bin 25, on can observe a small non-monotonicity at the top where wages level off and slightly decrease at the highest firm types. For medium-type workers in bin 50, wages are monotonically increasing everywhere. For high-type workers, the fixed effect model and the theoretical sorting model are observationally equivalent. Both frameworks predict the highest wages in matches with high-type firms. For workers in bins 75 and 100, monotonicity is only partly met. For the upper two thirds of firm types, wages increase as expected. In the bottom third, where high worker types are matched with firms of low types, we see stark deviations from monotonicity. Wages decrease at first, reach a minimum around firm bin 12 and increase only thereafter. One possible explanation for this pattern is that these workers are managers or executives at these low-type firms. Their wages might be determined in a way very different from the simple wage equations in CHK-AKM and HLM. For instance, managers could be able to extract rents from the firm in form of higher pay, even (or especially) at badly performing firms.[116] The analysis

[116]Indeed, rent extraction by managers might be a cause of bad firm performance. We see identifying different compensation schemes for different worker types within the firm as a fascinating avenue for further research, related to recent advances in our understanding of CEO compensation, see Gabaix and

of mean wages per worker bin across firm types confirms that the AKM assumption of monotonically increasing wages is met for a large part of our data, as we already suspected from Figure 2.8. For low-type workers, it is clearly violated. Match-specific effects appear to be a more important determinant of wages for low-type workers as compared to the majority of higher-ranked workers. In Section 2.5.6, we conjecture that increased sorting at the bottom and the non-monotonicity of wages are related to the tendency of firms to increasingly outsource workers in certain occupations.

Our results are in line with CHK in the sense that deviations from the two-way fixed effect model are small on average. We have shown, however, that deviations are large and systematic for low-type workers. Depending on the question at hand, the AKM model can do a good job in explaining the link between unobservable characteristics of medium and high-type workers, firms, and the wages paid. However, the model leads to wrong conclusions regarding low-type workers because the non-monotonicity of their wages is at odds with AKM. From a policy perspective, it appears crucial to take the negative relation between low-type workers' wages and the type of firm they are matched with into account, for instance to optimally design unemployment insurance.

2.5.5 Wage Inequality

Overall, wage dispersion in Germany has strongly increased during our period of observation. This is well-know, e.g. from Dustmann et al. (2009). CHK report that the variance of log real wages has increased from about 0.2 in 1995 to 0.34 in 2009.[117] This is an increase of 70%. The ranking tools we use allow us to analyze how wages change for single worker types and how these changes contribute to the overall trend of increasing wage dispersion. To this end, we plot the difference of the log wage variances over two time periods for the same worker bins as before (see Figure 2.10). We observe negative differences of log wage variances for low worker types (here bin 1) and firms in the bottom two thirds of the firm ranking. These negative differences indicate that the wage vari-

Landier (2008). Unfortunately, we cannot distinguish normal workers from managers and executives in German data. We find, however, that dropping a fraction of workers at the very top of the wage distribution within every firm reduces the prevalence of higher wages of high-type workers at low-type firms as compared to medium-type firms.

[117] These variances include imputed wages beyond the censoring threshold. For more information about the imputation procedure and wage variances, see our Section 2.2 and Table I on p. 975 in CHK.

Figure 2.10: Differences of the Variance of Log Wages over Time and across Worker Types

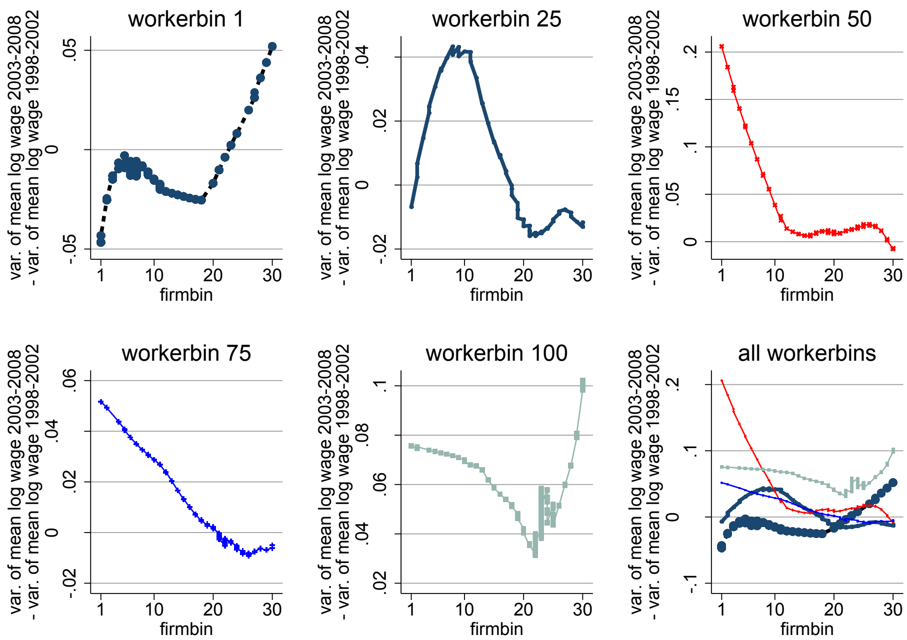

Note: Variances are calculated from log real daily wages of 30x100 bin-cells. The Figure shows differences between the variance in 2003-2008 and the variance in 1998-2002 for certain worker bins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function.

ance within these worker bins has decreased over time. This effect is most pronounced for workers in bin 1 who are employed at firms of type 1. For low-type workers, the variance has decreased stronger in the parts of the firm ranking most affected by the documented distributional shifts. In other words, increased sorting of low-type workers decreased the wage dispersion in these worker bins and, hence, compressed the wage distribution at the bottom. For worker bins 25, 50, and 75, the wage variance increased most for matches with low-type firms, particularly so for workers in bin 50. At the top of the firm ranking, wage dispersion did not change much or decreased even slightly. The wage distribution in bin 100 has become more unequal across all firm types. The lowest variance differences are found with firm types between bins 20 and 25, indicating that the slight downward shift of the match density we documented for these worker-firm combinations might have somewhat dampened wage growth (recall Figure 2.7m).

Since the observed changes of sorting and the related wage patterns have led to a reduction of within-bin wage inequality for low-type workers and an increase of within-

Figure 2.11: Wage Dynamics

(a) All Worker Bins (b) Worker Bins 1, 3, 5, 7, and 9

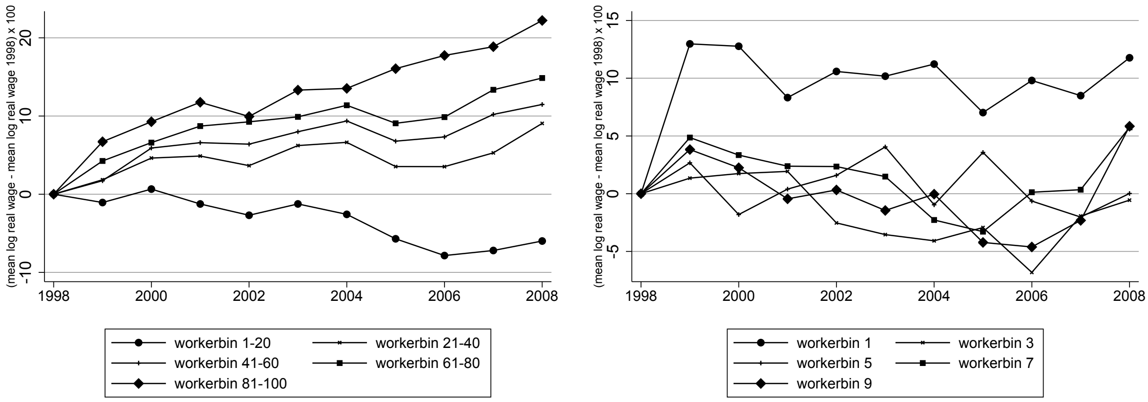

Note: The Figure shows deviations of mean log real daily wages for single and grouped worker bins relative to 1998 (multiplied by 100).

bin wage inequality for high-type workers, it is worth asking the question how these developments relate to the overall increase of wage dispersion in Germany. Therefore, in Figure 2.11 we plot how the mean wages for our worker bins evolved over time. Panel (a) confirms, in line with CHK and Dustmann et al. (2009), that overall wage inequality in Germany has increased over time during our period of observation. For the highest worker types in bins 81-100, real log daily wages grew by more than 20%. For medium worker types in bins 21-80, wages increased by around 10% to 15%. On the other hand, real log daily wages for the 20 lowest worker bins decreased by about 6% relative to 1998.

Dustmann et al. (2014) point out that a large part of the recent favorable development of the German labor market[118], can be traced back to decreasing real wages at the lower end of the wage distribution. They argue that this wage moderation was largely a result of a decentralization of the wage-setting process.[119] Zooming into the lowest worker bins as depicted in Panel (b) of Figure 2.11 reveals a novel fact relative to Dustmann et al. (2014). Using our rankings, we observe that wages for low-type workers have not been uniformly decreasing over time. In fact, the wages of workers in bin 1 increased

[118]The German aggregate unemployment rate decreased from 11.7% in July 2005 to 6% in July 2016 according to data from the Federal Statistical Office.

[119]Traditionally, wages in Germany are set by employer associations, trade unions, and works councils, typically at the industry level. The share of workers covered by this kind of industry-wide agreement decreased sharply because firms started opting-out increasingly from the mid 1990s, primarily to make the wage setting process more flexible.

by about 10%. The wages of workers in the other bins depicted here (3, 5, 7, 9) largely evolved between -5% and 5% without a clear time trend.[120] Hence, not all low worker types contributed to increasing wage dispersion in Germany. The decentralization of the wage-setting process emphasized by Dustmann et al. (2014) supposedly affected low-type workers in the manufacturing sector in particular because collective bargaining is traditionally strong in this sector. This wage moderation, or negative real wage growth, quantitatively dominates in worker bins 1-20. Low-type workers in, for example, the service sector might have very different wage dynamics. In line with our finding of non-monotonic wage patterns for workers of the lowest type and their increased propensity to sort into their wage-maximizing jobs at the bottom of the firm ranking, workers in bins 10 and below exhibit wage dynamics which are not in line with the wage moderation hypothesis, see Panel (b). However, the increased sorting of low-type workers into their wage-maximizing jobs, which is in line with wage growth for low-type workers and their reduced within-bin wage dispersion, was apparently not strong enough to revert the overall trend of increasing wage dispersion in Germany.

2.5.6 Sorting and Domestic Outsourcing

Finally, we present an explanation for the increased sorting of low-type workers into low-type firms and their non-monotonic wage patterns. We observe that a large share of the new matches of low-type workers and low-type firms is concentrated in a specific sector of the economy, the so-called business service firms. These firms are at the center of a trend towards domestic labor service outsourcing in Germany.[121] They provide services that many firms traditionally organized internally, for instance cleaning, logistics, security, and food services. Temporary work agencies are another type of business service firms. They allow firms to hire workers for tasks directly related to production, for instance assembly line workers in manufacturing, without directly employing them. By outsourcing workers, firms can exclude them from high wage premia within the firm and reduce their overall

[120] We selected bins 1, 3, 5, 7, and 9 for visibility reasons only. The same plot for bins 2, 4, 6, 8, and 10 can be found in Figure A2.5 with the patterns of wage growth being very similar.

[121] Outsourcing is commonly understood as the process of relocating tasks beyond the boundary of the firm. The term domestic implies that this happens within a country rather than internationally (offshoring).

wage bill. This point is made by Goldschmidt and Schmieder (2015), who provide a thorough analysis of domestic outsourcing in Germany using an AKM model and the universe of social security records. They show that the wages of displaced workers drop by around 10 log points due to foregone wage premia.[122] Note that in an AKM model, the establishment-fixed effect is the only channel capable of explaining a wage change when a worker switches employers. The estimated worker-fixed effect does not change by definition and the additive separability assumption excludes additional interactions of unobserved heterogeneity.

We re-evaluate the interplay between domestic outsourcing and observed wages using our more flexible worker and firm rankings. This makes a difference as compared to an AKM model because our framework allows a worker's wage to increase when he switches jobs to a firm with a lower rank. A positive match-specific component of the wage may outweigh a negative change of the firm component and lead to an overall increase. As we demonstrated in Section 2.5.4, these patterns can be observed empirically. Wages do indeed rise on average when low-type workers move down the firm ranking. We find that this wage pattern and the increased sorting of low-type workers are primarily driven by business service firms. In other words, displaced workers do not necessarily suffer a wage loss and the reason lies in possible production complementarities as highlighted by sorting theory. For example, imagine a manufacturing plant with a small number of directly employed cleaning personnel. Since cleaning is not at the core of its business, the establishment does not have an incentive to invest in increasing the productivity of these workers. Once the cleaners are displaced, they might be employed at a highly specialized cleaning service provider. This business service firm has a large incentive to invest into the productivity of these workers because it directly affects its output. This is an example of a complementarity between the worker and the firm type, which determines wages in theoretical sorting models.

To show that domestic outsourcing is a quantitatively important phenomenon, we first introduce a distinction between four broad sectors in the economy: business service firms, manufacturing, consumer services, and other services. Business service firms are at the

[122] By identifying on-site outsourcing events, where the worker and the location of the job stay the same but the employer changes, Goldschmidt and Schmieder (2015) can analyze the causal effect of outsourcing on wages. These are, however, only a small share of all outsourcing events.

heart of the outsourcing hypothesis and include temporary work agencies, security services providers, cleaning companies, and other supplementary services specifically offered to firms. Manufacturing comprises all companies engaged in the industrial production of goods. Consumer services include trading firms and the retail sector. Insurance companies, leasing providers, and other firms supplying more complex services to firms are included in the other services category.[123] For simplicity, we condense our rankings and group all worker bins into four and all firm bins into three broad groups.

In our full sample, more than 20% of all new matches in the lowest 10 firm bins involve a business service firm. Most of these firms in our sample are temporary work agencies. The majority of them has been founded after 1998, so these are young firms. For medium and high firm bins, the percentages of new matches with business service firms are much smaller, only 4.6% (firm bins 11-20) and 2.5% (firm bins 21-30), respectively. Additionally conditioning on the worker type, leads to a share of new matches with business service firms in excess of 36% for worker bins 1-25. Hence, more than a third of all matches between low-type workers and low-type firms in our sample is likely a result of domestic outsourcing activities. The German labor market reforms[124] play an important part in the observed increase of domestic outsourcing, particularly due to the deregulation of temporary work agencies. The Hartz I reform package, which came into effect on January 1, 2003, greatly liberalized temporary employment and subcontracted labor.[125] Accordingly, we observe that of all matches which low-type workers (bins 1-25) formed with low-type firms (bins 1-10) in the business service sector, about 77% were formed after 2002.

Table 2.6 provides more details about the changing distribution of matches across our four sectors. The listed numbers are percentage point differences of the shares of new

[123] To be precise, we allocate companies to the four broad sectors using the WZ93/WZ03 classification of industries available in the IAB Establishment Panel. The respective WZ codes are 1-500 for manufacturing, 501-700 for consumer services, 701-744 for other services, and 745-748 for business service firms.

[124] The so-called Hartz reforms consist of four "Acts for Modern Services on the Labor Market" ('Gesetze für moderne Dienstleistungen am Arbeitsmarkt'), which came into effect on January 1, 2003 (Hartz I and II), on January 1, 2004 (Hartz III), and on January 1, 2005 (Hartz IV).

[125] Before the reform, it was prohibited to repeatedly hire a worker on temporary contracts, to rehire a temporary worker within 3 months, to synchronize the length of the contract with an agency and the length of the assignment to a firm, and to assign a worker to a firm for more than 24 months. All these rules were abolished as a part of the Hartz I reform.

Table 2.6: Percentage Point Differences of Industry-Shares in New Matches by Worker Bins and Firm Bins (1998-2002 vs. 2003-2008)

	(a) Business Services			(b) Manufacturing		
Firm Bins	1-10	11-20	21-30	1-10	11-20	21-30
Worker Bins 1-25	10.82	1.24	0.08	-1.04	1.98	2.53
Worker Bins 26-50	2.28	0.50	0.23	-1.90	1.33	3.77
Worker Bins 51-75	1.06	0.09	0.31	-3.76	0.61	3.03
Worker Bins 76-100	0.32	-0.08	0.15	-6.51	-4.25	-8.42

	(c) Consumer Services			(d) Other Services		
Firm Bins	1-10	11-20	21-30	1-10	11-20	21-30
Worker Bins 1-25	1.50	1.02	0.88	0.20	0.68	0.06
Worker Bins 26-50	-1.01	-0.28	-0.09	-0.25	-0.05	0.28
Worker Bins 51-75	-0.65	-0.26	-0.15	-0.18	-0.90	-0.16
Worker Bins 76-100	-0.94	-0.52	-1.04	0.07	-1.12	-1.46

Note: WZ Codes 1-500 = "Manufacturing", WZ Codes 501-700 = "Consumer Services", WZ Codes 701-744 = "Other Services", WZ Codes 745-748 = "Business Services". All the percentage point differences in this Table sum to zero by construction.

matches in every sector between the two sub periods used before, 1998-2002 and 2003-2008, for different combinations of worker and firm types. For instance, we observe that the share of new matches of low-type workers (bins 1-25) and low-type firms (bins 1-10) in the business service sector has increased by 10.82 percentage points. It rose from 8.90% (1998-2002) to 19.72% (2003-2008). This is by far the largest change we observe in our sample. Similarly, but to a lesser extent, new matches of low-type workers with medium-type firms and of medium-type workers with low-type firms has increased by 1.24 and 2.28 percentage points, respectively. Conversely, manufacturing firms have reduced their direct hiring, in line with the outsourcing hypothesis. Particularly, low-type manufacturing firms have negative percentage point differences with all worker types. Moreover, manufacturing firms particularly reduced the hiring of high-type workers between the two sub periods. For consumer and other services, the percentage point differences are small.

Low-type consumer service firms increased their share of matches with low-type workers by 1.5 percentage points, the biggest change for these two sectors. It rose from 10.11% to 11.61%, so the share of new matches in this cell is substantial and it contributes to observed sorting patterns, but it did not increase as much as in the business service sector. We conclude that most of the dynamics of sorting this research documents is driven by domestic outsourcing of manufacturing firms to business service firms. Manufacturing firms reduce their direct hiring of workers and, apparently, increasingly rely on business service firms to satisfy their labor demand. Interpreting this trend in relation to low-type workers' non-monotonic wage profiles leads to the conclusion that this development and the liberalization of subcontracted labor did not necessarily lead to wage losses for low-type workers as the AKM model would suggest.

2.6 Conclusions

This Chapter reconciles empirical models of wage dispersion in the spirit of AKM with recent structural work emphasizing the importance of production complementarities and match-specific effects for wage determination and the sorting of heterogeneous workers into heterogeneous firms. In the German case, the difference of estimated rank correlations using either an AKM model (CHK) or a structural equilibrium search model (HLM) is huge. We start out from the structural identification procedure proposed by HLM and test its main identifying assumption, wage bargaining, which allows them to rank both workers and firms based on wages. We use additional firm data to construct an efficiency-based alternative firm ranking which is independent of wage information. Correlating the HLM worker ranking with the independent firm ranking delivers rank correlation which are slightly higher but broadly in line with CHK. While this confirms that AKM models generate a good fit, we document important deviations from wage monotonicity and a tendency towards more labor market sorting in the portion of the type space where the AKM model produces the largest residuals. Moreover, this research provides a detailed empirical analysis of labor market sorting in Germany which reveals a number of novel empirical facts: first, we find evidence for an increasing degree of positive sorting in the German labor market throughout a period of profound institutional change. Sorting has increased particularly for low-type workers out of unemployment, who show an increased propensity to match with low-type firms, their theoretically predicted optimal match. Second, we present direct empirical evidence that the wages of low-type workers decline in the type of the firm they are matched with. This wage pattern drives the distributional shift of low-type workers towards low-type firms. Higher wages in these matches guide the sorting of low-type workers, supporting the non-monotonicity prediction of sorting theory. Such wage patterns are at odds with the AKM fixed effect model of wage dispersion, which assume that firms pay the same wage premium to all their workers. Third, many of the new matches that contributed to the increased sorting of low-type workers involve business service firms. Increased domestic outsourcing, which has been liberalized as part of the German labor market reforms, is an important driving force behind the observed shift of the distribution of low-type workers across firms. In contrast to an AKM analysis of this trend, we conclude that domestic outsourcing does not necessarily imply wage losses for low-type workers. Fourth, our results are in line with previous studies documenting the rising overall wage inequality in Germany. We see deviations from this trend for low-type workers using our ranking method. Their wages grew by about 10% over the eleven years of our data (1998-2008), in line with the wage non-monotonicity and increased sorting. However, this wage growth for low-type workers could not overturn the overall trend of increasing wage dispersion. We did not attempt to make causal claims linking recent changes of labor market policy in Germany to the changing allocation of workers to jobs. It would appear far-fetched to assert, however, that the changes occurred in isolation. Taking our data at face value, we do not see a clear-cut discontinuity in the sorting patterns which could have been triggered by, for example, the Hartz IV reform in Germany, which reduced unemployment benefits for long-term unemployed workers. Rather, our analysis supports the conclusion of Dustmann et al. (2014), who find that the recent trends in the German labor market—the severe reduction of aggregate unemployment as well as increasing wage dispersion—can be traced back well into the 1990s and are not a direct result of the Hartz reforms. Institutional changes which take effect over prolonged time periods—for instance the decline of collective bargaining at the industry level or the trend towards domestic outsourcing—have been determinative for the development of the wage distribution and labor market sorting in Germany.

Appendix to Chapter 2

A.1 Details of Data Preparation

Our analysis is based on German matched employer-employee data. We use the "LIAB Mover Model" (file: `LIAB_MM_9308`). This section serves to detail the various data preparation and imputation procedures we apply. The LIAB Mover Model is based upon the IAB Establishment Panel. First, establishments are selected which employ at least one employee who is employed at least at two different establishments of the IAB Establishment Panel. Op to 500 additional employees per establishment are chosen randomly. The sampling procedure includes a robustness check regarding the number of employees in a certain establishment, i.e. whenever the information in the IAB Establishment Panel survey data deviates by more than 50% from the information in the register data the establishment is excluded.

Education Imputation

The employee education information is reported by employers after every year and whenever a job ends. Its quality may suffer because employers do not face consequences for non- and misreporting. However, the existence of a reporting rule allows for correction. It prescribes that only the highest educational degree of an employee needs to be reported. Therefore the individual educational attainment should not decline over consecutive job spells. The imputation procedure (IP1) as suggested by Fitzenberger et al. (2006) builds upon this reporting rule by assuming that there is any over-reporting in the data.

The original education variable distinguishes the following four different educational degrees: high school, vocational training, technical college and university. By imputing following the IP1 procedure we extrapolate both back and forwards and do some additional adjustments using individual information on age and occupational status. As a result we get six education categories which can be ranked in increasing order. However, we still observe missing entries of about 2 percent of the initial data after imputation. We drop these observations because we simply cannot make any statement about their true educational background.

Table A2.1: Summary Statistics of the Wage Distribution (1998-2008)

	Mean	Std. Dev.	Min	Max
Censored	4.582	0.393	2.411	5.153
Imputed	4.618	0.455	2.411	7.132

Note: Summary statistics of the distribution of daily real log wages.

Wage Imputation

In the LIAB data earning are right censored at the contribution assessment ceiling ('Beitragsbemessungsgrenze'). We use the pension insurance of workers and employees. This earning limit is given by the statutory pension fund and is adjusted annually due to changes in earnings. First we deflate daily wages by using the CPI with base year 2005. Then we identify censored wage observations by comparing wages with the contribution assessment ceiling. We define a wage observation as censored whenever the reported wage is higher than 99% of the censoring threshold. On average about 13% among all wage observations are censored according to our definition. Then we follow Dustmann et al. (2009) and fit a series of Tobit regression on age-education-year-combinations to impute the right tail of the wage distribution. In all of these regressions we control for eight five-year age-categories, six education categories and all possible interactions between these two categories. The imputation methodology assumes that the error term in the Tobit regression is normally distributed and each education and each age category can have different variances. Hence for each year, we impute censored wages as the sum of the predicted wage and a random component which is obtained from the standard error of the forecast. This component is separately drawn from a normal distribution with mean zero and a different variance for each education and age category. Table A2.1 shows moments of the imputed wage distributions compared to the censored wage distribution.

Table A2.2: Additional Variance-Covariance Matrices

(a) Without Top-Coded Wages

	$\ln w_{it}$	$x'_{it}\hat{\gamma}$	$\hat{\alpha}_i$	$\hat{r}_{it}$
$\ln w_{it}$	0.126			
$x'_{it}\hat{\gamma}$	0.006	0.005		
$\hat{\alpha}_i$	0.091	0.002	0.089	
$\hat{r}_{it}$	0.029	0.000	0.000	0.029

Notes: Variance-Covariance matrix of regression model 2.1 without imputation of the censored part of the wage distribution. Top-coded wages are dropped. The variance of log wages ($\ln w_{it}$) is decomposed into the variance of observable characteristics ($x'_{it}\hat{\gamma}$), the person-fixed effect ($\hat{\alpha}_i$), and the residual ($\hat{r}_{it}$). Rounded to three decimal places.

(b) Including Occupational Controls

	$\ln w_{it}$	$x'_{it}\hat{\gamma}$	$\hat{\alpha}_i$	$\hat{r}_{it}$
$\ln w_{it}$	0.207			
$x'_{it}\hat{\gamma}$	0.031	0.016		
$\hat{\alpha}_i$	0.143	0.014	0.128	
$\hat{r}_{it}$	0.034	0.000	0.000	0.034

Notes: Variance-Covariance matrix of regression model 2.1 with 32 additional occupational controls, interacted with education and time effects. The variance of log wages ($\ln w_{it}$) is decomposed into the variance of observable characteristics ($x'_{it}\hat{\gamma}$), the person-fixed effect ($\hat{\alpha}_i$), and the residual ($\hat{r}_{it}$). Rounded to three decimal places.

A.2 Further Results

Rank Correlations

Table A2.3: Rank Correlations for Different Time Intervals by Match Type

	Spearman's ρ (Number of Observations)		
	all matches	out of unemp.	job-to-job
1998-2008	0.235 (183,156)	0.257 (75,831)	0.195 (107,325)
1998-2002	0.213 (83,348)	0.218 (32,399)	0.190 (50,949)
2003-2008	0.250 (99,808)	0.281 (43,432)	0.200 (56,376)
1998-1999	0.211 (28,493)	0.196 (11,062)	0.189 (17,431)
2000-2001	0.206 (40,698)	0.217 (15,408)	0.188 (25,290)
2002-2003	0.229 (28,860)	0.239 (12,051)	0.197 (16,809)
2004-2005	0.245 (33,913)	0.257 (14,925)	0.211 (18,988)
2006-2008	0.259 (51,192)	0.306 (22,385)	0.191 (28,807)
1998	0.198 (13,924)	0.171 (5,739)	0.204 (8,185)
1999	0.221 (14,569)	0.220 (5,323)	0.170 (9,246)
2000	0.233 (21,741)	0.198 (8,589)	0.238 (13,152)
2001	0.176 (18,957)	0.240 (6,819)	0.135 (12,138)
2002	0.233 (14,157)	0.256 (5,929)	0.188 (8,228)
2003	0.225 (14,703)	0.221 (6,122)	0.206 (8,581)
2004	0.215 (17,442)	0.244 (7,622)	0.177 (9,820)
2005	0.276 (16,471)	0.270 (7,303)	0.247 (9,168)
2006	0.271 (18,177)	0.301 (8,202)	0.214 (9,975)
2007	0.250 (18,187)	0.310 (7,857)	0.169 (10,330)
2008	0.260 (14,828)	0.306 (7,857)	0.190 (8,502)

Notes: We test the null hypothesis that worker and firm bins are statistically independent. All rank correlation coefficients are different from 0 at 1% level of significance. Rounded to 3 decimal places. Numbers of observations (new matches according to our definition) are reported in brackets.

Distribution of Observables Across Worker and Firm Bins

Figure A2.1: Observable Characteristics of Workers by Bin (Full Sample, 1998-2008)

(a) Age (b) Education

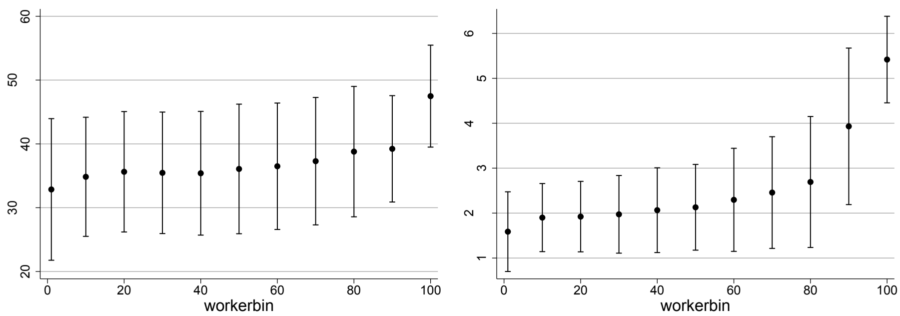

Note: Panels (a) and (b) show the means +/– one standard deviation of workers' age and education across worker bins. The age of individual workers in our sample ranges from 20 to 60. There are 6 education categories: 1 = "no degree", 2 = "vocational training", 3 = "high school", 4 = "high school and vocational training", 5 = "technical college", 6 = "university".

Figure A2.2: Observable Characteristics of Firms by Bin (Full Sample, 1998-2008)

(a) Number of Workers (b) Sector

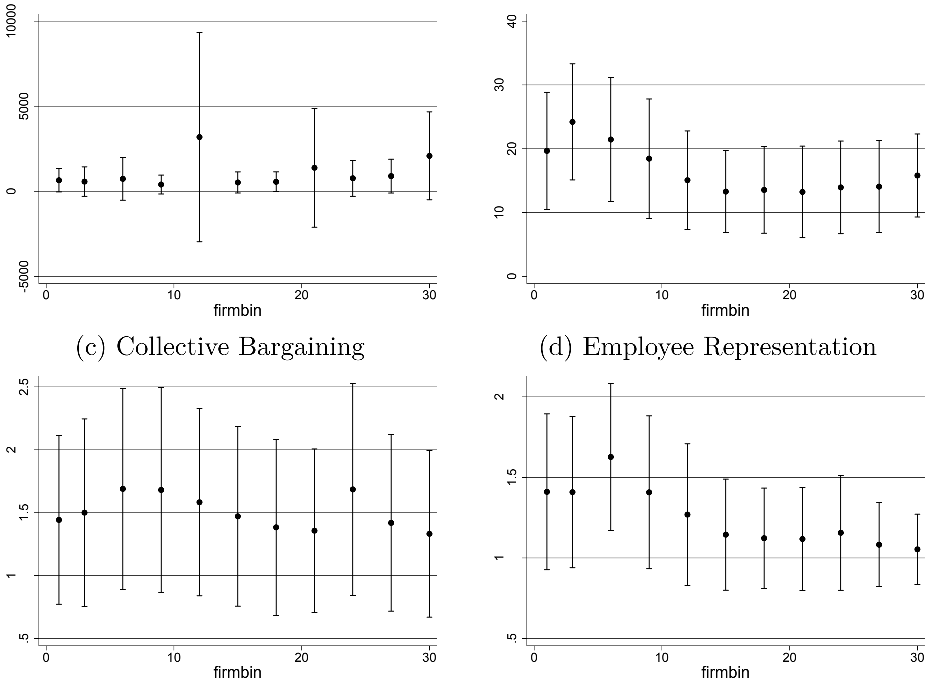

Note: Panel (a) shows means +/– one standard deviation of the average workforce size of firms within the bins. Panel (b) shows means +/– one standard deviation of the firms' sectoral classification within the bins. We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel, which is compatible to the common international classifications of industries, NACE and ISIC. We use 32 industries, roughly classified as follows: 1-2 = "Agriculture & Mining", 3-18 = "Manufacturing", 19-20 = "Construction", 21-23 = "Retail Trade", 24-32 = "Service Sector". Panel (c) shows means +/– one standard deviation of the variable indicating the application of a collective bargaining agreement within the bins: 1 = "sectorwide bargaining", 2 = "firmwide bargaining", 3 = "no collective bargaining". Panel (d) shows means +/– one standard deviation of the variable indicating the presence of formal employee representation within the bins: 1 = "employee representation exists", 2 = "no employee representation".

Histograms and Kernel Density Estimates

Figure A2.3: Histograms of Raw Data and Estimated Density Functions

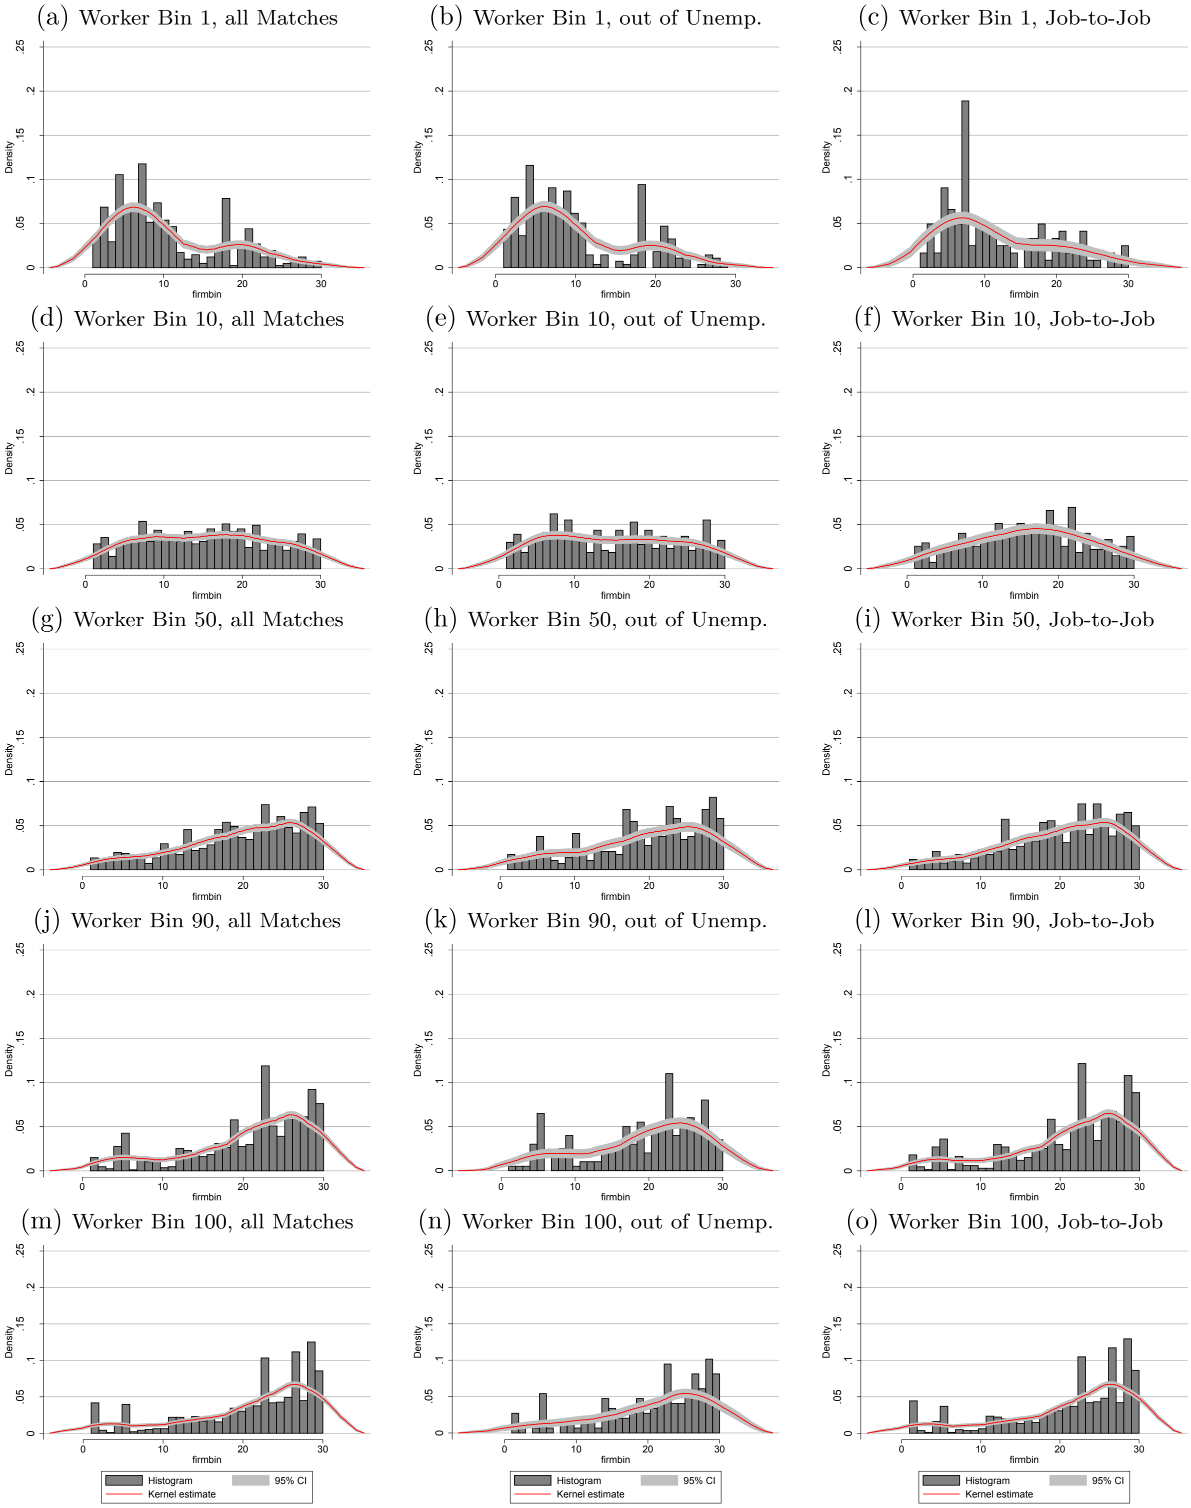

(a) Worker Bin 1, all Matches (b) Worker Bin 1, out of Unemp. (c) Worker Bin 1, Job-to-Job

(d) Worker Bin 10, all Matches (e) Worker Bin 10, out of Unemp. (f) Worker Bin 10, Job-to-Job

(g) Worker Bin 50, all Matches (h) Worker Bin 50, out of Unemp. (i) Worker Bin 50, Job-to-Job

(j) Worker Bin 90, all Matches (k) Worker Bin 90, out of Unemp. (l) Worker Bin 90, Job-to-Job

(m) Worker Bin 100, all Matches (n) Worker Bin 100, out of Unemp. (o) Worker Bin 100, Job-to-Job

Note: Histograms of raw data and estimated univariate kernel densities of matches conditional on worker bins, time, and match type. Kernel: Epanechnikov. The bandwidth is calculated by Silverman's rule of thumb. Pointwise confidence intervals are calculated using a quantile of the standard normal distribution.

Density Differences

Figure A2.4: Density Differences: 2003-2008 – 1998-2002

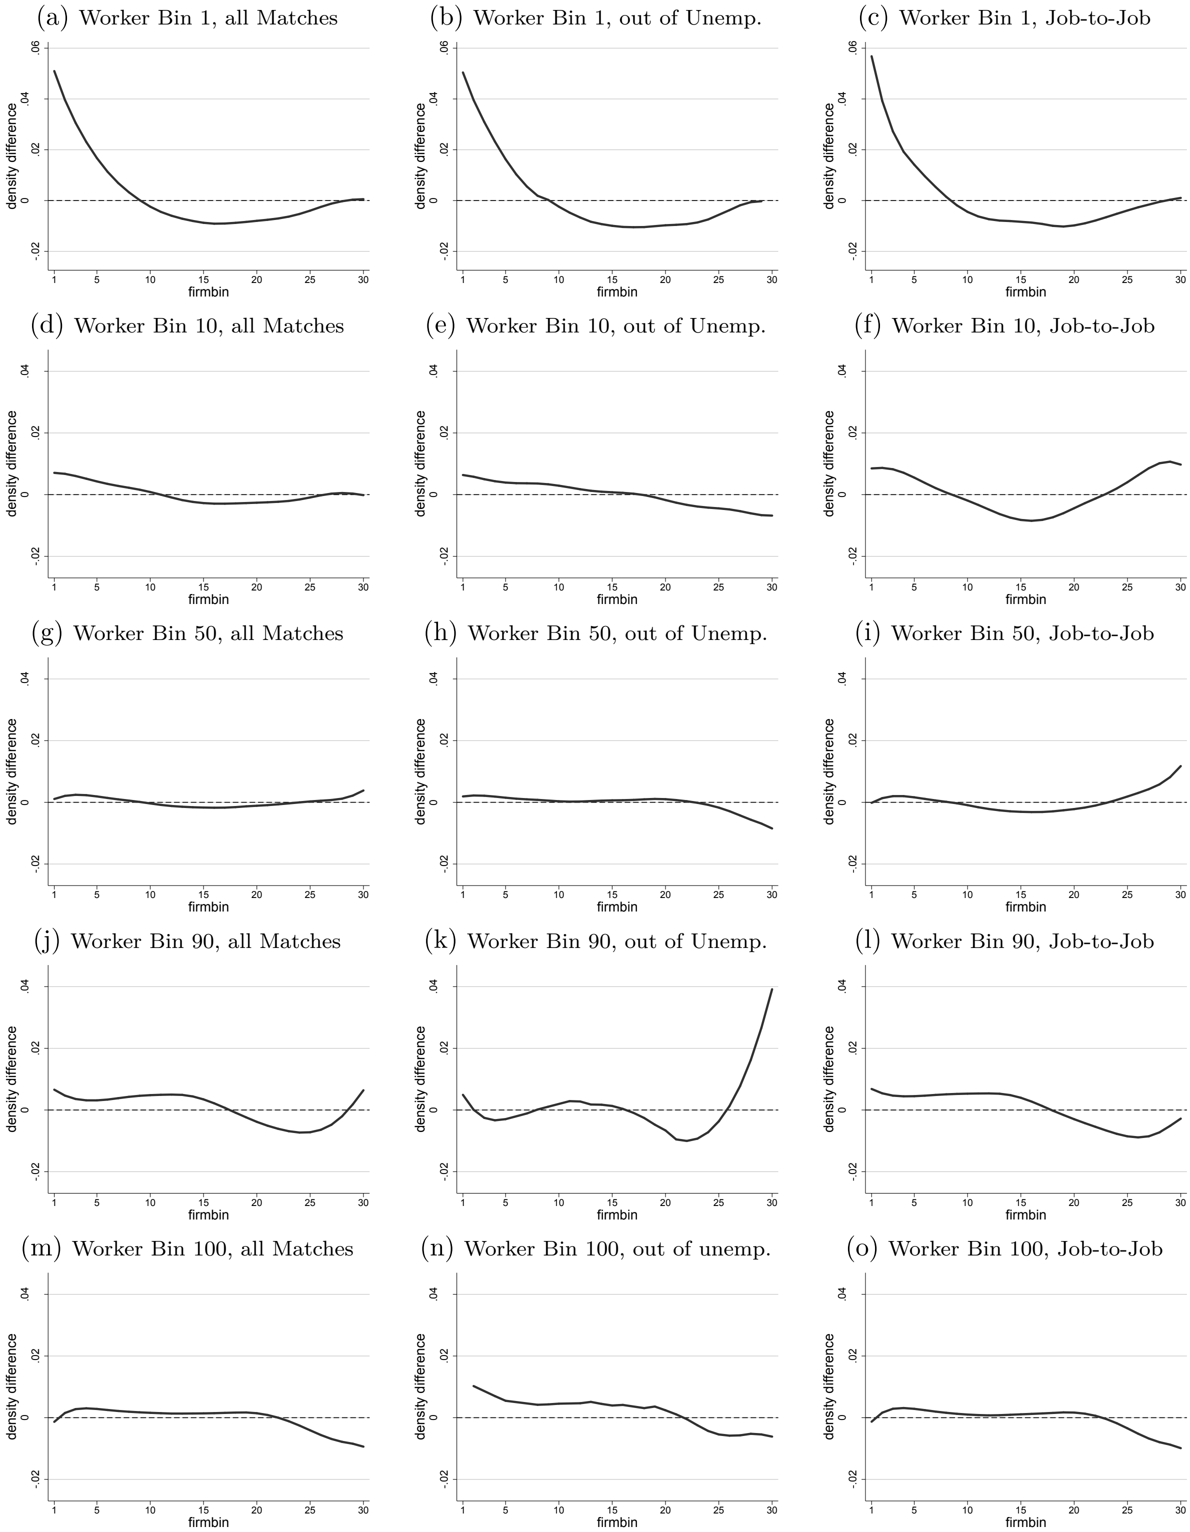

(a) Worker Bin 1, all Matches
(b) Worker Bin 1, out of Unemp.
(c) Worker Bin 1, Job-to-Job
(d) Worker Bin 10, all Matches
(e) Worker Bin 10, out of Unemp.
(f) Worker Bin 10, Job-to-Job
(g) Worker Bin 50, all Matches
(h) Worker Bin 50, out of Unemp.
(i) Worker Bin 50, Job-to-Job
(j) Worker Bin 90, all Matches
(k) Worker Bin 90, out of Unemp.
(l) Worker Bin 90, Job-to-Job
(m) Worker Bin 100, all Matches
(n) Worker Bin 100, out of unemp.
(o) Worker Bin 100, Job-to-Job

Note: Density differences are calculated by subtracting the raw match frequencies by bin. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function.

Wage Dynamics

Figure A2.5: Wage Dynamics in Worker Bins 2, 4, 6, 8, and 10

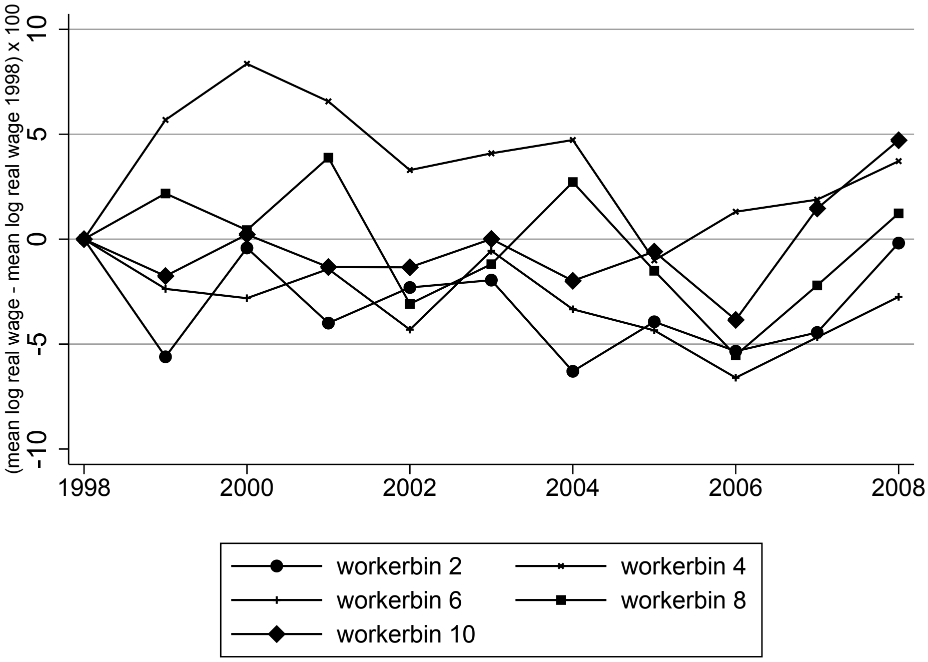

Note: The Figure shows deviations of mean log real daily wages for single worker bins relative to 1998 (multiplied by 100).

A.3 Robustness

Alternative Firm Ranking based on Profits per Worker

The construct a firm ranking based on average profits per worker, we build on Bartolucci et al. (2015) who use very detailed firm data (balance sheets) to study labor market sorting in the Italian region of Veneto. Using the IAB Establishment Panel, we can compute economic profits per worker of firm k in period t, π_{kt}, by simply subtracting the reported costs from firms' revenues: $\pi_{kt} = \frac{\Pi_{kt}}{N_{kt}} = R_{kt} - C_{kt} - W_{kt} - K_{kt}$. Π_{kt} denotes aggregate profits, N_{kt} the size of the workforce, R_{kt} revenues, C_{kt} input costs, and W_{kt} the wage bill. K_{kt} represents the capital cost of firm k in year t, which we compute by multiplying the capital stock with a yearly interest rate of 7.1%.[126] We do not observe the capital stock directly and approximate it for each establishment-year observation using a perpetual inventory method as outlined in section 2.2.3. We then rank firms based on the average of profits per worker over time, $\bar{\pi}_k$. We drop the most extreme outliers and group firms into 30 bins of equal size, denoting the estimated rank of all firms within one group $\hat{y}(k)$. Table A2.4 shows correlations of the profit based firm ranking with other firm-level statistics: log value added, log value added per worker, the log of the firms' workforce size (all in firm-level means), as well as the estimated firm-fixed effect. Table A2.5 decomposes the variance of some key firm variables in our data into the shares explained between and within the firm bins based on the profit ranking.

[126]This number is taken from Evers et al. (2015).

Table A2.4: Properties of Firm Ranking based on Profits

	$\ln \bar{v}_k$	$\ln \frac{\bar{v}_k}{\bar{N}_k}$	$\hat{\phi}_k$	$\ln \bar{N}_k$
Correlation with $\hat{y}(k)$	0.52	0.70	0.75	0.15

Note: The table shows correlations of our alternative firm ranking based on profits per worker with other statistics that could be used to rank firms: log mean value added ($\bar{v}_k$), log mean value added per worker ($\ln \bar{v}_k/\bar{N}_k$), and estimated firm-fixed effects (extracted from running regression 2.5, $\hat{\phi}_k$), the log of the average size of a firm's workforce, $\bar{N}_k$.

Table A2.5: Properties of Firm Bins based on Profits

	$\bar{\pi}_k$	$\ln \bar{v}_k/\bar{N}_k$	$\bar{N}_k$	Sector
Overall Variance	4.78E+09	0.873	1.92E+07	65.64
Between bins	4.06E+09 (85%)	0.477 (55%)	2.02E+06 (11%)	5.99 (9%)
Within bins	7.27E+08 (15%)	0.396 (45%)	1.72E+07 (89%)	59.65 (91%)

Note: The table decomposes the overall variance of average profits per worker ($\bar{\pi}_k$), log value added per worker ($\ln \bar{v}_k/\bar{N}_k$), the size of the firms' workforce ($\bar{N}_k$), and of the sectors the firms operate in into the respective shares explained within and between the firm bins based on the alternative profit ranking. We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel, which is compatible to the common international classifications of industries, NACE and ISIC. We use 32 industries, roughly classified as follows: 1-2 = "Agriculture & Mining", 3-18 = "Manufacturing", 19-20 = "Construction", 21-23 = "Retail Trade", 24-32 = "Service Sector".

Rank Correlations using the Profit Ranking

Table A2.6: Rank Correlations for Different Time Intervals by Match Type

	Spearman's ρ (Number of Observations)		
	all matches	out of unemp.	job-to-job
1998-2008	0.254 (167,764)	0.253 (69,398)	0.228 (98,366)
1998-2002	0.251 (77,892)	0.228 (30,236)	0.242 (47,656)
2003-2008	0.255 (89,872)	0.270 (39,162)	0.215 (50,710)
1998-1999	0.227 (26,743)	0.196 (10,386)	0.215 (16,357)
2000-2001	0.272 (37,949)	0.238 (14,303)	0.271 (23,646)
2002-2003	0.234 (26,694)	0.241 (11,174)	0.211 (15,520)
2004-2005	0.251 (31,003)	0.251 (13,675)	0.228 (17,328)
2006-2008	0.264 (45,375)	0.293 (19,860)	0.209 (25,515)
1998	0.195 (13,070)	0.177 (5,411)	0.200 (7,659)
1999	0.254 (13,673)	0.214 (4,975)	0.225 (8,698)
2000	0.250 (20,171)	0.225 (7,906)	0.249 (12,265)
2001	0.295 (17,778)	0.253 (6,397)	0.294 (11,381)
2002	0.236 (13,200)	0.259 (5,547)	0.203 (7,653)
2003	0.232 (13,494)	0.222 (5,627)	0.220 (7,867)
2004	0.228 (15,974)	0.248 (6,983)	0.198 (8,991)
2005	0.275 (15,029)	0.253 (6,692)	0.259 (8,337)
2006	0.269 (16,549)	0.281 (7,467)	0.226 (9,082)
2007	0.251 (15,935)	0.299 (6,916)	0.185 (9,019)
2008	0.275 (12,891)	0.302 (5,477)	0.218 (7,414)

Notes: Rank correlations using alternative firm ranking based on average profits per worker as detailed in Section A.3. We test the null hypothesis that worker and firm bins are statistically independent. All rank correlation coefficients are different from 0 at 1% level of significance. Rounded to 3 decimal places. Numbers of observations (new matches according to our definition) are reported in brackets.

Main Plots using the Profit Ranking

Figure A2.6: Empirical Bivariate Density of Matches in Germany (1998-2008)

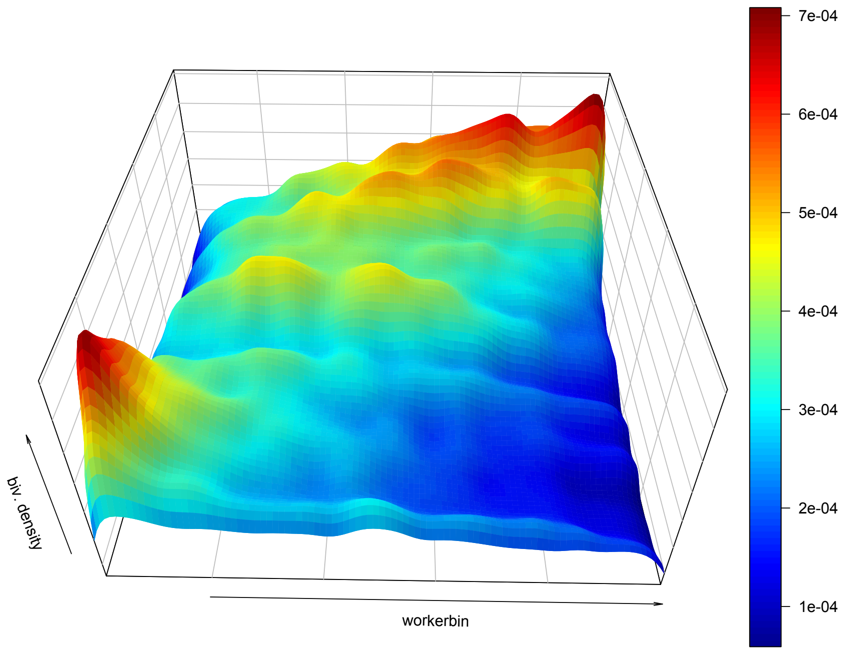

Note: Two-dimensional kernel density estimation with an axis-aligned bivariate normal kernel, evaluated on a grid with dimensions 100 × 30 (#worker types × #firm types). Alternative firm ranking based on average profits per worker as detailed in Section A.3.

Figure A2.7: Contour Plot of Empirical Bivariate Match Density

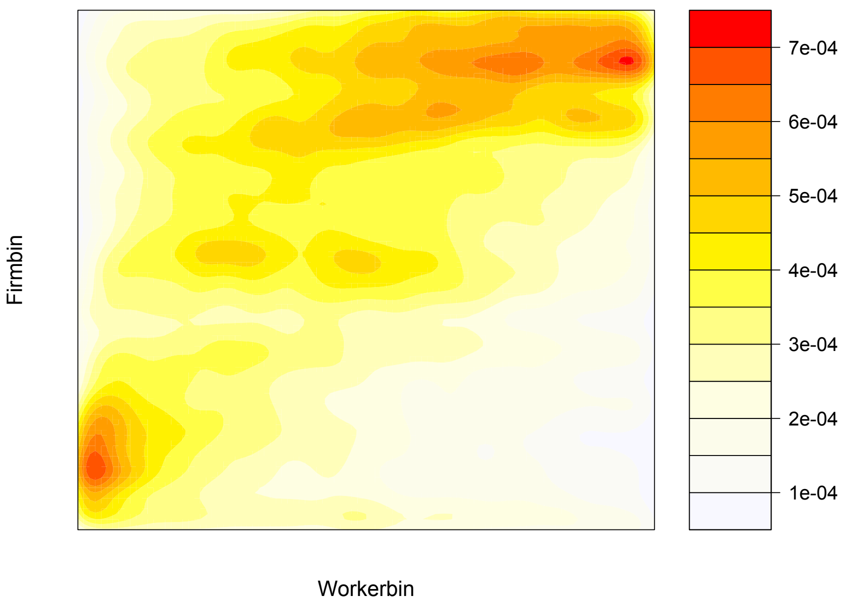

Note: Contour plot of Figure A2.6. Alternative firm ranking based on average profits per worker as detailed in Section A.3.

Figure A2.8: Rank Correlations by Match Type over Time (1998-2008)

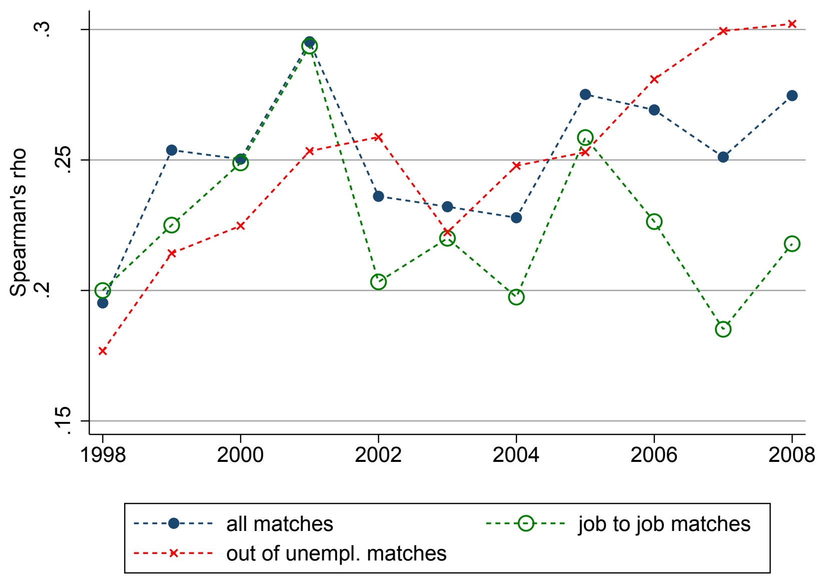

Note: Alternative firm ranking based on average profits per worker as detailed in Section A.3.

Figure A2.9: Estimated Mean Wages

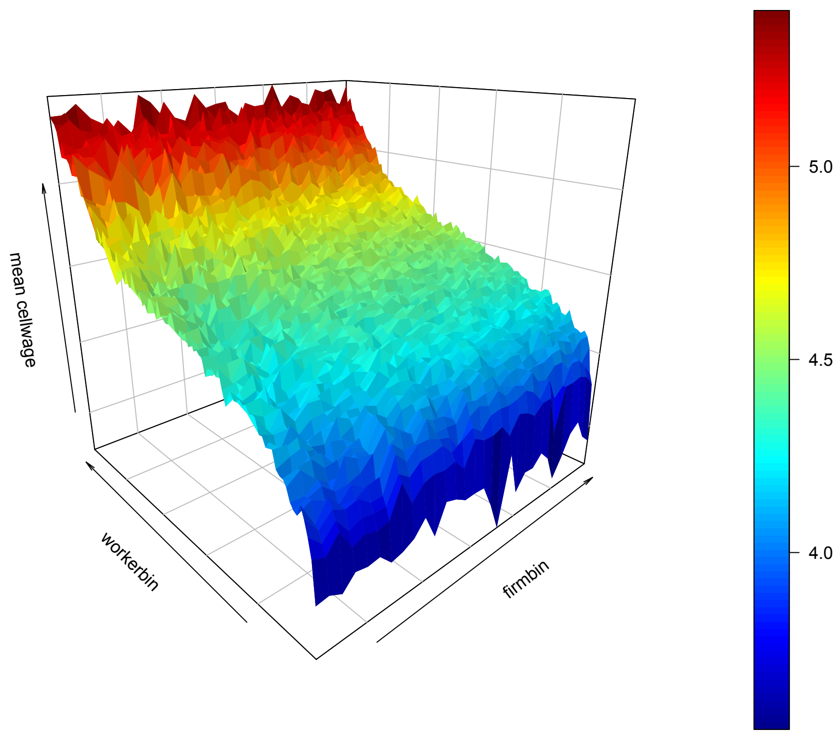

Note: Alternative firm ranking based on average profits per worker as detailed in Section A.3. The Figure shows the mean of the log real daily wage for all combinations of worker and firm types on a grid with dimensions 100 × 30 (#worker types × #firm types).

Figure A2.10: Estimated Mean Wages across Worker Types

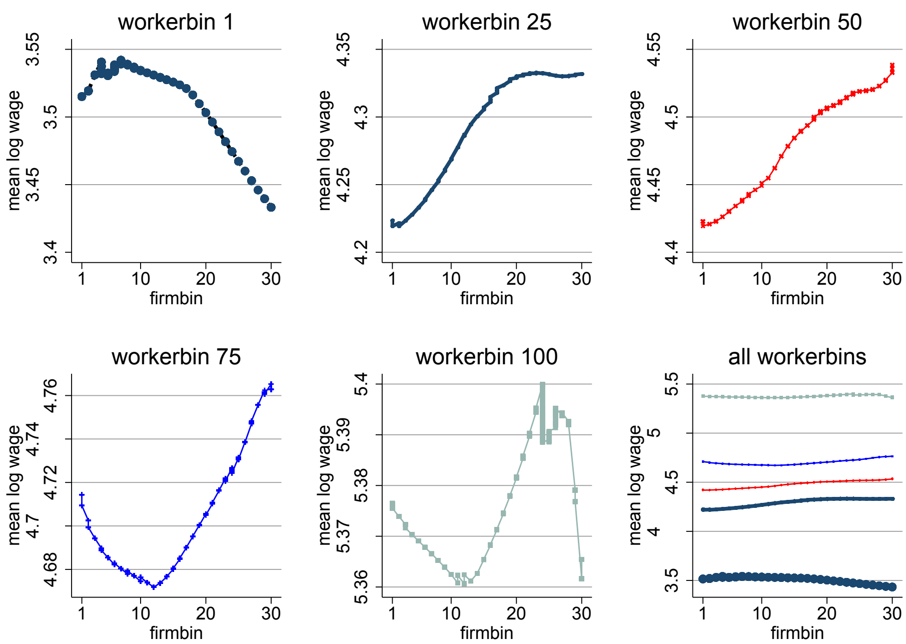

Note: Wages are means of real log cell-wages for the given workerbins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function. Alternative firm ranking based on average profits per worker as detailed in Section A.3.

Chapter 3

Marriage and Divorce under Labor Market Uncertainty*

*This research is joint work with Christian Holzner, University of Munich and ifo Institute. We thank Stéphane Bonhomme, Matthias Doepke, Greg Kaplan, Helmut Rainer, and Michèle Tertilt for inspiring discussions. This project has evolved as part of the ifo Institute's research group on "Economic Uncertainty and the Family" (EcUFam), led by Natalia Danzer. We are indebted to her and gratefully acknowledge financial support from the Leibniz Association. We also thank the ifo project "Calculation of returns on education in Germany" for helping with data access and Philipp Hochmuth and Patrick Schulze for excellent research assistance. Finally, Heiko Bergmann at the Research Data Center of the Bavarian Statistical Office has been immensely helpful during our data work.

3.1 Introduction

Among the many choices individuals make during their lifetime marriage is one of the most, if not the most, important decision. The marriage vow to be true to each other in good times and in bad, in sickness and in health and to love and honor each other for as long as one shall live reminds both partners that this is truly a decision taken under uncertainty. Strokes of fate like unexpected unemployment and severe sickness can stress a partnership and cause partners to drift apart and divorce.

The relative importance of economic shocks compared to other shocks disrupting a marriage is still poorly understood. The economic literature has documented that unemployment, especially male unemployment, is associated with an increase in the divorce rate.[127] Also, we know that marriage and divorce rates are negatively correlated with the unemployment rate over the business cycle.[128] Additionally, we know that marriage rates declined since the 1970s and that assortative matching with respect to education has increased.[129] Researchers have proposed explanations based on improvements in household technology since World War II and increased female labor supply,[130] as well as increased incentives for females to invest in education.[131] Very little is known, however, about the nature of the channels that connect marriage market decisions to the underlying source of economic shocks.

To investigate the importance of economic shocks, we integrate labor market shocks into a two-sided marriage market model with transferable utility and ex-ante heterogeneous men and women (Shimer and Smith, 2000; Jacquemet and Robin, 2013; Goussé et al., 2017). Individuals search for partners in the marriage market and, at the same time, switch between employment and unemployment. The employment statuses of both partners influence utility flows and the sharing of resources within the household.[132] A

[127] See Jensen and Smith (1990), Hansen (2005), and Amato and Beattie (2011) among others.

[128] See Schaller (2013) and González-Val and Marcén (2017a,b), among others.

[129] Both Doepke and Tertilt (2016) and Greenwood et al. (2017) offer excellent literature overviews, the latter with some cross-country facts.

[130] See Greenwood et al. (2005a) and Greenwood et al. (2005b). More recently, Greenwood et al. (2016) use a search model to analyze these trends empirically for the U.S. with an emphasis on sorting.

[131] See Nick and Walsh (2007); Chiappori et al. (2009)

[132] Guler et al. (2012) show that the joint search behavior of married couples in the labor market may be very different from the search of singles, for instance due to different reservation wage strategies or risk sharing motives. We must abstract from this important issue because it is impossible to estimate labor market transition rates conditional on marital status with our data.

negative shock, i.e., job loss of one spouse, may decrease the marital surplus sufficiently to trigger a divorce. Additionally, an idiosyncratic component captures non-economic factors of marriage (e.g. mutual affection). It is subject to shocks and may lead to separations as well. A complementarity in the household production function induces the tendency to sort positively. Given the German context, we also consider benefits from joint taxation, which have the opposite effect and encourage negative sorting. In the model, the balance between all these forces determines marriage and divorce flows, differentially across heterogeneous men and women.

The relative importance of each of these forces is an open empirical question. We thus take our model to the data. Using German micro data from various sources, we use our model as a tool to decompose marriage and divorce flows into the respective contributions of economic and non-economic forces. To this end, we develop a structural estimation procedure that allows us to back out key components of our model from the data. We estimate meeting rates, marriage probabilities, and separation rates, all differentiated according to individuals' education and labor market status.

The marriage rate depends on an individual's chance to meet somebody from a certain education group times the probability he/she is willing to marry. We show that the probabilities to marry (willingness to marry) upon meeting is highest for employed individuals with equally educated couples. A similar positive assortative matching pattern emerges for medium and highly educated individuals in all other labor market status combinations (male employed/female unemployed, male unemployed/female employed, and male unemployed/female unemployed). Low educated females still have reasonably high chances to marry with a medium or highly educated male if they stay out of the labor force (remain unemployed), most likely because of the high financial incentives provided by joint income taxation in Germany. Low educated, unemployed males have almost no chances to marry. Marriage rates are also driven by an individual's chance to meet somebody from a certain education group. Our estimates suggest that medium and highly educated individuals direct their search such that the number of meetings with individuals from the other sex with similar education level are higher than the number of meetings with lower educated individuals. Conversely, our estimates tend to suggest random meetings for low educated individuals.

We finally decompose the number of divorces into economic (labor market) factors and non-economic factors and show how their contributions evolve over time. The overall majority of divorces is driven by non-economic factors. Overall, less than 10% of divorces are due to labor market transitions of one spouse. However, the share of "labor market divorces" exhibits very interesting dynamics, it has increased by more than 20% since the mid 2000s. We take a granular view and investigate which types of heterogeneous couples have started to divorce more frequently in response to labor market transitions. Surprisingly, we find that positively sorted couples are the major contributor to this trend. In our sample, the largest and growing share of labor market divorces can be attributed to couples in which a previously unemployed highly educated female starts working. On the other hand, low education couples with a high likelihood of job loss contribute a shrinking number of labor market divorces. Both trends might be related to the booming German labor market. Low separation rates make marriages among low education individuals more stable. With high education, the option value of going on the marriage market with good employment perspectives can outweigh the value of staying married.

This Chapter is organized as follows. Section 3.2 introduces our data sources. Section 3.3 discusses descriptive empirical evidence relevant to our hypothesis and modeling choices. In Section 3.4, we introduce our marriage market model and solve it. Section 3.5 presents a model-based, structural decomposition of marriage and divorce flows into their sources. Section 3.6 concludes and further results as well as technical details can be found in the Appendix.

3.2 Data Sources

The empirical content of this Chapter is based on three sources of German micro data:

1. The German Microcensus, a household survey.

2. The Sample of Integrated Labor Market Biographies, a matched employer-employee data set.

3. The German marriage and divorce registers.

We briefly introduce and describe each data source before presenting and discussing the respective empirical results relevant to our analysis.

3.2.1 The German Microcensus (MC)

The German Microcensus (henceforth MC) is an annual survey that yields representative statistics on the German population and labor force. Data access is provided by the Research data center (FDZ) of the statistical offices of the German federal states.

It samples 1% of the population, consisting of all persons legally residing in Germany. It is the largest household survey in Europe. Participation is mandatory[133] and only a subset of questions can be answered on a voluntary basis.

The MC survey design relies on single-stage stratified cluster sampling. The primary sampling units are artificially delimited districts with a number of neighboring buildings. All households residing in these buildings are interviewed (principal residence).[134] Typically, one household member responds to the survey for all individuals living in the household, including the spouse, children, and other cohabitants if applicable. The survey program of the MC consists of a set of core questions that remains the same in each wave, covering general socio-demographic characteristics like marital status, education, employment status, individual and household income, and many other things.

Data Preparation

We restrict our attention to adults of ages 18 to 68 living in private households, either as singles (alone or with cohabitants) or as heterosexual married partners in the same household. Our definition of singles includes never-married, divorced, and widowed individuals. To reliably identify couples, we have to condition on legal marriage, since we cannot distinguish cohabitation of non-married couples from shared apartments in the earlier MC waves. Married couples are legally required to have the same principal residence, if they want to file a joint tax statement in order to enjoy the benefits of joint income taxation.

[133] According to the German Microcensus law, non-response may be fined.

[134] Since 1990 the average number of buildings has been 9, the targeted number of individuals is 15. Larger buildings are subdivided.

In principle we could use all MC waves from 1976 to 2013 for our analysis.[135] We carefully clean and properly weight the cross-sectional data sets in order to represent the German population. This enables us to study the composition of the German married and non-married population conditional on gender, education, and employment over time.

For our analysis we use the data starting from 1993. With this short sample we avoid complications related to the German reunification and, in turn, can analyze the German population as a whole. Another reason is that the SIAB data (see below) does not fully cover the East German labor market before 1993.

Unfortunately, the MC is not a panel. In contrast to Goussé et al. (2017), who use the British Household Panel Survey (BHPS), we cannot follow individuals over time and directly observe them switching between states of singlehood and marriage as well as employment and unemployment. This complicates connecting the model to German data. To tackle this issue, we categorize the individuals in each cross-section into 84 classes based on gender (male or female), education (low, medium, high), employment status (employed or unemployed), marital status and, if married, the education and employment status of the partner. We use these aggregated data to study the German marriage market.

Our theoretical sorting model is based on the presumption that the value of the spouse's labor in home and market production (labor productivity) is an important determinant of matching and separation decisions in the marriage market (in addition to non-economic forces like love and companionship). The empirical part of this analysis uses education and wages as a proxy for labor productivity. The MC includes detailed information on individuals' school education and vocational degrees. The way this information is collected in the survey varies across waves. To construct a time-consistent measure, we rely on the ISCED-1997 scale[136] and, accordingly, define three education categories:

[135]The MC has been conducted in West Germany since 1957 and in East Germany since 1991. The waves before 1976 contain no information on individual education. Before 1995 we have one wave every two or three years (1976, 1978, 1980, 1982, 1985, 1987, 1989, 1991, 1993). This is due to the fact that the MC was not always a yearly survey and, once it was, not all waves asked for education information. From 1995 onwards we have all waves at an annual frequency.

[136]The International Standard Classification of Education (ISCED) of the UNESCO intends to make education systems internationally comparable.

1. Low education: individuals, who at most graduated from lower secondary schools with or without a vocational degree (ISCED categories 1 & 2).

2. Medium education: individuals, who graduated from upper secondary schools with or without a vocational degree (ISCED categories 3 & 4).

3. High education: individuals with a tertiary degree (ISCED categories 5 & 6).

The second important dimension of individual heterogeneity in our analysis is the labor market status. Some details about job search behavior in the labor market are available in our data, but only for a subsample in the later years. In order to ensure that employment and unemployment are defined consistently over the whole time horizon, we pool unemployment and non-participation and do not subdivide the non-employed into job-seekers and inactive persons.

Our final MC data set (1993 - 2013) contains information on 8,426,756 individuals[137] of whom 47% are men and 53% women. 72% of men and 64% of women are married. Across all ages in our sample, from 18 to 68 years, the labor force participation rate is 62% for men and 46% for women, respectively.[138] In the period after German reunification, the individuals in our sample are representative of a roughly constant population of about 53 million adults.[139]

3.2.2 Sample of Integrated Labor Market Biographies (SIAB)

To construct labor market transition rates and wage measures, we rely on German matched employer-employee data. We use the Sample of Integrated Labor Market Biographies (henceforth SIAB) provided by the Institute for Employment Research (IAB) in Nuremberg, Germany.[140] These data cover the years 1975 to 2014. The SIAB is a

[137] The average number of observations per wave is 443,513.

[138] The participation-age profiles are hump-shaped. In the 2006 MC wave, participation for men is highest in the age bracket 35-39 (88%) and the maximum for women (77%) is reached for ages 40-44.

[139] We use the MC sample weights to scale our sample. The population increases somewhat after reunification and reaches a maximum of almost 55 million people in 2007, afterward it starts declining. The mean population between 1993 and 2013 is about 53 million people with a standard deviation of 1 million.

[140] We use the factually anonymous Sample of Integrated Labor Market Biographies (File: `SIAB_7514`). Data access is provided via a Scientific Use File supplied by the Research Data Center (FDZ) of the German Federal Employment Agency (BA) at the IAB, project no. 101693. See also Ganzer et al. (2016) for more details on the data set.

2 percent random sample drawn from the universe of employment and unemployment spells registered at the Federal Employment Agency (Bundesagentur für Arbeit) within the German social security system.[141] Individuals who are not subject to social insurance contributions i.e. self-employed workers, civil servants, students, and non-employed persons are not included in the sample.

Every observation in the SIAB corresponds to an employment or unemployment spell lasting between one day and one year in accordance with the reporting rules of the German social security system. This allows us to measure the employment status of an individual exact to the day. We observe individuals switching between different employers, employment and unemployment, as well as (with limitations) non-employment. For every employment spell we observe the nominal gross daily wage. In case of unemployment, the wage variable contains the amount of benefits paid to the worker. Since the SIAB is a sample of the labor force, we do not have the number of individuals not participating in the labor force. The SIAB is simply not representative for this part of the population. We are therefore unable to calculate transition rates out of inactivity and only use the transition rates from employment into unemployment and vice versa.[142]

The SIAB data contain a wide array of individual characteristics including gender, age, educational attainment, details on they type of employment (part/full-time, marginal/subject to social security) as well as occupation and some information on the employer. Unfortunately, only unemployment spells contain the information whether an individual is married or not. Since this is a non-representative fraction of the data, we cannot condition on marital status when estimating wage measures and labor market flows. The German social security data are collected at the individual level so they contain no information on the spouse that we could condition our estimations on.[143]

[141] The data consist of individuals which are characterized as follows: employed subject to social security, marginal part-time employed, unemployed benefit recipient, officially registered job-seeker, (planned) participant in programs of active labor market policy.

[142] Since we are interested in transition rates between labor market states, we have to divide the number of individuals changing the labor market status by the stock of individuals in the state from which the respective individuals exited.

[143] This limitation will be mitigated in a future version of the IAB data. Goldschmidt et al. (2017) develop a method for identifying married couples in the German matched employer-employee data by using confidential address and name data. The resulting couple identifier will be made available to other researchers and we will incorporate it in a future version. For now, we are forced to estimate wage distributions and labor market transitions without controlling for marital status and the partner's labor market attachment.

Data Preparation

In order to create a sample comparable to the MC data, we use data for both men and women in East and West Germany. The available age range in our SIAB sample, 17 to 62 years, is slightly narrower than in the MC. We conduct our estimations with the data from 1993 onwards, the same start year as in our primary MC sample. As mentioned before, the East German labor market was not completely covered by the institutional data sources before 1993. We drop all spells of marginal employment because they are not included in the data before 1999. In case an individual has multiple jobs at a given point in time we always define the highest paying job as the primary one and discard all other employment spells.

We use the SIAB education information to construct a variable that resembles the three ISCED categories in the MC data. Education information in the SIAB suffers from inconsistencies and missing values.[144] To solve this problem, we follow Fitzenberger et al. (2006) and impute missing and inconsistent values in the education variable. Spells with missing education after imputation are dropped.

Regarding wages, we start by deflating nominal gross daily wages using the German consumer price index with base year 2010. The German social security system tracks earnings only up to a certain limit. Beyond this threshold, further earnings are not taken into account for the calculation of social security contributions. We follow Dustmann et al. (2009) and impute the upper tail of the wage distribution by running a series of Tobit regressions, fitted separately for years, education levels, and age groups.

After data preparation, our SIAB sample consists of 18,623,471 employment spells from 968,215 individuals, 58% of which are male. The male share of all employment spells is similar (57%).

3.2.3 Marriage and Divorce Registers (MDR)

The marriage and divorce register data (henceforth MDR) originates from the German civil registry offices and the divorce courts. It is compiled by the Research Data Center of the statistical offices of the German federal states. The data are organized at the level

[144] Employers are not forced to report an employee's highest educational degree and might in some cases not even know about it, for instance when a worker switches occupations.

of the married couple and contain information on the exact birth dates of both spouses, the exact date of marriage, and, if applicable, the date of divorce. Additionally, the data contain various covariates including religion, citizenship, place of residence, number of children (before marriage and at the time of divorce), as well as who filed for divorce and the court's ruling. Unfortunately, there is no information about education, so we have to rely on age differences as a proxy for marital sorting.

Data Preparation

The marriage and divorce data are separate yearly files and we have access for the waves from 1991–2013 (marriage registers) and 1995–2013 (divorce registers). We clean the yearly files from missing and inconsistent observations. The waves of marriage and divorce data are then merged to get two big data sets. For one, we are interested in the aggregate yearly flows of marriages and divorces. We need these numbers for the structural decomposition of marriage and divorce flows in Section 3.5. We then proceed to link the two register data sets in order to estimate a series of hazard models.[145] 21.3% of the 17,166,070 marriages we observe ended in divorce. The rest survived until the end of our observation period. This data set enables us to study marriage duration conditional on the observable characteristics available to us, in particular age.

3.3 Empirical Results

3.3.1 Sorting in the Marriage Market

The aim of our analysis is to connect an equilibrium search model of the marriage market with heterogeneous men and women to German micro data. Our theoretical model allows for positive as well as negative assortative matching. A complementarity in the household production function induces homophily—the love of the same—and encourages positive assortative matching. Benefits from joint income taxation, which increase with the wage

[145] Due to the strict German data protection regulations it is not allowed to match the marriage and divorce observations at the level of the individual couple. We aggregate the marriage and divorce data to yearly cells containing the number of individuals with equal observable characteristics, particularly age. We then merge the cells based on the marriage year and "unpack" the linked data-set into individual marriage spells.

gap between spouses, encourages negative assortative matching. We use our MC and MDR data to document the extent of assortative matching on the German marriage market. Our model will have to match the empirical patterns we find.

Results from MC data

In this section we study how the homogeneity of married couples in terms of education, employment, and income has evolved over time. Overall, it has increased significantly. The correlations are depicted in Figure 3.1. In 1993, the correlation between spouses' education levels was just above 0.4. For comparison, this is somewhat lower in magnitude than what Greenwood et al. (2016) report for the U.S.[146] However, they have only two education categories, college and less than college, so the correlation should be somewhat higher than what we find with three education categories. The education correlation increases and reaches a maximum of 0.5 in 1999. Afterwards, it levels off and starts decreasing slightly in the second half of the 2000s. It remains well above its initial level, however. The leveling-off could be driven by supply factors, e.g. a limited number of highly educated individuals looking for a partner.[147] Also, the changing macroeconomic environment in Germany, especially the booming labor market, might have led to a change in matching patterns. We will pick up this thought in our structural empirical analysis in Section 3.5.

We do not observe a similar hump-shape for employment or income. The correlation between spouses labor market attachment (employed or unemployed) has increased steadily from just above 0.4 in 1993 to almost 0.6 in 2013. Today, it is more common that partners are either both employed or unemployed. In stark contrast to this observation, it is striking to see that the correlation between spouses' income levels has been negative until well into the 2000s. It turns positive in 2007 and increases further after 2009. This finding shows that the classical role model with one bread-winning individual, typically the husband, is still very dominant in Germany, even despite the high degree of education-based sorting. When both spouses are employed, however, earnings differences must be

[146]The long-run analysis of Greenwood et al. (2016) reveals a correlation of 0.41 in 1960 which rises to 0.52 in 2006.

[147]Figure A2.1 in the appendix shows the shares of education groups for men and women and their evolution over time.

Figure 3.1: Correlations of Education, Employment, and Income within Marriages

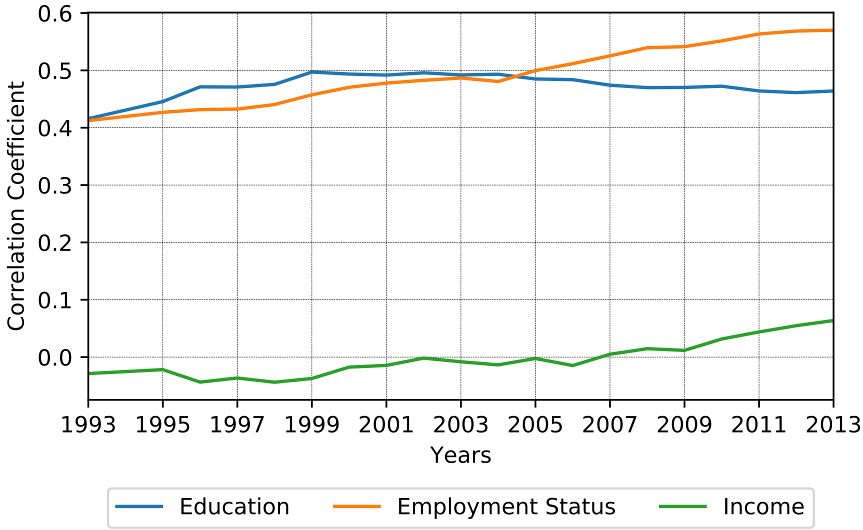

Note: Yearly Pearson correlation coefficients of the within-couple levels of education (three categories: low, medium, high), employment (two categories: employed, unemployed) and income (between 15 and 24 categories, depending on the wave). Source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations.

large in order to be consistent with an almost zero correlation between spouses' income. This is due to the fact that many working wives have a weak labor market attachment, work part-time or in marginal employment. One explanation for this observation is the German system of joint taxation for married couples, which provides strong disincentives for the secondary earner to increase labor supply.[148]

Figure 3.2 shows the extent of marriage market sorting in Germany in an alternative way. We now focus on education-based sorting only and use our MC aggregated data to calculate weighted population shares of men and women in partnerships of all possible education combinations. To set this into perspective, the population shares of all individuals by gender and education are depicted in Appendix Figure A2.1. For both men and women, the share of highly educated individuals has increased and the low education share has decreased. This trend is much more pronounced for women.

Each row in Figure 3.2 shows men (left Panel) and women (right Panel) for the same education category. For these individuals, we plot the share of marriages with partners

[148]See Gustafsson (1992) for a comparative empirical study of joint taxation and female labor supply in Germany and Sweden. The author exploits the Swedish switch to individual taxation in 1971 to identify the dampening effect of the high marginal tax rates on female labor supply under joint taxation.

Figure 3.2: Partner's Education by Education and Gender of Married Individuals

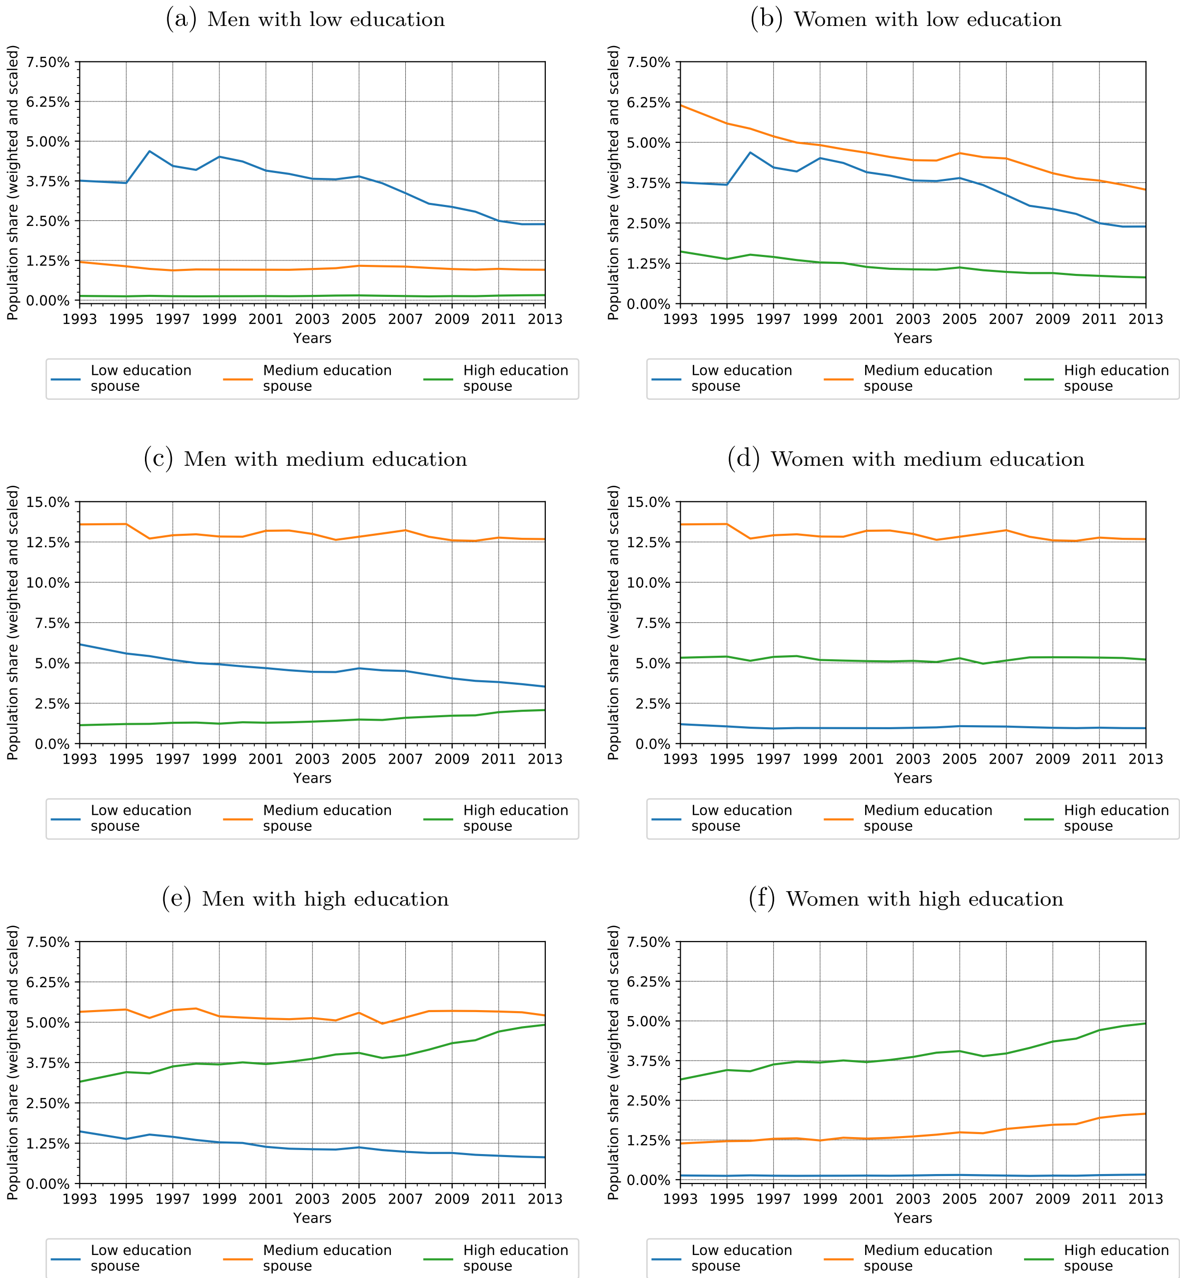

Note: Population shares are weighted and scaled using the MC sample weights. 100% on the y-axis corresponds to the full population, including married and unmarried individuals of both sexes. Source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations.

of each education category.[149] We see an increase of education-based sorting particularly for highly educated individuals. For men with university degrees (Panel 3.2e), the share

[149]Note that the lines representing marriages with partners of the same education level are identical by construction for men and women in each row (blue in the first, orange in the second, green in the third).

of marriages with highly educated women has almost doubled to 5% in 2013. At the same time, the share of marriages with women of the lowest education category has decreased significantly. The widening gap between the green and the blue line represents increasing homophily of highly-educated individuals. From the perspective of highly educated females (Panel 3.2f), the share of marriages with men of both high and medium education has increased. The medium education share starts increasing later and does not grow beyond 2%. Highly-educated women are almost never married to men from the lowest education group.

The increase of education-based sorting is not homogeneous across education groups. The first row of Figure 3.2 shows that the share of marriages among lowly educated individuals only increased somewhat in the beginning of our observation period but steadily decreased thereafter. The decline in the share of married low educated women across all groups of men is driven by the overall decrease in marriages of low educated women over the same time horizon. The same is true for men, albeit to a lesser extent. To show the increasing shares of singles, Figure A2.2 in the Appendix breaks down the population shares of the gender-education groups by marital status. The share of marriages between men and women of medium education (Panels 3.2c and 3.2d) is the highest in our sample with more than 12.5% and it is relatively stable over time. Sorting in this group is prevalent but not increasing. For the women, the share is constant across all groups. Men of medium education, however, become less likely to be married to women with low education and more likely to be married to highly educated women.

Results from MDR Data

We run a series of survival regressions on our linked marriage and divorce data in order to understand how the probability of divorce depends on the age difference in the married couple. Ideally, we would like to control for education. Since we do not have this information available in the MDR data, we use the age difference to proxy for the heterogeneity within married couples.

The correlation of the differences in education and age for married couples in the MC data is 0.66, indicating that spouses with small age differences also tend to be

Table 3.1: Hazard Ratio Estimation Results

Age difference	(1) Cox	(2) Weibull	(3) Exponential
2-5 years	1.075***	1.079***	1.074***
	(0.014)	(0.001)	(0.001)
5-10 years	1.197***	1.214***	1.193***
	(0.002)	(0.002)	(0.002)
10-15 years	1.298***	1.320***	1.291***
	(0.002)	(0.003)	(0.003)
>15 years	1.353***	1.370***	1.347***
	(0.004)	(0.004)	(0.004)
constant		0.0078***	0.0191***
		(0.000)	(0.000)
s (shape)		1.345***	
		(0.000)	
N	17166070	17166070	17166070

Note: Robust standard errors in parentheses. Hazard ratio relative to age difference < 2 years. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. s is the estimated shape parameter of the respective parametric distribution. Source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Marriage and Divorce Registers, 1991/1995-2013, own calculations.

homogeneous in terms of education.[150] To control for age differences, we classify couples as follows: the baseline group has an age difference of less than two years. Further, we group couples with more than two and up to five years, more than five and up to ten, more than ten and up to fifteen, and more than fifteen years of age difference. We estimate a semi-parametric Cox proportional hazard model as well as two parametric regressions, assuming Weibull and exponential distributions.

The estimated hazard ratios in Table 3.1 are precisely estimated and very close across specifications. Unanimously, we find that the hazard ratio is increasing in the age differ-

[150]This is calculated for couples with a maximum age difference of 11 years, representing more than 95% of couples in our data.

ence relative to the base category of less than two years. Couples with an age difference between two and five years have a 7.4% higher hazard rate compared to the baseline. For the remaining age-difference groups the hazard ratio relative to the baseline increases by 21.4%, 32%, and finally 37% for couples with the largest age differences of more than 15 years. We interpret these hazard ratios as strong evidence that—on average—the likelihood of a quick divorce increases when couples are very different in terms of age. Put differently, homogeneous couples have a higher probability of staying together for a long time.

Summary

We find an abundance of evidence for positive assortative matching on age, education, and employment status on the German marriage market. Only income was – probably due to the incentives caused by joint income taxation – negatively correlated before 2007 and became positively correlated after 2009. Still, homophily is prevalent and increasing. Particularly for education, however, we see that the tendency to sort is not uniform across groups. Sorting is increasing for the highly educated, rather constant for the middle group and even decreasing for men and women with low education. Additionally, increased sorting appears not just to be about matching with the right partner in the first place. The fact that hazard ratios increase in the age difference indicates that heterogeneous (i.e. non-sorted) couples, at least in terms of age, are on average more likely to divorce.

In the light of these empirical findings, the question remains *why* heterogeneous couples divorce more quickly. We suspect that the reasons for splitting up, economic and non-economic, must differ across couple types. This is where our structural model comes into play. It allows us to investigate this hypothesis by decomposing the flow of divorces into separations caused by idiosyncratic shocks and by economic reasons like a transition from employment to unemployment of one spouse. For this reason, we now turn to the labor market.

3.3.2 Wage Distributions

Following Goussé et al. (2017), we interpret individual labor productivity as the empirical counterpart of the heterogeneity of men and women in the model. We construct our measure of labor productivity from wage information in the SIAB data. We estimate the wage densities of men and women on the same domain in order to use them as the underlying type distributions when solving the structural model. In order to remove transitory components from wages and use both observable and unobservable determinants of individual labor productivity, we run a Mincerian wage regression including a person-fixed effect. Following Card et al. (2013), we regress log wages on a person-fixed effect, an unrestricted set of year dummies, and quadratic and cubic terms in age fully interacted with educational attainment:

$$\ln w_{it} = x_{it}'\gamma + \phi_i + r_{it}. \tag{3.1}$$

$\ln w_{it}$ denotes the log real daily wage of a worker i in year t, x_{it}' includes the time-varying observable characteristics, ϕ_i is a worker-fixed effect, and r_{it} is the residual . The explanatory power of this wage regression (adjusted R^2 of 72%) is high, albeit below the Card et al. (2013) benchmark (about 90%). There are two reasons for this difference: we include men and women from both East and West Germany, whereas Card et al. (2013) focus on men in West Germany in a smaller age bracket. Additionally, using the universe of social security records, they can include firm-fixed effects. We are unable to consistently estimate firm-fixed effects using the SIAB sample.[151] Wage differences across firms, however, are not the subject of our study.

Based on the estimated contributions of both observable and unobservable characteristics we predict individual wages as follows:

$$\hat{w}_{it} = \exp(x_{it}'\hat{\gamma} + \hat{\phi}_i). \tag{3.2}$$

We effectively remove the estimated residuals. The standard deviations of predicted wages is 0.615 for men and 0.611 for women, so male wages are somewhat more dispersed.

[151] See Andrews et al. (2008, 2012) on this topic.

Figure 3.3: Wage Distributions of Men and Women

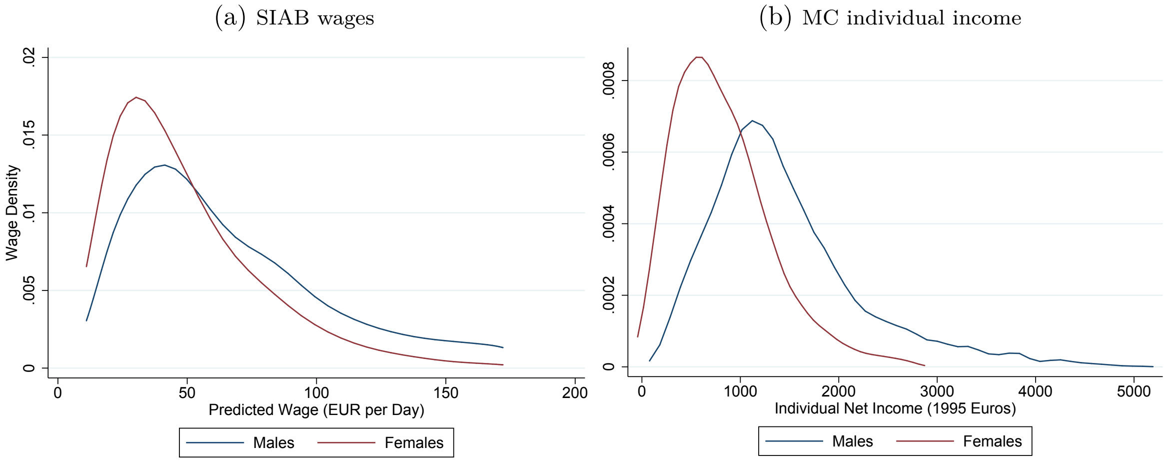

Note: Kernel densities of wages based on SIAB data (left Panel, kernel: Epanechnikov, bandwidth: 5) and individual income based on MC data (right Panel, kernel: Gaussian, bandwidth: 100). MC data source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations. SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB SUF 7514, 1993-2013, own calculations.

Next, we run kernel density estimations for both men and women using all wage observations between the 1st and the 99th percentile. Figure 3.3 depicts the resulting distributions of the predicted wage for men and women on a common domain. We compare the densities estimated from SIAB data to kernel density estimations of individual income from MC data (for both married and single individuals). The distributions share the same qualitative features, even though the wage and income information in the two data sets are very different.[152] Nevertheless, both male wage distributions have a fatter tail, male mean and median wages/income (see blue lines) lie to the right of the female ones (see red lines).

3.3.3 Labor Market Transitions

Our final set of empirical results concerns the job-finding and separation rates of men and women conditional on education. In our structural model, these rates are an important element because we suspect that labor market transitions trigger a sizable share of divorces.

[152]The MC contains only categorical information on individual and household income, which we make comparable across waves.

Table 3.2: Labor Market Transition Rates (%) by Gender and Education

Gender	Education	Job-Finding rate (UE)		Separation rate (EU)	
		Mean	STD	Mean	STD
All	-	4.579	0.158	0.639	0.083
Men	All	5.160	0.126	0.711	0.094
	Low	5.122	0.128	0.759	0.102
	Medium	5.689	0.328	0.453	0.033
	High	4.366	0.960	0.312	0.029
Women	All	3.880	0.260	0.552	0.067
	Low	3.654	0.272	0.565	0.069
	Medium	6.043	0.367	0.445	0.034
	High	5.773	1.019	0.450	0.048

Note: Mean and standard deviations of seasonally adjusted (X-13-ARIMA-SEATS) and HP-filtered ($\lambda = 900,000$) job-finding and separation rates by gender and education category. Rounded to three decimal places. SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB SUF 7514, 1993-2013, own calculations.

We estimate monthly transition rates between different labor market states following Jung and Kuhn (2014), additionally conditioning on gender and education. After data cleaning and wage/education imputations (as described above), we subdivide the spells in our data into periods of employment, unemployment, and inactivity. These data are then transformed into monthly slices. To get the transition rates, one simply has to count, for instance, how many employed individuals in a given month were unemployed in the previous month and divide this by the overall number of unemployed from the previous month. We compute the transition rates for unemployment to employment (UE) and employment to unemployment (EU) for the period of 1991-2014 and trim the series to 1993-2013. For one, to have a time frame similar to the MC data. Also, this allows for a burn-in of 24 months when computing flows. The resulting time series of labor market flows are highly seasonal and exhibit cyclical patterns. As we are seeking to connect an equilibrium search model to the data, we are (for now) not interested in the cyclical properties. We first apply the X-13-ARIMA-SEATS seasonal adjustment routine of the U.S. Census Bureau.[153] Afterwards, we use a Hodrick-Prescott filter with a penalty parameter of $\lambda = 900,000$ to remove the cyclical component from our monthly data. The statistics reported in the following are computed for the seasonally adjusted and filtered time series. For our analysis we take the yearly average of the monthly transition rates in order to make the data compatible with the MC-data.

Table 3.2 presents the means of the transition rate over the years 1993-2013. In the average month, 4.6% of unemployed workers find a job. The average job-finding rate is higher for men (5.2%) and lower for women (3.9%). There are sizable differences across education groups. Whereas low education women have the hardest time finding a job overall, medium and highly-educated women have higher job-finding rates than the corresponding men. Regarding separations, we note that on average men are more likely to separate in a given month (0.7%) than women (0.5%). This is not true, however, for all education groups. For both sexes, the overall level of the separation rate is mainly driven by low education individuals, who are still a sizable share of the German labor force (considering all age groups). Separation rates of better educated workers are lower and almost similar for men and women with medium education. Interestingly, however, highly-educated women have a 44% higher separation rate than similar men.

The monthly time series of job-finding and separation rates for men and women in all education groups are shown in Appendix Figure A2.3. One insight from the transition rates' changes over time is worth keeping in mind for our structural empirical analysis in Section 3.5: women with high and medium education have the highest job-finding rates overall and they have increased significantly since the first half of the 2000s.

3.4 The Model

We extend the frictional marriage and divorce model by Goussé et al. (2017), which itself is based on Shimer and Smith (2000) and Jacquemet and Robin (2013), by incorporating that single and married men and women change their labor market status l. For simplicity

[153] We use the R package "seasonal" developed by Sax (2017).

and compatibility with our data sources, we consider only employment (indexed by e) and voluntary or involuntary unemployment (indexed by u), i.e., $l \in \{e, u\}$. Besides the time varying labor market status, individuals differ in their level of education and the associated wages (or wage distributions). We capture the time-invariant heterogeneity of men and women by the indices i for men and j for women. We will be more specific when the model is taken to the data.

3.4.1 Preferences and Home Production

Individual utility depends on private consumption c, leisure h, and the public good q. The public good of a married couple depends on the time inputs of the spouses (d_i, d_j). We assume that these inputs are complements. The public good also depends on the types ij and on an idiosyncratic bliss shock z drawn from the cumulative probability distribution G, i.e., $q = zF^1_{ij}(d_i, d_j)$. A bliss shock arrives at the type specific rate δ_{ij}. The public good of a single is solely a function of the time input d_i, i.e., $q = F^0_i(d_i)$. Employment and unemployment enter the production of the public good indirectly by changing the time available for the production of the public good. The total time available for an individual is given by $T > 1$. We assume fixed working hours normalized to 1 for an employed worker. Thus, the time remaining for leisure h and household production d is given by,

$$d_i + h_i = T^l_i = \begin{cases} T - 1 & \text{for } l = e, \\ T & \text{for } l = u. \end{cases}$$

Working individuals will receive a wage w_i depending on their type. Unemployed individuals receive a replacement income bw_i if they are unemployed. Since working time cannot be adjusted at the intensive margin, private consumption of a single is given by

$$c_i = R^l_i = \begin{cases} w_i & \text{for } l = e, \\ bw_i & \text{for } l = u. \end{cases}$$

Following Goussé et al. (2017) we assume that households are subject to living cost C_{ij}, which is a function of the exogenous types ij. This fixed living cost is paid by the spouses through transfers, i.e., $t_i + t_j = C_{ij}$, which are determined by Nash-Bargaining.

The respective private consumption of a spouse is given by

$$c_i = R_i^l - t_i = \begin{cases} w_i - t_i \quad \text{for } l = e, \\ \\ bw_i - t_i \text{ for } l = u. \end{cases}$$

The flow utility of a single individual depends on the public good $F_i^0(d_i)$, consumption (equal income) $c_i = R_i^l$, and leisure $h_i = T_i^l - d_i$, i.e.,

$$u_i^l(d_i) = F_i^0(d_i)\left[R_i^l + T_i^l - d_i\right]. \tag{3.3}$$

The bliss variable z for singles is normalized to unity. The flow utility of a married individual depend on her own labor market status $l \in \{e, u\}$ and the labor market status of the partner $-l \in \{e, u\}$. It is assumed to have the following form,

$$u_i^{l,-l}(t_i, d_i, d_j|z) = zF_{ij}^1(d_i, d_j)\left[R_i^l + \frac{\iota}{2}\left(R_j^{-l} - R_i^l\right)^2 - t_i + T_i^l - d_i\right]. \tag{3.4}$$

The term $\frac{\iota}{2}\left(R_j^{-l} - R_i^l\right)^2$ takes into account the net income gain from joint taxation of couples. The utility depends on the time (d_i, d_j) devoted to public good production, the contribution t_i to the fixed cost of living, and the own and the partners labor market status l and $-l$ (via R_i^l, R_j^{-l} and T_i^l). Given the labor market status the time inputs to public good production $\{d_i, d_j\}$ are chosen to maximize joint surplus of the match and the transfers $\{t_i, t_j\}$ to ensure that each individual gets its respective fraction of the surplus.

3.4.2 Marriage Formation and Renegotiations

The present values of a marriage for a female (and the male respectively) depend on her own labor market status $l \in \{e, u\}$ and the labor market status of the partner $-l \in \{e, u\}$. We denote the flow utility of the married female for the optimal choices of $\{d_i, d_j, t_i, t_j\}$ by $u_j^{l,-l}$, where $(l, -l) \in \{(u, u), (u, e), (e, u), (e, e)\}$. The following Bellman equation,

$$\begin{aligned} rV_j^{l,-l} &= u_j^{l,-l} + \delta_{ij}\int\left[\max\left[V_j^l, V_j^{l,-l}(z')\right] - V_j^{l,-l}\right]dG(z') \\ &\quad +\tau_j(l)\left[\max\left[V_j^{l'}, V_j^{l',-l}\right] - V_j^{l,-l}\right] + \tau_i(-l)\left[\max\left[V_j^l, V_j^{l,-l'}\right] - V_j^{l,-l}\right], \end{aligned} \tag{3.5}$$

describes the present values of marriage. $\tau_j(l)$ denotes the exogenous transition rate from the current labor market status $l \in \{e, u\}$ into the labor market status $l' \neq l$ for an individual of type j. The last term in the Bellman equations describes the labor market transition of the partner of type i. Individuals do not make long-run commitments. If a labor market transition or a bliss shock occurs, both partners renegotiate their contributions (d_i, d_j) to the public good and the transfers (t_i, t_j) to finance the fixed living costs C_{ij}. If the outside option is higher than the surplus from remaining married, then the couple divorces. In the renegotiations $\{d_i, d_j, t_i, t_j\}$ are chosen such that the Nash-Product,

$$\left[V_j^{l,-l} - V_j^l\right]^{1-\beta} \left[V_i^{l,-l} - V_i^l\right]^{\beta}, \tag{3.6}$$

is maximized subject to the feasibility constraint $t_i + t_j = C_{ij}$ and the participation constraint,

$$V_j^{l,-l} - V_j^l \geq 0, \text{ and } V_i^{l,-l} - V_i^l \geq 0,$$

where V_i^l (V_j^l) is the outside option of the single male (female) individual. The present value of being a single female satisfies the Bellman equation,

$$rV_j^l = u_j^l + \lambda_{ij} \iiint \left[V_j^{l,-l'}(z') - V_j^l\right] W_{ij}^{ll}(z')\, dG(z')\, s(i, l')\, didl' + \tau_j(l) \left[V_j^{l'} - V_j^l\right]. \tag{3.7}$$

The maximized flow utility of a single is denoted by $u_j^l = \max_{d_i} u_i^l(d_i)$. λ_{ij} denotes the type specific meeting rate of a potential partner. A meeting only results in a marriage if the joint surplus is positive. The respective willingness to marry (or stay in the marriage) is denoted by the index $W_{ij}^{ll}(z)$. If a pair is willing to marry (stay together) then $W_{ij}^{ll}(z) = 1$, and zero otherwise. The willingness to marry depends on the types and the labor market status (the first l corresponds to the male's labor market status, the second to the female's) as well as the bliss shock z. We denote by α_{ij}^{ll} the probability that $W_{ij}^{ll}(z) = 1$.

The marriage surplus is defined as the gain from marriage for the female and the male of type ij and labor market status ll, where the first l corresponds to the male's labor market status, the second to the female's, i.e.,

$$S_{ij}^{ll} \equiv \left[V_i^{l,-l} - V_i^l\right] + \left[V_j^{l,-l} - V_j^l\right]. \tag{3.8}$$

Using the first order conditions for the transfers and the time devoted to public goods production $\{d_i, d_j, t_i, t_j\}$ – derived in Appendix A.2 – the surplus for any type ij and employment status ll is given by,

$$\begin{aligned}
& \left(r+\delta_{ij}+\tau_i(l)+\tau_j(l')\right) S_{ij}^{ll'}(z) \qquad (3.9)\\
= \; & u_{ij}^{ll'}(z)+\delta_{ij}\int \max\left[S_{ij}^{ll'}(z'),0\right] dG(z') \\
& +\tau_i(l)\max\left[S_{ij}^{l'l'}(z),0\right]+\tau_j(l')\max\left[S_{ij}^{ll}(z),0\right] \\
& -\lambda_{ij}(1-\beta)\iiint \max\left[S_{ij}^{ll''}(z'),0\right] dG(z')\, s(j,l'')\, dj dl'' \\
& -\lambda_{ij}\beta\iiint \max\left[S_{ij}^{l''l'}(z'),0\right] dG(z')\, s(i,l'')\, di dl'',
\end{aligned}$$

where $u_{ij}^{ll}(z)$ denotes the maximized joint flow surplus of both partners, i.e.,

$$\begin{aligned}
u_{ij}^{ll'}(z) &\equiv u_i^{l,-l'}+u_j^{l',-l}-u_i^l-u_j^{l'} \qquad (3.10)\\
&= z\kappa_{ij}\left[W_{ij}+\psi_{ij}^{ll'}\right]^{\kappa}-\kappa_i\left[w_i+\psi_i^l\right]-\kappa_j\left[w_j+\psi_j^{l'}\right]
\end{aligned}$$

with

$$\kappa=1+K_f^1+K_m^1, \qquad \kappa_{ij}=Z_{ij}K, \qquad \kappa_i=\left(K_i^0\right)^{K_i^0}, \qquad \kappa_j=\left(K_j^0\right)^{K_j^0},$$
$$\psi_{ij}^{ll'}=-C_{ij}+T_i^l-D_i^1+T_j^{l'}-D_j^1, \qquad \psi_i=T_i^l-D_i^0-K_i^0, \qquad \psi_j=T_j^{l'}-D_j^0-K_j^0,$$
$$W_{ij}=w_i+w_j+\iota\left(w_i-w_j\right)^2.$$

The maximized joint flow surplus $u_{ij}^{ll'}(z)$ is strictly increasing in z. This ensures that also the surplus functions are strictly increasing in z. The cutoff bliss values $z_{ji}^{ll'}$ for $ll' \in \{ee, ue, eu, uu\}$ are defined such that the surplus is equal to zero, i.e., $S_{ji}^{ll'}\left(z_{ji}^{ll'}\right)=0$. Since $u_{ji}^{ll'}(z)$ is increasing in z it follows that $S_{ji}^{ll'}(z)>0$ for $z>z_{ji}^{ll'}$. This allows us to write the probability $\alpha_{ij}^{ll'}$ that a couple of type ij and labor market status ll' is willing to marry upon meeting as,

$$\alpha_{ij}^{ll'}=\left(1-G\left(z_{ij}^{ll'}\right)\right). \qquad (3.11)$$

3.4.3 Steady State Flows and Measures

We denote by $m(i,j,l,l)$ the number of married couples of type ij and labor market status ll. The number of single males (females) of type i (j) and labor market status l is denoted by $s(i,l)$ ($s(j,l)$). The number of married couples of type ij and labor market status ll divorce, if a bliss shock reduces the bliss value below z_{ij}^{ll}, or change into another labor market status ll if one partner changes her/his employment status (at rate $\tau_i(l) + \tau_j(l)$). The inflow, i.e., the number of new marriages of type ij and labor market status ll formed, is given by $\lambda_{ij}\alpha_{ij}^{ll} s(i,l) s(j,l)$, where α_{ij}^{ll} denotes the probability that a couple of type ij and labor market status ll is willing to marry upon meeting. There are additional inflows into the group $m(i,j,l,l)$ from couples of labor market status $m(i,j,l',l)$ and $m(i,j,l,l')$. The probability that a couple stays together after a change of the labor market status from $l'l$ to ll is equal to 1 if $z_{ij}^{ll} \leq z_{ij}^{l'l}$ and equal to $\alpha_{ij}^{ll}/\alpha_{ij}^{l'l} < 1$ if $z_{ij}^{ll} > z_{ij}^{l'l}$, i.e., equal to $\min\left[\left(\alpha_{ij}^{ll}/\alpha_{ij}^{l'l}\right), 1\right]$. We therefore get,

$$\begin{aligned} & \left[\delta_{ij}\left(1-\alpha_{ij}^{ll}\right) + \left(\tau_i(l) + \tau_j(l)\right)\right] m(i,j,l,l) \\ = \ & \lambda_{ij}\alpha_{ij}^{ll} s(i,l) s(j,l) \\ & + \tau_i(l') \min\left[\left(\alpha_{ij}^{ll}/\alpha_{ij}^{l'l}\right), 1\right] m(i,j,l',l) \\ & + \tau_j(l') \min\left[\left(\alpha_{ij}^{ll}/\alpha_{ij}^{ll'}\right), 1\right] m(i,j,l,l'). \end{aligned} \tag{3.12}$$

Let us now consider the flow equations for the respective single groups. The outflow of a single male of type i with labor market status l is given by the rate at which the individual marries with single female of type j with labor market status l'', i.e., the rate $\lambda_{ij}\alpha_{ij}^{ll''} s(j,l'')$, plus the rate at which the single male changes her/his labor market status, i.e., the rate $\tau_i(l)$. The inflow is given by the rate at which the single males with the opposite labor market status l' change their status (at rate $\tau_i(l')$) plus the rate at which the respective marriages break up. This happens when a bliss shock occurs (at rate $\delta_{ij}\left(1-\alpha_{ij}^{ll''}\right) m(i,j,l,l'')$) or when the married male of type i or the married female of type j changes the labor market status (at rates $\widetilde{\tau}_i(l') \max\left[1-\left(\alpha_{ij}^{ll''}/\alpha_{ij}^{l'l''}\right), 0\right] \widetilde{m}(i,j,l',l'')$ or

$\widetilde{\tau}_j\left(l''\right)\max\left[1-\left(\alpha_{ij}^{ll'}/\alpha_{ij}^{ll''}\right),0\right]\widetilde{m}\left(i,j,l,l''\right))$. Formally,

$$\begin{aligned}
&\left[\iint\lambda_{ij}\alpha_{ij}^{ll''}s\left(j,l''\right)djdl''+\tau_i\left(l\right)\right]s\left(i,l\right) \quad (3.13)\\
=\ &\tau_i\left(l'\right)s\left(i,l'\right)+\iint\delta_{ij}\left(1-\alpha_{ij}^{ll''}\right)m\left(i,j,l,l''\right)djdl''\\
&+\iint\widetilde{\tau}_i\left(l'\right)\max\left[1-\left(\alpha_{ij}^{ll''}/\alpha_{ij}^{l'l''}\right),0\right]\widetilde{m}\left(i,j,l',l''\right)djdl''\\
&+\iint\widetilde{\tau}_j\left(l''\right)\max\left[1-\left(\alpha_{ij}^{ll'}/\alpha_{ij}^{ll''}\right),0\right]\widetilde{m}\left(i,j,l,l''\right)djdl'',
\end{aligned}$$

because $\max\left[1-\left(\alpha_{ij}^{ll''}/\alpha_{ij}^{l'l''}\right),0\right]=1-\min\left[\left(\alpha_{ij}^{ll''}/\alpha_{ij}^{l'l''}\right),1\right]$. To get number of singles of a certain type and labor market status we can use the aggregate labor market transitions, e.g.,

$$\tau_i(l)s(i,l)+\tau_i(l)\iint m\left(i,j,l,l''\right)djdl''=\tau_i(l')s(i,l')+\tau_i(l')\iint m\left(i,j,l',l''\right)djdl'',$$

and the market clearing conditions for the different types of males and females, e.g.,

$$\begin{aligned}
n\left(i\right) &= s\left(i,l'\right)+s\left(i,l\right)+\iiint m\left(i,j,l,l''\right)djdl''dl,\\
n\left(j\right) &= s\left(j,l'\right)+s\left(j,l\right)+\iiint m\left(i,j,l,l''\right)didl''dl.
\end{aligned}$$

Substituting and rearranging then implies the following formula for singles of type i and labor market status l,

$$s\left(i,l\right)=\frac{\tau_i(l')}{\tau_i(l)+\tau_i(l')}n\left(i\right)-\iiint\frac{\tau_i(l)}{\tau_i(l)+\tau_i(l')}m\left(i,j,l,l''\right)djdl''dl. \quad (3.14)$$

3.4.4 Equilibrium

The equilibrium is characterized by a set of surplus functions $S_{ij}^{ll}\left(z\right)$, cutoff bliss values z_{ij}^{ll}, and joint distributions of married couples $m\left(i,j,l,l\right)$ for each type ij and labor market status ll. We compute the equilibrium in the following way: Given a set of initial conditions, the cutoff bliss values z_{ij}^{ll} determine $\alpha_{ij}^{ll}\equiv\left(1-G\left(z_{ij}^{ll}\right)\right)$. Given α_{ij}^{ll} we can use equations (3.12) and (3.14), i.e., a set of four equations for $m\left(i,j,l,l\right)$ for each $ll\in\{ee,ue,eu,uu\}$ and a set of two equations determining $s\left(i,l\right)$ and $s\left(j,l\right)$ for each $l\in\{e,u\}$, respectively, to compute $s\left(i,l\right)$ and $s\left(j,l\right)$. The number of singles $s\left(i,l\right)$ and

$s(j,l)$ of type i (j) and labor market status l (l) determine the surplus functions $S_{ij}^{ll}(z)$ given by equation (3.9) for all types ij and labor market status ll. The bliss values z_{ij}^{ll} for all types ij and labor market status combinations ll are then pinned-down at a value such that the respective surplus is zero, i.e., $S_{ij}^{ll}\left(z_{ij}^{ll}\right) = 0$. The problem involves alternating between solving the two fixed-point systems of $S_{ij}^{ll}(z)$ and z_{ij}^{ll} until convergence. Appendix A.2 describes in detail how the fixed point systems are solved numerically.

3.4.5 Model Solution

Figure 3.4: z Cutoffs

(a) z_{ij}^{ee} (b) z_{ij}^{ue} (c) z_{ij}^{eu} (d) z_{ij}^{uu}

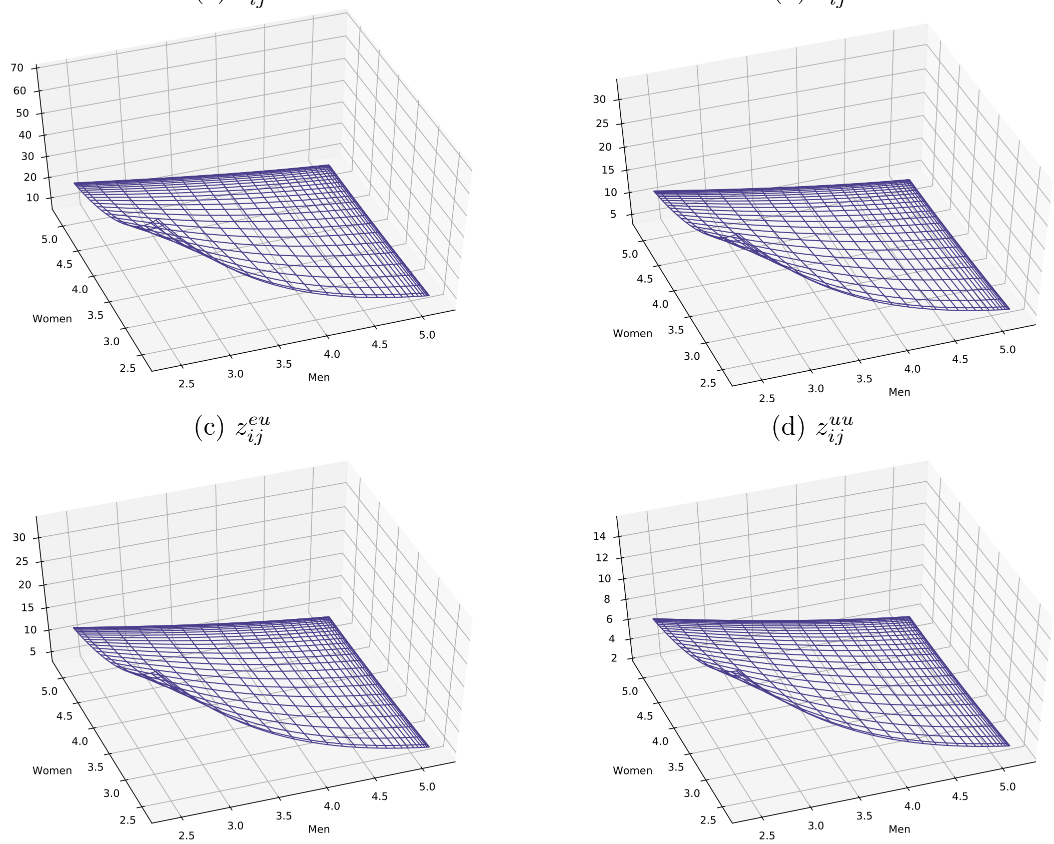

We solve the model on a Chebyshev grid with 50×50 nodes. We use the empirical wage distribution functions estimated for men and women in Section 3.3.2 to set up the underlying distributions of men and women, $n(i)$ and $n_(j)$. Given a first parametrization (see Appendix A.2), the model's stationary equilibrium exhibits a number of interesting properties. Here, we focus on the distribution of married couples across type and employment status combinations because this is what we see in our data.

Figure 3.5: Marriage Probabilities α_{ij}^{ll}

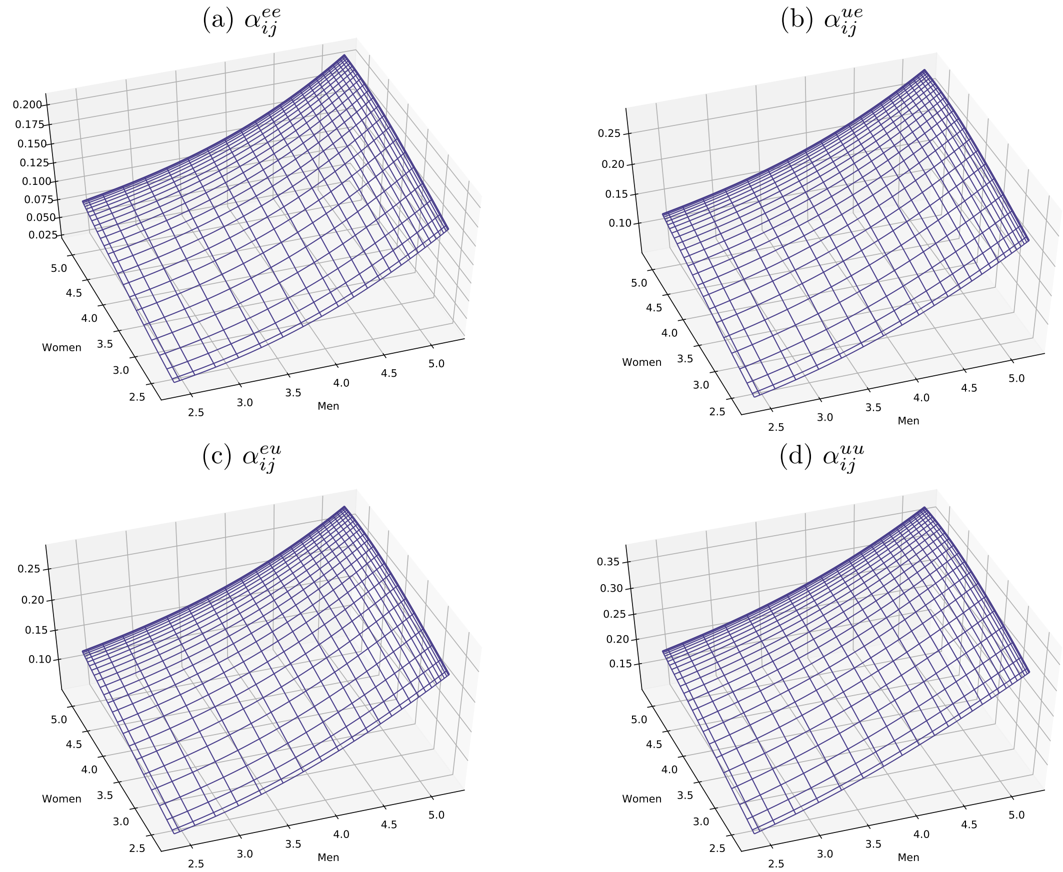

Figure 3.4 shows the minimum realizations of the bliss shock a couple needs to draw upon meeting in order to form a marriage. This value is highest for the lowest types of men and women. The higher the types of the man and woman who are meeting the lower is the z they need to draw in order to have a positive marriage surplus. This pattern holds across all employment status combinations, the levels, however, are very different. The outside option of continued search in the marriage market is higher for employed individuals, hence the necessity to draw higher z values in order to compensate both parties. Meetings between unemployed men and women need the lowest z overall to result in wedlock. The model-generated marriage probabilities α_{ij}^{ll} are the mirror image of the z values. They are depicted in Figure 3.5. Again, the general pattern is the same across all employment status combinations, the marriage probability increased in both partners' types. Our structural estimation in the next Section will enable us to compare the model generated marriage probabilities to their empirical counterparts.

3.5 Decomposition of Marriage and Divorce

We now take our structural model of the marriage market to the data. This connection allows us to go beyond the descriptive analysis in Section 3.3. We use our model to uncover the different contributing factors to matching and separation decisions at the individual level from the data. This allows us to decompose and explain the observed aggregate dynamics of marriage and divorce. We first highlight the relevant matching mechanisms in our model and then describe how to identify their relative importance from the data.

According to our model, matching in the marriage market has two components: The meeting rate, λ_{ij}, determines the likelihood of meeting a certain type of partner in the frictional marriage market. The ij dependence resembles the idea that individuals, depending on their type, have different probabilities to meet with heterogeneous members of the other sex. This is likely to occur, since many couples meet at education institutions or at the workplace. The second component of matching decisions are the acceptance probabilities, α_{ij}. Conditional on meeting, they capture the likelihood of wedlock. It differs across ij combinations because, according to our model, the option value of continued search for another partner may dominate forming the marriage. The willingness to marry α_{ij} and the meeting rate λ_{ij} may also differ across labor market statuses.

Regarding divorce, two things can happen: First, a negative update of the match-specific bliss shock z occurs, decreasing home production and flow utilities. This may drive the marriage surplus below zero and lead to a divorce. Second, as we have emphasized throughout, labor market transitions may trigger divorces. These "labor market divorces" may happen for two reasons: First, an employed spouse becomes unemployed. Depending on the combination of types in the couple, the drop in household income may outweigh the increase in home production (via the time input) and, hence, decrease utility flows and lead to divorce. Second, a previously unemployed spouse may find it optimal to divorce after finding a new job. Theoretically, the outside option of starting over in the marriage market as an employed person can dominate the option value of staying in the current match.

Looking at the data through the lens of our model, we now let the data decide which channels drive marriage and divorce in Germany and analyze how their respective contributions have evolved over time.

3.5.1 Meetings & Marriages

Using our three sources of micro data (MC, SIAB, MDR), we estimate the simultaneous flow equation system from our model, summarized in Equation (3.13), using variation across time and single/couple types. The details of our estimation procedure are included in Appendix A.3. In short, we construct the empirical counterparts of the (joint) distributions of singles and married couples from MC data and define them as follows: $\widetilde{s}^{l}_{it} = s\left(i, l\right)$ and $\widetilde{m}^{ll}_{ijt} = m\left(i, j, l, l\right)$. The observed labor market transition rates $\widetilde{\tau}^{l}_{it} = \tau_i\left(l\right)$ are the second data input. Due to the nature of our aggregate data, variation is limited and we need to discipline the estimation. We derive a large set of equality, inequality, and non-linear constraints from our model and impose them on the parameters to be estimated. In particular, our constraints guarantee that all estimated values of $\hat{\alpha}^{ll}_{ij}$, the estimated matching probabilities, lie in the unit interval. Given the flow equations and the constraints, we estimate a set of composite parameters using a non-linear least squares method.[154] It is then possible to back out the model parameters from the estimated composite parameters.

Table 3.3 shows our estimates of the four α^{ll}_{ij} matrices, one for each combination of the spouses' labor market status ($ll \in \{ee, ue, eu, uu\}$). Our descriptive analysis of education-based sorting in Section 3.3.1 has revealed that the tendency to sort is not uniform across education groups. Based on our model, we can now refine this statement by additionally taking into account the estimated matching probabilities across types and labor market statuses of the spouses.

In each Panel, the horizontal dimension of the table represents the female education types (j) and the vertical dimension the male types (i). In many cases, the restrictions we impose on the data are binding; we get matching probabilities of one. Meetings of two employed singles (Panel 3.3a) have the overall highest probabilities of ending in wedlock. These couples also contribute to marriage market sorting. Estimated probabilities are

[154] We rely on the excellent lmfit package for Python by Newville et al. (2014).

Table 3.3: Estimates of Matching Probability α_{ij}^{ll}

(a)

$\hat{\alpha}_{ij}^{ee}$		j low	medium	high
	low	0.96	0.62	0.53
i	medium	0.99	0.72	1.00
	high	0.53	1.00	0.78

(b)

$\hat{\alpha}_{ij}^{eu}$		j low	medium	high
	low	0.01	0.79	0.89
i	medium	0.88	1.00	1.00
	high	0.89	0.01	1.00

(c)

$\hat{\alpha}_{ij}^{ue}$		j low	medium	high
	low	0.45	0.14	0.05
i	medium	0.20	1.00	1.00
	high	0.05	1.00	0.24

(d)

$\hat{\alpha}_{ij}^{uu}$		j low	medium	high
	low	0.00	0.05	0.45
i	medium	0.06	1.00	0.07
	high	0.45	0.00	1.00

Note: Estimated marriage probabilities as derived from our model for men and women with three education categories (low, medium, high) and two employment categories (employed, unemployed). Source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations. SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB, 1993-2013, own calculations.

high on the main diagonal and low for the low/high and high/low combinations. Probabilities are one for medium/high and high/medium couples. Panel 3.3b contains marriage probabilities of employed men and unemployed women of all types. The matching probability of two low education individuals with this employment status combination is almost zero, the other two values on the main diagonal, however, are one. Hence, these combinations also sort positively. Interestingly, the combination of a high-type man and a medium-type woman has a very low estimated matching probability while in the inverse case, high-type women (unemployed) and medium-type men (employed) are very likely to match. Panel 3.3c shows the case of an unemployed man and an employed woman. For two highly-educated individuals, this labor market status combination is not likely to lead to marriage (24%). Marriage is ensured conditional on meeting, however, for the medium/medium, high/medium and medium/high type combinations, what is in line with positive sorting. Unions involving low type individuals have matching probabilities which are monotonically decreasing in the partner's education type. Finally, we show in Panel 3.3d that two unemployed singles who meet are ensured to mate if both partners have medium or high education. All other probabilities are very low, with the exception of the two high-low combinations. Both in Panels 3.3b and 3.3d the high matching probabilities for the high-low combinations can be rationalized with the prevalence of joint taxation in Germany, which creates incentives for negative sorting.

While the overall picture of a tendency to sort positively is confirmed across the alphas, our prior from the descriptive analysis can be updated. It is true that the tendency to sort varies across education types but there is also interesting variation across employment status combinations. While two employed individuals will get married with a very high probability in all education categories, dating couples with an unemployed man are very unlikely to match if one of the potential partners has low education. Unemployment of the woman, however, seems to matter much less, matching probabilities are high almost everywhere.

Two unemployed individuals are least likely to get married overall but they still contribute to positive sorting, the upper two main diagonal elements are one. Note also that the solution of our model depicted in Section 3.4.5 can match the overall picture in the data of higher matching probabilities for higher education types but it cannot reproduce the heterogeneity across education cells. More work is needed to calibrate the model and to map the finer wage grid on which the theoretical model is solved into the education categories we have as a proxy in our data.

Given the estimates for α_{ij}^{ll} and our constraints, we can go one step further and calculate the estimated number of meetings per month across partner types. The number of meetings per month for a given education-pair ij with labor market status combination ll is given by multiplying the mean of the respective single shares, $\bar{s}_i^l$ and $\bar{s}_j^l$, with the estimated $\hat{\lambda}_{ij}$ parameter. Comparing the number of meetings across education groups and labor market status allows us to analyze whether search in the marriage market is random or directed.

Table 3.4 presents the estimated number of meetings per month again across all marriage and labor market type combinations. We note that for employed singles (Panel 3.4a) the number of meetings are highest for medium and high type individuals, with the surprising exception that high type singles have a rather low number of meetings. One has to take into account, however, that both male and female singles with high education are rare in the marriage market. Panel 3.4b reveals that unemployed low type women have a high number of meetings with all types of employed men. The ranking of matching

Table 3.4: Estimates of Meeting Rates λ_{ij}^{ll}

(a)

$\hat{\lambda}_{ij}^{ee}\bar{s}_i^e\bar{s}_j^e$		j low	medium	high
	low	1.33e-04	5.50e-05	4.19e-04
i	medium	8.51e-04	2.80e-03	1.90e-01
	high	5.46e-04	3.38e-02	8.24e-04

(b)

$\hat{\lambda}_{ij}^{eu}\bar{s}_i^e\bar{s}_j^u$		j low	medium	high
	low	2.45e-02	2.44e-05	1.39e-04
i	medium	1.01e-03	7.73e-04	9.55e-03
	high	1.59e-03	1.00e-04	7.29e-05

(c)

$\hat{\lambda}_{ij}^{ue}\bar{s}_i^u\bar{s}_j^e$		j low	medium	high
	low	1.56e-04	6.44e-05	4.91e-04
i	medium	4.58e-04	1.51e-03	2.83e-03
	high	2.53e-04	5.42e-04	3.84e-04

(d)

$\hat{\lambda}_{ij}^{uu}\bar{s}_i^u\bar{s}_j^u$		j low	medium	high
	low	2.88e-02	2.86e-05	1.63e-04
i	medium	5.42e-04	1.31e-04	5.19e-03
	high	7.37e-04	4.65e-05	5.94e-06

Note: Estimated meeting rates for men and women with three education categories (low, medium, high) and two employment categories (employed, unemployed), multiply by the respective share of single in the population. Source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations. SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB, 1993-2013, own calculations.

probabilities for the same cells, however, was opposite (recall Table 3.3). There are a lot of meetings between low-type men and low type women in the (eu) category but only a small fraction of them ends in marriage. Between low type women and medium/high type men, the number of meetings are lower by a factor of 10 but, conditional on meeting, marriage is very likely. This pattern seems to be consistent with random search of low type individuals in the marriage market, that is, search is not directed towards partner types with a high likelihood of marriage. For the high/medium type combinations, however, search appears to be directed as high values of $\hat{\alpha}_{ij}^{eu}$ coincide with a high number of meetings $\hat{\lambda}_{ij}^{eu}\bar{s}_i^e\bar{s}_j^u$. This pattern of random search of low type individuals and directed search of medium and high type individuals is repeated in Panel 3.4c. It seems to be a consistent feature of the German marriage market. Panel 3.4d, however, is a special case. For two unemployed individuals, the highest number of meetings (low/low) essentially never leads to marriage. Conversely unemployed high type singles are very unlikely to meet but very likely to marry.

Figure 3.6: Divorce Rate and Share of Idiosyncratic Divorces

(a) Divorce rate

(b) Share of "Labor Market Divorces"

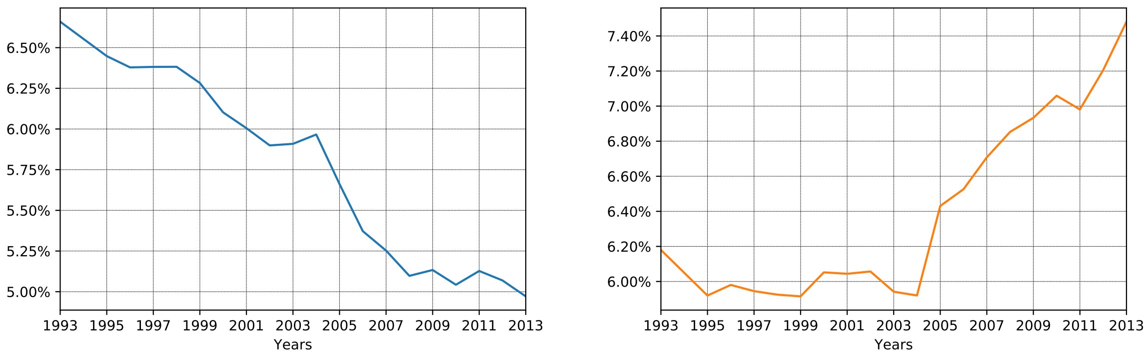

Note: MC data source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations. SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB SUF 7514, 1993-2013, own calculations.

3.5.2 Divorces

Our model implies that the aggregate flow of divorces must be consistent with the following aggregated flow equation:

$$\begin{aligned}\widetilde{\Delta}_t \;=\;& \delta \iiiint \left(1-\alpha_{ij}^{l''l}\right) \widetilde{m}_{ijt}^{l''l} di dj dl'' dl \\ &+ \iiiint \widetilde{\tau}_{it}^{l''} \max\left[1-\left(\alpha_{ij}^{l'l}/\alpha_{ij}^{l''l}\right),0\right] \widetilde{m}_{ijt}^{l''l} di dj dl'' dl \\ &+ \iiiint \widetilde{\tau}_{jt}^{l} \max\left[1-\left(\alpha_{ij}^{l''l'}/\alpha_{ij}^{l''l}\right),0\right] \widetilde{m}_{ijt}^{l''l} di dj dl'' dl.\end{aligned} \tag{3.15}$$

$\widetilde{\Delta}_t$ is the aggregate number of divorces in the data. By plugging in our estimated $\hat{\alpha}_{ij}^{ll}$ matrices on the RHS of equation (3.15) we can decompose the divorce flow into the shares of divorces caused by idiosyncratic shocks (first term on the RHS), the share caused by male labor market transitions (second term on the RHS), and the share caused by female labor market transitions (third term on the RHS).

We are interested in the share of divorces induced by labor market transitions, that is, one spouse transferring from employment to unemployment or vice versa. For succinctness, we refer to them as "labor market divorces". We can further differentiate these divorces for our four types of couples by labor market status combination before the transition and the underlying heterogeneity of types (education).

As a first step, we look at the aggregate divorce rate, see the left Panel of Figure 3.6. The overall number of divorces has declined significantly during our period of observation, it fell from 7% to below 5%. The right Panel shows the share of labor market divorces in all separations. The majority of divorces is not triggered by labor market shocks. According to our theory, the “residual”, between about 92% and 94% of all divorces, are triggered by an update of the bliss shock z. Remarkably, however, the share of labor market divorces has increased over time, against the overall trend of a declining divorce rate. The share was quite stable at around 6% until 2004 and started increasing rapidly thereafter, with a small correction in 2011, reaching almost 7.5% in 2013.

To understand which couple types contributed to the increasing share of labor market divorces, we now further differentiate it according to our gender-education-marriage-type cells. We look at married men and women in marriages with all four employment status combinations and across all 9 education types. The data for all cells can be found in Appendix A.1.

For males affected by a job loss, the share of separations has most drastically increased for couples in which both spouses are highly educated and employed. It increased from 0.12% to 0.25% of all divorces. This corresponds to roughly 10% of the overall increase we measure. Interestingly, this education combination never divorces upon job loss of the man when the woman is already unemployed. Conversely, a sizable number of divorces occur for similar couples with the only difference being that the husband has a lower education type (medium) than the woman (high). Male education seems to matter a lot for the survival probability of marriages. Other important contributions to divorces triggered by male job loss come from employed couples with low education males and low or medium educated females, respectively. This is true irrespectively of the female employment status. The share has not increased in the second half of our sample, however. Rather, it decreased towards the end.

We find that a male finding a job, so a transition from unemployment to employment, almost never triggers a divorce. The only exception: when both partners have medium education and the wife is already employed, the husband’s new job leads to 0.33% of divorces in the beginning of our period and to 0.74% in the end. This is roughly a third of the overall increase of labor market divorces.

Let us now turn to female labor market transitions. Overall, the likelihood of a divorce is much lower when women lose their job as compared to men. Two groups of couples have sizable propensities to divorce upon female job loss, however: employed couples with two low education spouses and employed couples with a high education husband and a medium education wife. For the latter group, the share of divorces has increased from 0.19% to 0.30%, about 10% of the overall effect. The share of the low-education employed couple is also more than 0.10% but it has not increased over time.

Finally, we look at married females who exit from unemployment. There are four striking cases which, in combination, make up most of the time dynamics of labor market divorces we observe. First, highly-educated couples in which both spouses were unemployed before the woman finds a job are responsible for 2.10% of the overall number of divorces in 2013. This share has almost quintupled over time and alone accounts for most of the observed aggregate increase of labor market divorces. Second, the divorce share of unemployed couples with a low education husband and a highly educated wife has also increased significantly over time, it grew by almost 50%. Conversely, the share of the opposite couple type, high education man and low education woman who are both unemployed and divorce when the woman starts working is also sizable (0.32% in 2013) but decreasing over time. Third, now looking at couples where the man is already employed and the woman starts working, the share of divorces between, again, two highly educated spouses has almost doubled to 0.45% in 2013. Fourth, the share of the biggest contributor to labor market divorces in the beginning of our sample (almost 2%) has decreased significantly. Couples who share medium education and the woman joins the labor market in addition to a working husband, however, are still responsible for 1.31% of all divorces in 2013.

The analysis of the shares of labor market divorces reveals an interesting general picture: it is mostly the group of couples with two high or medium education spouses that drives labor market divorces, so the *sorted couples*. Some combinations of low education couples are also affected by labor market uncertainty but their share of divorces has decreased significantly over time.

Strikingly, in many cases these sorted high education couples divorce when the woman starts working. The increasing share of this kind of divorces can be connected to an earlier observation we made based on the labor market flow data. The job-finding rate of high and medium education women has increased significantly in the second half of our sample, much more strongly than the respective transition rates of males.

Finally, recall Figure 3.1. We have observed that formerly increasing educational sorting has leveled off in the second half of the 2000s. Our structural decomposition of divorce flows has enabled us to make sense of this observation. If sorted couples show the tendency to react more strongly to both job-finding and separation shocks it would be natural to expect that the correlation between education values will not surpass a certain point.

3.6 Conclusions

This piece of research has connected the two-sided marriage market model of Goussé et al. (2017) to the labor market. The uncertainty that singles and married couples face regarding their labor market status is, as we show theoretically and empirically, an important driving force of matching decisions in the marriage market. Using three sources of German micro data, we document that the German marriage market is coined by positive sorting of in the marriage market based on education, income, and employment status. The trend towards more educational sorting, however, has stalled in recent years.

We perform a structural empirical analysis that allows us to back out key elements of our marriage market model from the data, specifically meeting rates and matching probabilities. We find that search patterns in the marriage market appear to be directed for highly educated individuals while single with low education search randomly. Based on our data and the estimated model parameters, we decompose the aggregate flow of divorces into the share induced by labor market transitions and by match-specific shocks. Transitions from employment to unemployment or vice versa make up only a minor fraction of all divorces. This fraction, however, shows interesting dynamics. The share of labor market divorces has grown by more than 20% since the mid 2000s and most of the additional divorcées are highly educated and were married to highly educated individuals. Most of these marriages break up when a previously unemployed woman starts working, especially if the husband stays unemployed. In 2013, 5.3% of all divorces happened when a previously unemployed woman started to work. This percentage share equals 27,968 divorces. The case that the literature has previously analyzed, divorce upon male job loss, accounts only for a shrinking fraction of all divorces in Germany.

One possible explanation for the differential changes of different couple types in the overall number of labor market divorce relates to the booming German labor market in the second half of the 2000s. Many low education couples divorce for reasons of economic hardship and related stress in the relationship when they become unemployed. This divorce hazard may have been mitigated by the shrinking unemployment rate in Germany and the good general macroeconomic environment. High education couples who are the source of assortative matching in the marriage market, however, seem to divorce for other reasons. When a high education women starts working, for instance, this might change the balance of power and the resource sharing in a household. Due to favorable outside options of two employed persons, the option value of searching for a new partner in the marriage market might become dominant.

Appendix to Chapter 3

A.1 Additional Results

MC Data

Figure A2.1: Education by Gender

(a) Men (b) Women

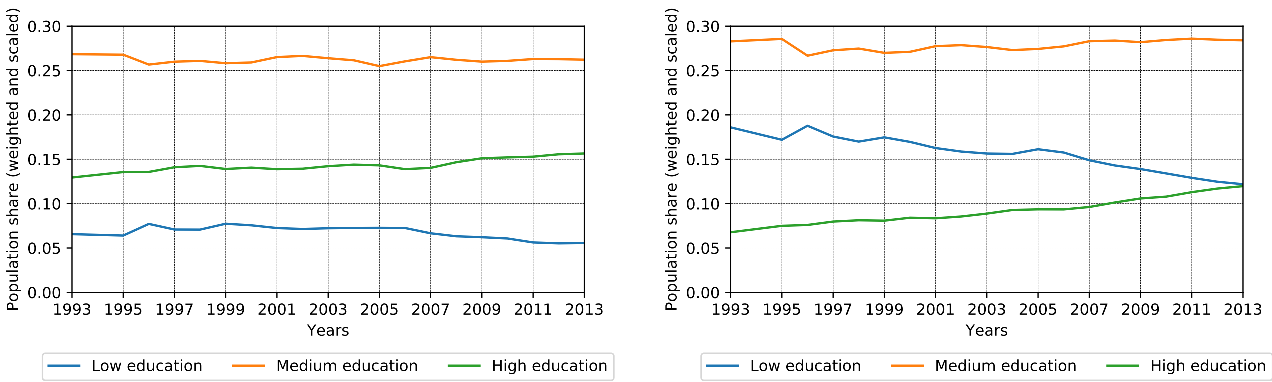

Note: Population shares are weighted and scaled using the MC sample weights. 100% on the y-axis corresponds to the full population, including married and unmarried individuals of both sexes. MC source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations.

Figure A2.2: Education and Marital status by Gender

(a) Men (b) Women

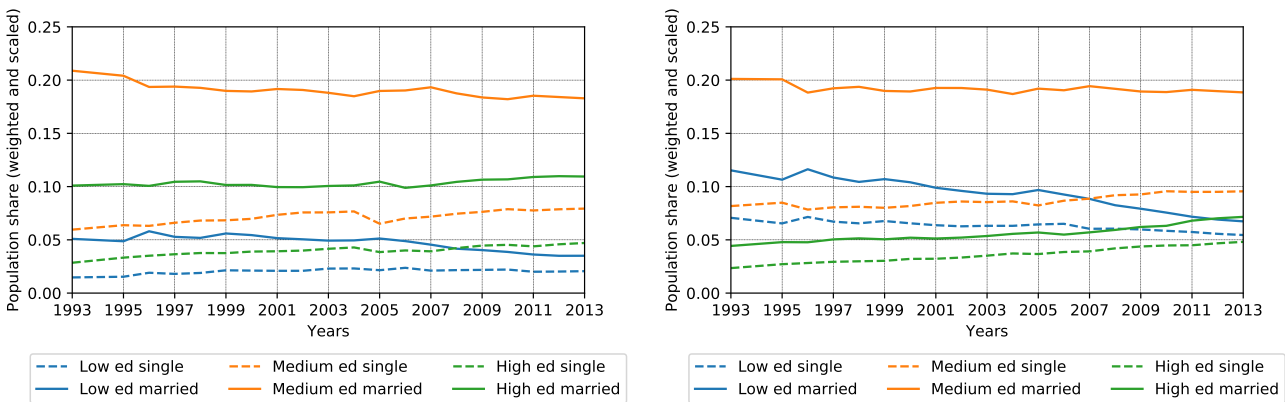

Note: Population shares are weighted and scaled using the MC sample weights. 100% on the y-axis corresponds to the full population, including married and unmarried individuals of both sexes. MC source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations.

Figure A2.1 depicts the respective population shares of men and women in each of three education categories: Panel (a) shows that for men the education distribution has not changed much between 1991 and 2013. The share of highly educated men increased somewhat and surpassed 15% in 2009. For women in Panel (b), the share of women with high education has increased much stronger, from 7.2% in 1991 to almost 12% in 2013. The share of low education women has decreased accordingly. For both men and women the share with medium education is roughly constant over time. Figure A2.2 further breaks down the three education shares into married and single individuals.

SIAB Data

Figure A2.3: Job-Finding and Separation Rates

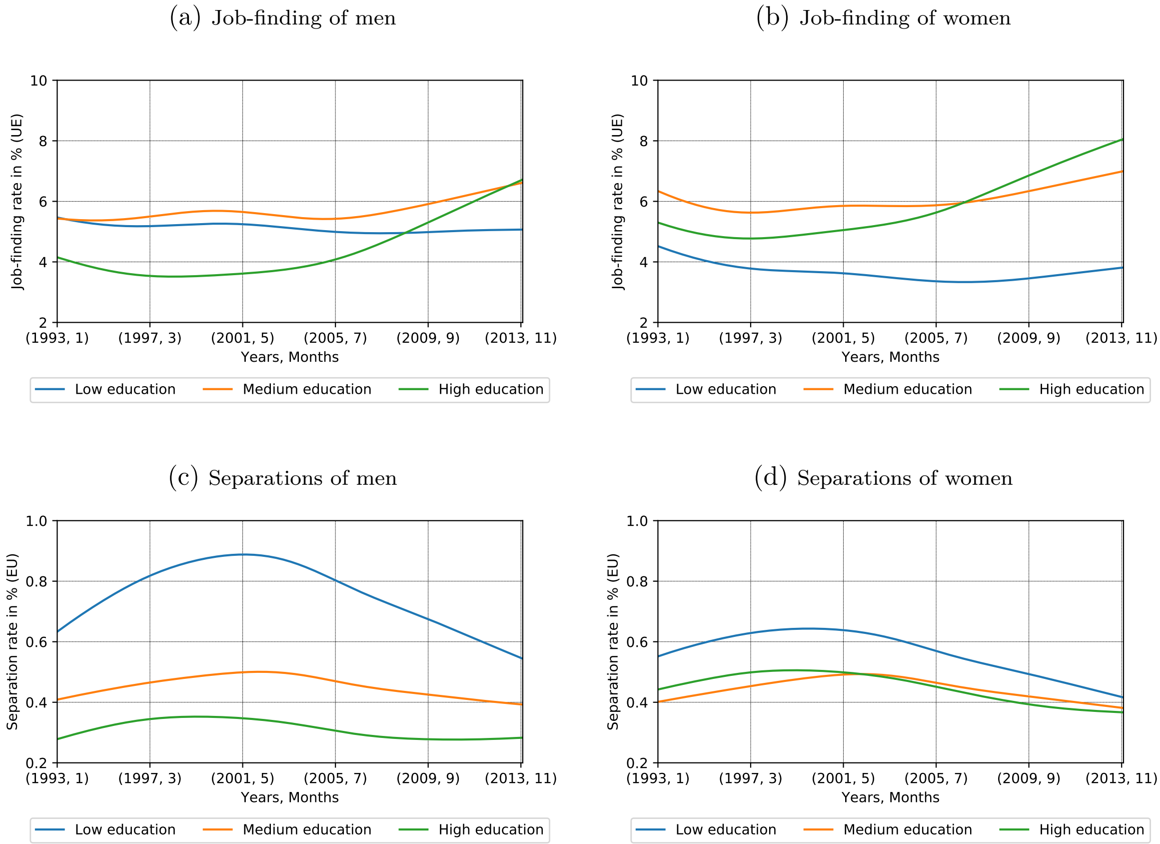

Note: seasonally adjusted (X-13-ARIMA-SEATS) and HP-filtered ($\lambda = 900,000$) monthly job-finding and separation rates by gender and education category. SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB SUF 7514, 1993-2013, own calculations.

Shares Labor Market Transitions in Divorce Flows

The following four tables show the decomposition of the aggregate number of divorces into shares for all gender-education-marriage cells, using equation (3.15).

Table A2.1: Contribution of Male Separations (EU) to the Overall Divorce Flow

	$\widetilde{m}^{ee}_{11t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ee}_{12t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ee}_{13t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ee}_{21t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ee}_{22t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ee}_{23t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ee}_{31t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{ee}_{32t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{ee}_{33t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{eu}_{11t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{eu}_{12t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{eu}_{13t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{eu}_{21t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{eu}_{22t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{eu}_{23t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{eu}_{31t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{eu}_{32t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{eu}_{33t}\widetilde{\tau}^{e}_{3t}$
1993	0.10%	0.14%	0.04%	0.09%	0.00	0.00	0.02%	0.00	0.12%	0.16%	0.18%	0.03%	0.12%	0.00	0.11%	0.01%	0.06%	0.00
1995	0.10%	0.14%	0.04%	0.09%	0.00	0.00	0.02%	0.00	0.15%	0.18%	0.19%	0.02%	0.11%	0.00	0.11%	0.01%	0.06%	0.00
1996	0.14%	0.15%	0.05%	0.09%	0.00	0.00	0.02%	0.00	0.16%	0.22%	0.19%	0.03%	0.11%	0.00	0.11%	0.01%	0.06%	0.00
1997	0.13%	0.14%	0.05%	0.08%	0.00	0.00	0.02%	0.00	0.17%	0.20%	0.18%	0.03%	0.10%	0.00	0.11%	0.01%	0.06%	0.00
1998	0.12%	0.14%	0.05%	0.08%	0.00	0.00	0.02%	0.00	0.17%	0.21%	0.18%	0.03%	0.10%	0.00	0.11%	0.01%	0.06%	0.00
1999	0.16%	0.15%	0.05%	0.09%	0.00	0.00	0.02%	0.00	0.18%	0.22%	0.17%	0.02%	0.10%	0.00	0.11%	0.01%	0.06%	0.00
2000	0.17%	0.16%	0.05%	0.09%	0.00	0.00	0.02%	0.00	0.19%	0.22%	0.17%	0.02%	0.09%	0.00	0.11%	0.01%	0.06%	0.00
2001	0.17%	0.16%	0.05%	0.10%	0.00	0.00	0.02%	0.00	0.18%	0.21%	0.16%	0.02%	0.09%	0.00	0.11%	0.01%	0.06%	0.00
2002	0.16%	0.16%	0.04%	0.09%	0.00	0.00	0.02%	0.00	0.19%	0.20%	0.16%	0.02%	0.09%	0.00	0.11%	0.01%	0.06%	0.00
2003	0.14%	0.14%	0.04%	0.09%	0.00	0.00	0.02%	0.00	0.18%	0.18%	0.15%	0.02%	0.08%	0.00	0.10%	0.01%	0.05%	0.00
2004	0.13%	0.13%	0.04%	0.08%	0.00	0.00	0.01%	0.00	0.18%	0.17%	0.14%	0.02%	0.08%	0.00	0.10%	0.01%	0.05%	0.00
2005	0.13%	0.14%	0.04%	0.09%	0.00	0.00	0.02%	0.00	0.18%	0.18%	0.14%	0.02%	0.08%	0.00	0.10%	0.01%	0.05%	0.00
2006	0.13%	0.14%	0.04%	0.09%	0.00	0.00	0.01%	0.00	0.18%	0.17%	0.14%	0.02%	0.08%	0.00	0.10%	0.01%	0.05%	0.00
2007	0.11%	0.15%	0.03%	0.09%	0.00	0.00	0.01%	0.00	0.18%	0.15%	0.13%	0.02%	0.08%	0.00	0.10%	0.01%	0.05%	0.00
2008	0.10%	0.14%	0.03%	0.09%	0.00	0.00	0.01%	0.00	0.20%	0.14%	0.13%	0.01%	0.08%	0.00	0.09%	0.01%	0.05%	0.00
2009	0.10%	0.13%	0.03%	0.08%	0.00	0.00	0.01%	0.00	0.20%	0.12%	0.11%	0.01%	0.07%	0.00	0.09%	0.01%	0.04%	0.00
2010	0.09%	0.12%	0.03%	0.08%	0.00	0.00	0.01%	0.00	0.21%	0.11%	0.10%	0.01%	0.07%	0.00	0.08%	0.00%	0.04%	0.00
2011	0.08%	0.12%	0.03%	0.08%	0.00	0.00	0.01%	0.00	0.23%	0.10%	0.09%	0.01%	0.06%	0.00	0.08%	0.00%	0.04%	0.00
2012	0.07%	0.11%	0.03%	0.08%	0.00	0.00	0.01%	0.00	0.24%	0.09%	0.08%	0.01%	0.06%	0.00	0.08%	0.00%	0.04%	0.00
2013	0.07%	0.11%	0.03%	0.08%	0.00	0.00	0.01%	0.00	0.25%	0.09%	0.08%	0.01%	0.06%	0.00	0.07%	0.00%	0.04%	0.00

Note: SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB SUF 7514, 1993-2013, own calculations. MC source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations.

Table A2.2: Contribution of Male Job-Findings (UE) to the Overall Divorce Flow

	$\widetilde{m}^{ue}_{11t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{ue}_{12t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{ue}_{13t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{ue}_{21t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{ue}_{22t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{ue}_{23t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{ue}_{31t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{ue}_{32t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{ue}_{33t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{uu}_{11t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{uu}_{12t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{uu}_{13t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{uu}_{21t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{uu}_{22t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{uu}_{23t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{uu}_{31t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{uu}_{32t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{uu}_{33t}\widetilde{\tau}^{u}_{3t}$
1993	0.00	0.00	0.00	0.00	0.33%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1995	0.00	0.00	0.00	0.00	0.37%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1996	0.00	0.00	0.00	0.00	0.40%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1997	0.00	0.00	0.00	0.00	0.43%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1998	0.00	0.00	0.00	0.00	0.46%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1999	0.00	0.00	0.00	0.00	0.48%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2000	0.00	0.00	0.00	0.00	0.49%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2001	0.00	0.00	0.00	0.00	0.53%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2002	0.00	0.00	0.00	0.00	0.55%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2003	0.00	0.00	0.00	0.00	0.56%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2004	0.00	0.00	0.00	0.00	0.55%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2005	0.00	0.00	0.00	0.00	0.58%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2006	0.00	0.00	0.00	0.00	0.60%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2007	0.00	0.00	0.00	0.00	0.61%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2008	0.00	0.00	0.00	0.00	0.61%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2009	0.00	0.00	0.00	0.00	0.66%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2010	0.00	0.00	0.00	0.00	0.68%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2011	0.00	0.00	0.00	0.00	0.67%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2012	0.00	0.00	0.00	0.00	0.68%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2013	0.00	0.00	0.00	0.00	0.74%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Note: SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB SUF 7514, 1993-2013, own calculations. MC source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations.

Table A2.3: Contribution of Female Separations (EU) to the Overall Divorce Flow

	$\widetilde{m}^{ee}_{11t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ee}_{21t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ee}_{31t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ee}_{12t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ee}_{22t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ee}_{32t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ee}_{13t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{ee}_{23t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{ee}_{33t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{ue}_{11t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ue}_{21t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ue}_{31t}\widetilde{\tau}^{e}_{1t}$	$\widetilde{m}^{ue}_{12t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ue}_{22t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ue}_{32t}\widetilde{\tau}^{e}_{2t}$	$\widetilde{m}^{ue}_{13t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{ue}_{23t}\widetilde{\tau}^{e}_{3t}$	$\widetilde{m}^{ue}_{33t}\widetilde{\tau}^{e}_{3t}$
1993	0.16%	0.02%	0.00	0.00	0.00	0.19%	0.00	0.00	0.00	0.04%	0.02%	0.00	0.02%	0.00	0.02%	0.00	0.02%	0.00
1995	0.16%	0.02%	0.00	0.00	0.00	0.22%	0.00	0.00	0.00	0.04%	0.03%	0.00	0.02%	0.00	0.02%	0.00	0.02%	0.00
1996	0.21%	0.02%	0.00	0.00	0.00	0.22%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.02%	0.00	0.03%	0.00
1997	0.18%	0.02%	0.00	0.00	0.00	0.23%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.03%	0.00	0.03%	0.00
1998	0.17%	0.02%	0.00	0.00	0.00	0.24%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.03%	0.00	0.03%	0.00
1999	0.22%	0.02%	0.00	0.00	0.00	0.24%	0.00	0.00	0.00	0.07%	0.03%	0.00	0.02%	0.00	0.03%	0.00	0.03%	0.00
2000	0.23%	0.02%	0.00	0.00	0.00	0.26%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.03%	0.00	0.03%	0.00
2001	0.23%	0.02%	0.00	0.00	0.00	0.26%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.03%	0.00	0.03%	0.00
2002	0.21%	0.02%	0.00	0.00	0.00	0.26%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.03%	0.00	0.03%	0.00
2003	0.19%	0.01%	0.00	0.00	0.00	0.26%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.03%	0.00	0.03%	0.00
2004	0.17%	0.01%	0.00	0.00	0.00	0.25%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2005	0.17%	0.01%	0.00	0.00	0.00	0.27%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2006	0.17%	0.01%	0.00	0.00	0.00	0.26%	0.00	0.00	0.00	0.06%	0.03%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2007	0.15%	0.02%	0.00	0.00	0.00	0.28%	0.00	0.00	0.00	0.05%	0.03%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2008	0.14%	0.01%	0.00	0.00	0.00	0.30%	0.00	0.00	0.00	0.04%	0.03%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2009	0.14%	0.01%	0.00	0.00	0.00	0.30%	0.00	0.00	0.00	0.04%	0.03%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2010	0.12%	0.01%	0.00	0.00	0.00	0.30%	0.00	0.00	0.00	0.04%	0.02%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2011	0.10%	0.01%	0.00	0.00	0.00	0.30%	0.00	0.00	0.00	0.03%	0.02%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2012	0.10%	0.01%	0.00	0.00	0.00	0.31%	0.00	0.00	0.00	0.03%	0.02%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00
2013	0.10%	0.01%	0.00	0.00	0.00	0.30%	0.00	0.00	0.00	0.03%	0.02%	0.00	0.02%	0.00	0.04%	0.00	0.03%	0.00

Note: SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB SUF 7514, 1993-2013, own calculations. MC source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations.

Table A2.4: Contribution of Female Job-Findings (UE) to the Overall Divorce Flow

	$\widetilde{m}^{eu}_{11t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{eu}_{21t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{eu}_{31t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{eu}_{12t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{eu}_{22t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{eu}_{32t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{eu}_{13t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{eu}_{23t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{eu}_{33t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{uu}_{11t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{uu}_{21t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{uu}_{31t}\widetilde{\tau}^{u}_{1t}$	$\widetilde{m}^{uu}_{12t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{uu}_{22t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{uu}_{32t}\widetilde{\tau}^{u}_{2t}$	$\widetilde{m}^{uu}_{13t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{uu}_{23t}\widetilde{\tau}^{u}_{3t}$	$\widetilde{m}^{uu}_{33t}\widetilde{\tau}^{u}_{3t}$
1993	0.00	0.00	0.14%	0.40%	1.98%	0.00	0.17%	0.00	0.25%	0.00	0.00	0.39%	0.00	0.02%	0.00	0.46%	0.00	0.38%
1995	0.00	0.00	0.11%	0.34%	1.77%	0.00	0.13%	0.00	0.24%	0.00	0.00	0.33%	0.00	0.02%	0.00	0.40%	0.00	0.45%
1996	0.00	0.00	0.11%	0.31%	1.57%	0.00	0.14%	0.00	0.22%	0.00	0.00	0.36%	0.00	0.02%	0.00	0.44%	0.00	0.47%
1997	0.00	0.00	0.10%	0.29%	1.54%	0.00	0.13%	0.00	0.24%	0.00	0.00	0.34%	0.00	0.02%	0.00	0.43%	0.00	0.53%
1998	0.00	0.00	0.09%	0.27%	1.52%	0.00	0.12%	0.00	0.25%	0.00	0.00	0.33%	0.00	0.02%	0.00	0.42%	0.00	0.58%
1999	0.00	0.00	0.08%	0.26%	1.44%	0.00	0.11%	0.00	0.24%	0.00	0.00	0.31%	0.00	0.02%	0.00	0.41%	0.00	0.56%
2000	0.00	0.00	0.08%	0.26%	1.45%	0.00	0.11%	0.00	0.25%	0.00	0.00	0.31%	0.00	0.02%	0.00	0.42%	0.00	0.62%
2001	0.00	0.00	0.07%	0.25%	1.44%	0.00	0.10%	0.00	0.26%	0.00	0.00	0.30%	0.00	0.03%	0.00	0.42%	0.00	0.66%
2002	0.00	0.00	0.07%	0.24%	1.43%	0.00	0.10%	0.00	0.26%	0.00	0.00	0.28%	0.00	0.03%	0.00	0.41%	0.00	0.73%
2003	0.00	0.00	0.06%	0.23%	1.34%	0.00	0.10%	0.00	0.27%	0.00	0.00	0.28%	0.00	0.03%	0.00	0.43%	0.00	0.79%
2004	0.00	0.00	0.06%	0.23%	1.31%	0.00	0.10%	0.00	0.28%	0.00	0.00	0.29%	0.00	0.03%	0.00	0.46%	0.00	0.87%
2005	0.00	0.00	0.06%	0.24%	1.38%	0.00	0.11%	0.00	0.29%	0.00	0.00	0.32%	0.00	0.03%	0.00	0.54%	0.00	1.04%
2006	0.00	0.00	0.06%	0.25%	1.43%	0.00	0.11%	0.00	0.29%	0.00	0.00	0.32%	0.00	0.03%	0.00	0.56%	0.00	1.09%
2007	0.00	0.00	0.06%	0.25%	1.47%	0.00	0.10%	0.00	0.31%	0.00	0.00	0.32%	0.00	0.03%	0.00	0.58%	0.00	1.24%
2008	0.00	0.00	0.05%	0.25%	1.42%	0.00	0.10%	0.00	0.34%	0.00	0.00	0.32%	0.00	0.03%	0.00	0.61%	0.00	1.42%
2009	0.00	0.00	0.06%	0.24%	1.35%	0.00	0.11%	0.00	0.36%	0.00	0.00	0.32%	0.00	0.03%	0.00	0.64%	0.00	1.55%
2010	0.00	0.00	0.05%	0.24%	1.35%	0.00	0.10%	0.00	0.38%	0.00	0.00	0.32%	0.00	0.03%	0.00	0.65%	0.00	1.69%
2011	0.00	0.00	0.05%	0.23%	1.32%	0.00	0.11%	0.00	0.40%	0.00	0.00	0.31%	0.00	0.03%	0.00	0.63%	0.00	1.75%
2012	0.00	0.00	0.05%	0.23%	1.32%	0.00	0.11%	0.00	0.42%	0.00	0.00	0.31%	0.00	0.04%	0.00	0.65%	0.00	1.93%
2013	0.00	0.00	0.05%	0.24%	1.31%	0.00	0.11%	0.00	0.45%	0.00	0.00	0.32%	0.00	0.04%	0.00	0.67%	0.00	2.10%

Note: SIAB source: Research Data Centre (FDZ) of the Federal Employment Agency at the Institute for Employment Research, SIAB SUF 7514, 1993-2013, own calculations. MC source: Research Data Center of the Statistical Offices of the Länder and the Federal State, Microcensus, 1993-2013, own calculations.

A.2 Theoretical Appendix

Surplus function

To obtain the surplus function, consider first the gain from marriage for a female of type j with labor market status l. i.e.,

$$\begin{aligned}
& \left(r+\delta_{ij}+\tau_j(l)+\tau_i(-l)\right)\left[V_j^{l,-l}-V_j^l\right] \qquad (3.16)\\
= \; & u_j^{l,-l}+\delta_{ij}\int \max\left[V_j^{l,-l}(z')-V_j^l,0\right]dG(z') \\
& +\tau_j(l)\max\left[V_j^{l',-l}-V_j^{l'},0\right]+\tau_i(-l)\max\left[V_j^{l,-l'}-V_j^l,0\right] \\
& -rV_j^l+\tau_j(l)\left[V_j^{l'}-V_j^l\right] \\
= \; & u_j^{l,-l}+\delta_{ij}\int \max\left[V_j^{l,-l}(z')-V_j^l,0\right]dG(z') \\
& +\tau_j(l)\max\left[V_j^{l',-l}-V_j^{l'},0\right]+\tau_i(-l)\max\left[V_j^{l,-l'}-V_j^l,0\right] \\
& -u_j^l-\lambda_{ij}\iiint\left[V_j^{l,-l'}(z')-V_j^l\right]W_{ij}^{ll}(z')\,dG(z')\,s(i,l')\,didl',
\end{aligned}$$

where the second equality follows from substituting rV_j^l using equation (3.7). The gain from marriage for the male partner is defined similarly. Take the per period utility function $u_i^l(t_i,d_i,d_j|z)$ for couples as defined in equation (3.4)

$$\begin{aligned}
u_i^{l,-l}(t_i,d_i,d_j|z) &= zF_{ij}^1(d_i,d_j)\left[R_i^l+\frac{\iota}{2}\left(R_j^{-l}-R_i^l\right)^2-t_i+T_i^l-d_i\right], \\
u_j^{l,-l}(t_j,d_j,d_i|z) &= zF_{ij}^1(d_i,d_j)\left[R_j^l+\frac{\iota}{2}\left(R_j^l-R_i^{-l}\right)^2-C_{ij}+t_i+T_j^l-d_j\right],
\end{aligned}$$

where we used $t_i+t_j=C_{ij}$ to substitute t_j. Maximizing the Nash-Product (3.6) with respect to t_i impies,

$$\begin{aligned}
\frac{1-\beta}{V_j^{l,-l}-V_j^l}\frac{\partial u_j^l(t_j,d_j,d_i|z)}{\partial t_i} &= \frac{\beta}{V_i^{l,-l}-V_i^l}\frac{\partial u_i^l(t_i,d_i,d_j|z)}{\partial t_i}, \\
(1-\beta)\left[V_i^{l,-l}-V_i^l\right] &= \beta\left[V_j^{l,-l}-V_j^l\right].
\end{aligned}$$

Using the definition of the surplus in equation (3.8) allows us to write $\left[V_i^{l,-l}-V_i^l\right]=\beta S_{ij}^{ll}$ and $\left[V_j^{l,-l}-V_j^l\right]=(1-\beta)S_{ij}^{ll}$. Using the gain from marriage for a female of type j in equation (3.16) and the respective equation for a male of type i allows us to write the surplus for any employment status ll as stated in equation (3.9).

The optimal time input into the public good's production follows from differentiating the Nash-Product (3.6), i.e.,

$$\begin{aligned}\frac{1-\beta}{V_f^{ll}-V_j^l}\frac{\partial u_j^l\left(t_j,d_j,d_i|z\right)}{\partial d_i} &= -\frac{\beta}{V_m^{ll}-V_i^l}\frac{\partial u_j^l\left(t_j,d_j,d_i|z\right)}{\partial d_i}\\ z\frac{\partial F_{ij}^1\left(d_i,d_j\right)}{\partial d_i}\left[R_i^l+R_j^l+\iota\left(R_j^l-R_i^l\right)^2-C_{ij}+T_i^l-d_i+T_j^l-d_j\right] &= zF_{ij}^1\left(d_i,d_j\right)\end{aligned}$$

where the second line follows from taking into account that $(1-\beta)\left[V_m^{ll}-V_i^l\right]=\beta\left[V_f^{ll}-V_j^l\right]$. Under the assumption that the public good production functions are Stone-Geary, i.e.,

$$F_{ij}^1\left(d_i,d_j\right)=Z_{ij}\left(d_i-D_i^1\right)^{K_i}\left(d_j-D_j^1\right)^{K_j},$$

with $0<K_i+K_j<1$, we get,

$$\frac{\partial F_{ij}^1\left(d_i,d_j\right)}{\partial d_i}=\frac{K_i}{d_i-D_i^1}F_{ij}^1\left(d_i,d_j\right).$$

The optimal time inputs for the female and the male are given by,

$$\begin{aligned}\left(d_i-D_i\right) &= \frac{K_i}{1+K_j+K_i}\left[R_i^l+R_j^l+\iota\left(R_j^l-R_i^l\right)^2-C_{ij}+T_i^l-D_i^1+T_j^l-D_j^1\right],\\ \left(d_j-D_j\right) &= \frac{K_j}{1+K_j+K_i}\left[R_i^l+R_j^l+\iota\left(R_j^l-R_i^l\right)^2-C_{ij}+T_i^l-D_i^1+T_j^l-D_j^1\right].\end{aligned}$$

In equilibrium the public good is therefore given by,

$$\begin{aligned}F_{ij}^1\left(d_i,d_j\right) &= Z_{ij}\left(1+K_j+K_i\right)K\left[R_i^l+R_j^l+\iota\left(R_j^l-R_i^l\right)^2-C_{ij}+T_i^l-D_i^1+T_j^l-D_j^1\right]^{K_j+K_i}\\ \text{where } K &= \frac{\left(K_i\right)^{K_i}\left(K_j\right)^{K_j}}{\left(1+K_j+K_i\right)^{1+K_j+K_i}}.\end{aligned}$$

Maximized joint flow utility in a marriage is hence given by,

$$u_i^{l,-l}+u_j^{l,-l}=zZ_{ij}K\left[R_i^l+R_j^l+\iota\left(R_j^l-R_i^l\right)^2-C_{ij}+T_i^l-D_i^1+T_j^l-D_j^1\right]^{1+K_j+K_i}.$$

Given the flow utility function for singles in equation (3.3), and the public good production function $F_i^0\left(d_i\right)=\left(d_i-D_i^0\right)^{K_i^0}$ the optimal time input for the public good is hence given by $d_i - D_i^0 = K_i^0$ and the maximized flow utility for a single male of type i by

$$u_i^l = \left(K_i^0\right)^{K_i^0} \left[R_i^l + T_i^l - D_i^0 - K_i^0\right].$$

The maximized joint flow utility in a marriage $u_i^{l,-l} + u_j^{l,-l}$ and the maximized flow utility for the respective singles u_i^l and u_j^l give the maximized joint flow surplus of both partners in equation (3.10).

Computation of the fixed point

The first step to determine the surplus functions $S_{ij}^{ll}(z)$ and the cutoff values z_{ij}^{ll} is to compute integrated surpluses $\overline{S}_{z_{ij}^{ll}}^{l'l}$, where the subindex z_{ij}^{ll} indicates the support over which the surplus is integrated, i.e.,

$$\overline{S}_{z_{ij}^{ll}}^{l'l} \equiv \int_{z_{ij}^{ll}}^{\infty} S_{ij}^{l'l}(z'')\, dG(z'').$$

Integrating the surplus functions (3.9) for the different labor market status combinations $\{ll, l'l, ll', l'l'\}$ over the support under consideration, i.e., support $\left[z_{ij}^{ll}, \infty\right]$ if the considered cutoff value is z_{ij}^{ll}, gives the following fixed-point equation for each labor market status combinations $\{ll, l'l, ll', l'l'\}$ given z_{ij}^{ll}, i.e.,

$$\begin{aligned}
\left(r + \delta_{ij} + \tau_i(l) + \tau_j(l)\right) \overline{S}_{z_{ij}^{ll}}^{ll} \quad = \quad & \kappa_{ij}\left[W_{ij} + \psi_{ij}^{ll}\right]^{\kappa} \Phi\left(\frac{\sigma^2 - \ln z_{ij}^{ll}}{\sigma}\right) e^{\frac{1}{2}\sigma^2} \\
& + \left(\delta_{ij}\overline{S}_{z_{ij}^{ll}}^{ll} - \widehat{\Theta}_{ij}^{ll}\right)\left[1 - \Phi\left(\frac{\ln z_{ij}^{ll}}{\sigma}\right)\right] \\
& + \tau_i(l)\, \eta_{ij}^{(ll,l'l)} \overline{S}_{z_{ij}^{ll}}^{l'l} + \tau_i(l)\left(1 - \eta_{ij}^{(ll,l'l)}\right) \overline{S}_{z_{ij}^{l'l}}^{l'l} \\
& + \tau_j(l)\, \eta_{ij}^{(ll,ll')} \overline{S}_{z_{ij}^{ll}}^{ll'} + \tau_j(l)\left(1 - \eta_{ij}^{(ll,ll')}\right) \overline{S}_{z_{ij}^{ll'}}^{ll'},
\end{aligned}$$

$$\begin{aligned}\left(r+\delta_{ij}+\tau_i\left(l'\right)+\tau_j\left(l\right)\right)\overline{S}^{l'l}_{z^{ll}_{ij}} &= \kappa_{ij}\left[W_{ij}+\psi^{l'l}_{ij}\right]^{\kappa}\Phi\left(\frac{\sigma^2-\ln z^{ll}_{ij}}{\sigma}\right)e^{\frac{1}{2}\sigma^2}\\ &+\left(\delta_{ij}\overline{S}^{l'l}_{z^{l'l}_{ij}}-\widehat{\Theta}^{l'l}_{ij}\right)\left[1-\Phi\left(\frac{\ln z^{ll}_{ij}}{\sigma}\right)\right]\\ &+\tau_i\left(l'\right)\overline{S}^{ll}_{z^{ll}_{ij}}\\ &+\tau_j\left(l\right)\eta^{(ll,l'l')}_{ij}\overline{S}^{l'l'}_{z^{ll}_{ij}}+\tau_j\left(l\right)\left(1-\eta^{(ll,l'l')}_{ij}\right)\overline{S}^{l'l'}_{z^{l'l'}_{ij}},\end{aligned}$$

$$\begin{aligned}\left(r+\delta_{ij}+\tau_i\left(l\right)+\tau_j\left(l'\right)\right)\overline{S}^{ll'}_{z^{ll}_{ij}} &= \kappa_{ij}\left[W_{ij}+\psi^{ll'}_{ij}\right]^{\kappa}\Phi\left(\frac{\sigma^2-\ln z^{ll}_{ij}}{\sigma}\right)e^{\frac{1}{2}\sigma^2}\\ &+\left(\delta_{ij}\overline{S}^{ll'}_{z^{ll'}_{ij}}-\widehat{\Theta}^{ll'}_{ij}\right)\left[1-\Phi\left(\frac{\ln z^{ll}_{ij}}{\sigma}\right)\right]\\ &+\tau_i\left(l\right)\eta^{(ll,l'l')}_{ij}\overline{S}^{l'l'}_{z^{ll}_{ij}}+\tau_i\left(l\right)\left(1-\eta^{(ll,l'l')}_{ij}\right)\overline{S}^{l'l'}_{z^{l'l'}_{ij}}\\ &+\tau_j\left(l'\right)\overline{S}^{ll}_{z^{ll}_{ij}},\end{aligned}$$

$$\begin{aligned}\left(r+\delta_{ij}+\tau_i\left(l'\right)+\tau_j\left(l'\right)\right)\overline{S}^{l'l'}_{z^{ll}_{ij}} &= \kappa_{ij}\left[W_{ij}+\psi^{l'l'}_{ij}\right]^{\kappa}\Phi\left(\frac{\sigma^2-\ln z^{ll}_{ij}}{\sigma}\right)e^{\frac{1}{2}\sigma^2}\\ &+\left(\delta_{ij}\overline{S}^{l'l'}_{z^{l'l'}_{ij}}-\widehat{\Theta}^{l'l'}_{ij}\right)\left[1-\Phi\left(\frac{\ln z^{ll}_{ij}}{\sigma}\right)\right]\\ &+\tau_i\left(l'\right)\eta^{(ll,ll')}_{ij}\overline{S}^{ll'}_{z^{ll}_{ij}}+\tau_i\left(l'\right)\left(1-\eta^{(ll,ll')}_{ij}\right)\overline{S}^{ll'}_{z^{ll'}_{ij}}\\ &+\tau_j\left(l'\right)\eta^{(ll,l'l)}_{ij}\overline{S}^{l'l}_{z^{ll}_{ij}}+\tau_j\left(l'\right)\left(1-\eta^{(ll,l'l)}_{ij}\right)\overline{S}^{l'l}_{z^{l'l}_{ij}},\end{aligned}$$

where

$$\eta^{(ll,l'l)}_{ij}=\begin{cases}0 & \text{if } z^{ll}_{ij}\leq z^{l'l}_{ij},\\ 1 & \text{if } z^{ll}_{ij}>z^{l'l}_{ij},\end{cases}$$

and

$$\begin{aligned}
\widehat{\Theta}_{ij}^{ll} &= \kappa_i \left[w_i + \psi_i^l\right] + \kappa_j \left[w_j + \psi_j^l\right] \\
&+ \lambda_{ij} (1-\beta) \iint \overline{S}_{z_{ij}^{ll''}}^{ll''} s(j, l'')\, dj dl'' + \lambda_{ij} \beta \iint \overline{S}_{z_{ij}^{l''l}}^{l''l} s(i, l'')\, di dl'', \\
\widehat{\Theta}_{ij}^{l'l} &= \kappa_i \left[w_i + \psi_i^{l'}\right] + \kappa_j \left[w_j + \psi_j^l\right] \\
&+ \lambda_{ij} (1-\beta) \iint \overline{S}_{z_{ij}^{l'l''}}^{l'l''} s(j, l'')\, dj dl'' + \lambda_{ij} \beta \iint \overline{S}_{z_{ij}^{l''l}}^{l''l} s(i, l'')\, di dl'', \\
\widehat{\Theta}_{ij}^{ll'} &= \kappa_i \left[w_i + \psi_i^l\right] + \kappa_j \left[w_j + \psi_j^{l'}\right] \\
&+ \lambda_{ij} (1-\beta) \iint \overline{S}_{z_{ij}^{ll''}}^{ll''} s(j, l'')\, dj dl'' + \lambda_{ij} \beta \iint \overline{S}_{z_{ij}^{l''l'}}^{l''l'} s(i, l'')\, di dl'', \\
\widehat{\Theta}_{ij}^{l'l'} &= \kappa_i \left[w_i + \psi_i^{l'}\right] + \kappa_j \left[w_j + \psi_j^{l'}\right] \\
&+ \lambda_{ij} (1-\beta) \iint \overline{S}_{z_{ij}^{l'l''}}^{l'l''} s(j, l'')\, dj dl'' + \lambda_{ij} \beta \iint \overline{S}_{z_{ij}^{l''l'}}^{l''l'} s(i, l'')\, di dl''.
\end{aligned}$$

These equations have to be solved simultaneously given a set of (initial) cutoff values $\left\{z_{ij}^{ee}, z_{ij}^{ue}, z_{ij}^{eu}, z_{ij}^{uu}\right\}$. The values $\overline{S}_{z_{ij}^{ll}}^{l'l}, \overline{S}_{z_{ij}^{ll}}^{ll'}$, and $\overline{S}_{z_{ij}^{ll}}^{l'l'}$ for each $z_{ij}^{ll} \in \left\{z_{ij}^{ee}, z_{ij}^{ue}, z_{ij}^{eu}, z_{ij}^{uu}\right\}$ are not needed for further analysis. They are only required to find the fixed-points $\overline{S}_{z_{ij}^{ll}}^{ll}$ for each labor market status $ll \in \{ee, eu, ue, uu\}$. Given the fixed-points $\overline{S}_{z_{ij}^{ll}}^{ll}$ for each labor market status ll, we can use the following equation system based on the surplus function given in equation (3.9) to find the z_{ij}^{ll} associated with each labor market status ll, i.e.,

$$\begin{aligned}
0 &= z_{ij}^{ll} \kappa_{ij} \left[W_{ij} + \psi_{ij}^{ll}\right]^{\kappa} + \delta_{ij} \overline{S}_{z_{ij}^{ll}}^{ll} - \widehat{\Theta}_{ij}^{ll} \\
&\quad + \tau_i(l) \max\left[S_{ij}^{l'l}\left(z_{ij}^{ll}\right), 0\right] \\
&\quad + \tau_j(l) \max\left[S_{ij}^{ll'}\left(z_{ij}^{ll}\right), 0\right], \\
\left(r + \delta_{ij} + \tau_i(l') + \tau_j(l)\right) S_{ij}^{l'l}\left(z_{ij}^{ll}\right) &= z_{ij}^{ll} \kappa_{ij} \left[W_{ij} + \psi_{ij}^{l'l}\right]^{\kappa} + \delta_{ij} \overline{S}_{z_{ij}^{l'l}}^{l'l} - \widehat{\Theta}_{ij}^{l'l} \\
&\quad + \tau_j(l) \max\left[S_{ij}^{l'l'}\left(z_{ij}^{ll}\right), 0\right], \\
\left(r + \delta_{ij} + \tau_i(l) + \tau_j(l')\right) S_{ij}^{ll'}\left(z_{ij}^{ll}\right) &= z_{ij}^{ll} \kappa_{ij} \left[W_{ij} + \psi_{ij}^{ll'}\right]^{\kappa} + \delta_{ij} \overline{S}_{z_{ij}^{ll'}}^{ll'} - \widehat{\Theta}_{ij}^{ll'} \\
&\quad + \tau_i(l) \max\left[S_{ij}^{l'l'}\left(z_{ij}^{ll}\right), 0\right], \\
\left(r + \delta_{ij} + \tau_i(l') + \tau_j(l')\right) S_{ij}^{l'l'}\left(z_{ij}^{ll}\right) &= z_{ij}^{ll} \kappa_{ij} \left[W_{ij} + \psi_{ij}^{l'l'}\right]^{\kappa} + \delta_{ij} \overline{S}_{z_{ij}^{l'l'}}^{l'l'} - \widehat{\Theta}_{ij}^{l'l'} \\
&\quad + \tau_i(l') \max\left[S_{ij}^{ll'}\left(z_{ij}^{ll}\right), 0\right] \\
&\quad + \tau_j(l') \max\left[S_{ij}^{l'l}\left(z_{ij}^{ll}\right), 0\right],
\end{aligned}$$

where the zero in the first equation follows from $S_{ij}^{ll}\left(z_{ij}^{ll}\right)=0$. Again the values $S_{ij}^{l'l}\left(z_{ij}^{ll}\right)$, $S_{ij}^{ll'}\left(z_{ij}^{ll}\right)$, and $S_{ij}^{l'l'}\left(z_{ij}^{ll}\right)$ for each $z_{ij}^{ll} \in \left\{z_{ij}^{ee}, z_{ij}^{ue}, z_{ij}^{eu}, z_{ij}^{uu}\right\}$ are not needed for further analysis. Iterating between the two equation systems while updating the (joint) distributions of married individuals as well as singles in every iteration using Equations (3.12) and (3.14) determines the fixed-point of the system for $\overline{S}_{z_{ij}^{ll}}^{ll}$ and z_{ij}^{ll} and each combination of labor market statuses $ll \in \{ee, eu, ue, uu\}$. In practice, a Python implementation of the model converges pretty fast, in less than one minute on a Chebyshev grid with 50×50 nodes.

Calibration

Table A2.5: Calibration of the Marriage Market Model

Parameter	Symbol	Value	Source
Discount rate	r	0.05	-
Women's bargaining power	β	0.7	-
Value of nonmarket activity	b	0.8	-
Joint taxation	ι	0.1	-
Meeting rate	λ	8.5	-
Bliss shock updates	δ	0.1	-
Mean of z distribution	$\mu(z)$	1.65.	-
Standard deviation of z distribution	$\sigma(z)$	4.67.	-
Male job-finding rate	$\tau_i(u)$	5.16	SIAB data
Female job-finding rate	$\tau_j(u)$	3.88	SIAB data
Male separation rate	$\tau_i(e)$	0.71	SIAB data
Female separation rate	$\tau_j(u)$	0.60	SIAB data
Home production, single	K_f^0	0.02	Goussé et al. (2017)
Home production, single	K_m^0	0.00	Goussé et al. (2017)
Home production, married	K_f^1	0.02	Mean of Goussé et al. (2017)
Home production, married	K_m^1	0.01	Mean of Goussé et al. (2017)
Home production, married	D_f^1	0.06	Mean of Goussé et al. (2017)
Home production, married	D_m^1	0.06	Mean of Goussé et al. (2017)

A.3 Structural Estimation

The probabilities $\alpha_{ij}^{ll'}$ that a meeting of a type i male and type j female with labor market status l and l' leads to a marriage is estimated with a constrained linear equation system using the flow equations of $m(i,j,l,l')$ in equation (3.12). To simplify the notation below we define the data of period t as follows, $\widetilde{\tau}_{it}^{l} = \tau_i(l)$, $\widetilde{m}_{ijt}^{ll} = m(i,j,l,l)$ and $\widetilde{s}_{it}^{l} = s(i,l)$. The system of four equations, one for each labor market status $ll' \in \{ee, eu, ue, uu\}$ can be written in the following matrix notation, $\mathbf{y}_{ijt} = \mathbf{Z}_{ijt}\mathbf{b}_{ij} + \epsilon_{ijt}$. Given the data on labor market transition rates for males and females, $\widetilde{\tau}_{it}^{l}$ and $\widetilde{\tau}_{jt}^{l}$, and the number of singles and married couples, $\widetilde{s}_{it}^{l}$,$\widetilde{s}_{jt}^{l}$ and $\widetilde{m}_{ijt}^{ll}$, the LHS of the equation system is given by the vector,

$$\mathbf{y}_{ijt} = \begin{pmatrix} y_{ijt}^{1} \\ y_{ijt}^{2} \\ y_{ijt}^{3} \\ y_{ijt}^{4} \end{pmatrix} = \begin{pmatrix} \left(\widetilde{\tau}_{it}^{e} + \widetilde{\tau}_{jt}^{e}\right)\widetilde{m}_{ijt}^{ee} \\ \left(\widetilde{\tau}_{it}^{u} + \widetilde{\tau}_{jt}^{e}\right)\widetilde{m}_{ijt}^{ue} \\ \left(\widetilde{\tau}_{it}^{e} + \widetilde{\tau}_{jt}^{u}\right)\widetilde{m}_{ijt}^{eu} \\ \left(\widetilde{\tau}_{it}^{u} + \widetilde{\tau}_{jt}^{u}\right)\widetilde{m}_{ijt}^{uu} \end{pmatrix}.$$

The RHS is given by

$$\mathbf{Z}_{ijt} = \begin{pmatrix} \mathbf{Z}_{ijt}^{1} & 0 & 0 & 0 \\ 0 & \mathbf{Z}_{ijt}^{2} & 0 & 0 \\ 0 & 0 & \mathbf{Z}_{ijt}^{3} & 0 \\ 0 & 0 & 0 & \mathbf{Z}_{ijt}^{4} \end{pmatrix}$$

with

$$\begin{aligned}
\mathbf{Z}_{ijt}^{1} &= \begin{pmatrix} z_{ijt}^{11} & z_{ijt}^{12} & z_{ijt}^{13} & z_{ijt}^{14} \end{pmatrix} = \begin{pmatrix} \widetilde{\tau}_{it}^{u}\widetilde{m}_{ijt}^{ue} & \widetilde{\tau}_{jt}^{u}\widetilde{m}_{ijt}^{eu} & \widetilde{s}_{it}^{e}\widetilde{s}_{jt}^{e} & \widetilde{m}_{ijt}^{ee} \end{pmatrix} \\
\mathbf{Z}_{ijt}^{2} &= \begin{pmatrix} z_{ijt}^{21} & z_{ijt}^{22} & z_{ijt}^{23} & z_{ijt}^{24} \end{pmatrix} = \begin{pmatrix} \widetilde{\tau}_{it}^{e}\widetilde{m}_{ijt}^{ee} & \widetilde{\tau}_{jt}^{u}\widetilde{m}_{ijt}^{uu} & \widetilde{s}_{it}^{u}\widetilde{s}_{jt}^{e} & \widetilde{m}_{ijt}^{ue} \end{pmatrix} \\
\mathbf{Z}_{ijt}^{3} &= \begin{pmatrix} z_{ijt}^{31} & z_{ijt}^{32} & z_{ijt}^{33} & z_{ijt}^{34} \end{pmatrix} = \begin{pmatrix} \widetilde{\tau}_{it}^{u}\widetilde{m}_{ijt}^{uu} & \widetilde{\tau}_{jt}^{e}\widetilde{m}_{ijt}^{ee} & \widetilde{s}_{it}^{e}\widetilde{s}_{jt}^{u} & \widetilde{m}_{ijt}^{eu} \end{pmatrix} \\
\mathbf{Z}_{ijt}^{4} &= \begin{pmatrix} z_{ijt}^{21} & z_{ijt}^{22} & z_{ijt}^{23} & z_{ijt}^{24} \end{pmatrix} = \begin{pmatrix} \widetilde{\tau}_{it}^{e}\widetilde{m}_{ijt}^{eu} & \widetilde{\tau}_{jt}^{e}\widetilde{m}_{ijt}^{ue} & \widetilde{s}_{it}^{u}\widetilde{s}_{jt}^{u} & \widetilde{m}_{ijt}^{uu} \end{pmatrix}
\end{aligned}$$

and the coefficient vector,

$$\mathbf{b}_{ij} = \begin{pmatrix} {\beta_{ij}^{1}}' \\ {\beta_{ij}^{2}}' \\ {\beta_{ij}^{3}}' \\ {\beta_{ij}^{4}}' \end{pmatrix}$$

with

$$\begin{aligned}
\beta_{ij}^{1} &= \begin{pmatrix} \beta_{ij}^{11} & \beta_{ij}^{12} & \beta_{ij}^{13} & \beta_{ij}^{14} \end{pmatrix} \\
&= \begin{pmatrix} \min\left[\left(\alpha_{ij}^{ee}/\alpha_{ij}^{ue}\right),1\right] & \min\left[\left(\alpha_{ij}^{ee}/\alpha_{ij}^{eu}\right),1\right] & \lambda_{ij}\alpha_{ij}^{ee} & -\delta_{ij}\left(1-\alpha_{ij}^{ee}\right) \end{pmatrix} \\
\beta_{ij}^{2} &= \begin{pmatrix} \beta_{ij}^{21} & \beta_{ij}^{22} & \beta_{ij}^{23} & \beta_{ij}^{24} \end{pmatrix} \\
&= \begin{pmatrix} \min\left[\left(\alpha_{ij}^{ue}/\alpha_{ij}^{ee}\right),1\right] & \min\left[\left(\alpha_{ij}^{ue}/\alpha_{ij}^{uu}\right),1\right] & \lambda_{ij}\alpha_{ij}^{ue} & -\delta_{ij}\left(1-\alpha_{ij}^{ue}\right) \end{pmatrix} \\
\beta_{ij}^{3} &= \begin{pmatrix} \beta_{ij}^{31} & \beta_{ij}^{32} & \beta_{ij}^{33} & \beta_{ij}^{34} \end{pmatrix} \\
&= \begin{pmatrix} \min\left[\left(\alpha_{ij}^{eu}/\alpha_{ij}^{uu}\right),1\right] & \min\left[\left(\alpha_{ij}^{eu}/\alpha_{ij}^{ee}\right),1\right] & \lambda_{ij}\alpha_{ij}^{eu} & -\delta_{ij}\left(1-\alpha_{ij}^{eu}\right) \end{pmatrix} \\
\beta_{ij}^{4} &= \begin{pmatrix} \beta_{ij}^{41} & \beta_{ij}^{42} & \beta_{ij}^{43} & \beta_{ij}^{44} \end{pmatrix} \\
&= \begin{pmatrix} \min\left[\left(\alpha_{ij}^{uu}/\alpha_{ij}^{eu}\right),1\right] & \min\left[\left(\alpha_{ij}^{uu}/\alpha_{ij}^{ue}\right),1\right] & \lambda_{ij}\alpha_{ij}^{uu} & -\delta_{ij}\left(1-\alpha_{ij}^{uu}\right) \end{pmatrix}
\end{aligned}$$

The coefficient matrix implies the following sets of constraints- which ensure $\alpha_{ij}^{ll} \leq 1$,

$$\begin{gathered}
1 \geq \beta_{ij}^{11} > 0, \quad 1 \geq \beta_{ij}^{12} > 0, \quad \beta_{ij}^{13} > 0, \quad \beta_{ij}^{14} < 0, \\
1 \geq \beta_{ij}^{21} > 0, \quad 1 \geq \beta_{ij}^{22} > 0, \quad \beta_{ij}^{23} > 0, \quad \beta_{ij}^{24} < 0, \\
1 \geq \beta_{ij}^{31} > 0, \quad 1 \geq \beta_{ij}^{32} > 0, \quad \beta_{ij}^{33} > 0, \quad \beta_{ij}^{34} < 0, \\
1 \geq \beta_{ij}^{41} > 0, \quad 1 \geq \beta_{ij}^{42} > 0, \quad \beta_{ij}^{43} > 0, \quad \beta_{ij}^{44} < 0,
\end{gathered}$$

$$
\begin{aligned}
\beta_{ij}^{11} &= \frac{\beta_{ij}^{13}}{\beta_{ij}^{23}} \text{ if } \frac{\beta_{ij}^{14}}{\beta_{ij}^{24}} > 1 \text{ and } \beta_{ij}^{11} = 1 \text{ otherwise,} \\
\beta_{ij}^{21} &= \frac{\beta_{ij}^{23}}{\beta_{ij}^{13}} \text{ if } \frac{\beta_{ij}^{24}}{\beta_{ij}^{14}} > 1 \text{ and } \beta_{ij}^{21} = 1 \text{ otherwise,} \\
\beta_{ij}^{31} &= \frac{\beta_{ij}^{33}}{\beta_{ij}^{43}} \text{ if } \frac{\beta_{ij}^{34}}{\beta_{ij}^{44}} > 1 \text{ and } \beta_{ij}^{31} = 1 \text{ otherwise,} \\
\beta_{ij}^{41} &= \frac{\beta_{ij}^{43}}{\beta_{ij}^{33}} \text{ if } \frac{\beta_{ij}^{44}}{\beta_{ij}^{34}} > 1 \text{ and } \beta_{ij}^{41} = 1 \text{ otherwise,} \\
\beta_{ij}^{12} &= \frac{\beta_{ij}^{13}}{\beta_{ij}^{33}} \text{ if } \frac{\beta_{ij}^{14}}{\beta_{ij}^{34}} > 1 \text{ and } \beta_{ij}^{12} = 1 \text{ otherwise,} \\
\beta_{ij}^{22} &= \frac{\beta_{ij}^{23}}{\beta_{ij}^{43}} \text{ if } \frac{\beta_{ij}^{24}}{\beta_{ij}^{44}} > 1 \text{ and } \beta_{ij}^{22} = 1 \text{ otherwise,} \\
\beta_{ij}^{32} &= \frac{\beta_{ij}^{33}}{\beta_{ij}^{13}} \text{ if } \frac{\beta_{ij}^{34}}{\beta_{ij}^{14}} > 1 \text{ and } \beta_{ij}^{32} = 1 \text{ otherwise,} \\
\beta_{ij}^{42} &= \frac{\beta_{ij}^{43}}{\beta_{ij}^{23}} \text{ if } \frac{\beta_{ij}^{44}}{\beta_{ij}^{24}} > 1 \text{ and } \beta_{ij}^{42} = 1 \text{ otherwise,}
\end{aligned}
$$

$$
\begin{aligned}
& \beta_{ij}^{14} \text{ is free,} \\
& \beta_{ij}^{24} \text{ is free,} \\
\beta_{ij}^{34} &= \frac{\beta_{ij}^{14}\beta_{ij}^{23} - \beta_{ij}^{24}\beta_{ij}^{13}}{\beta_{ij}^{23} - \beta_{ij}^{13}} - \frac{\beta_{ij}^{24} - \beta_{ij}^{14}}{\beta_{ij}^{23} - \beta_{ij}^{13}}\beta_{ij}^{33}, \\
\beta_{ij}^{44} &= \frac{\beta_{ij}^{14}\beta_{ij}^{23} - \beta_{ij}^{24}\beta_{ij}^{13}}{\beta_{ij}^{23} - \beta_{ij}^{13}} - \frac{\beta_{ij}^{24} - \beta_{ij}^{14}}{\beta_{ij}^{23} - \beta_{ij}^{13}}\beta_{ij}^{43}. \\
\lambda_{ij} &= \frac{\beta_{ij}^{14}\beta_{ij}^{23} - \beta_{ij}^{24}\beta_{ij}^{13}}{\beta_{ij}^{24} - \beta_{ij}^{14}} \\
-\delta_{ij} &= \frac{\beta_{ij}^{14}\beta_{ij}^{23} - \beta_{ij}^{24}\beta_{ij}^{13}}{\beta_{ij}^{23} - \beta_{ij}^{13}}
\end{aligned}
$$

Or alternatively,

$$\begin{aligned}
&\beta_{ij}^{24} \text{ is free,} \\
&\beta_{ij}^{34} \text{ is free,} \\
\beta_{ij}^{14} &= \frac{\beta_{ij}^{24}\beta_{ij}^{33} - \beta_{ij}^{34}\beta_{ij}^{23}}{\beta_{ij}^{33} - \beta_{ij}^{23}} - \frac{\beta_{ij}^{24} - \beta_{ij}^{34}}{\beta_{ij}^{33} - \beta_{ij}^{23}}\beta_{ij}^{13}, \\
\beta_{ij}^{44} &= \frac{\beta_{ij}^{24}\beta_{ij}^{33} - \beta_{ij}^{34}\beta_{ij}^{23}}{\beta_{ij}^{33} - \beta_{ij}^{23}} - \frac{\beta_{ij}^{24} - \beta_{ij}^{34}}{\beta_{ij}^{33} - \beta_{ij}^{23}}\beta_{ij}^{43}, \\
\lambda_{ij} &= \frac{\beta_{ij}^{24}\beta_{ij}^{33} - \beta_{ij}^{34}\beta_{ij}^{23}}{\beta_{ij}^{24} - \beta_{ij}^{34}}, \\
-\delta_{ij} &= \frac{\beta_{ij}^{24}\beta_{ij}^{33} - \beta_{ij}^{34}\beta_{ij}^{23}}{\beta_{ij}^{33} - \beta_{ij}^{23}},
\end{aligned}$$

or

$$\begin{aligned}
&\beta_{ij}^{34} \text{ is free,} \\
&\beta_{ij}^{44} \text{ is free,} \\
\beta_{ij}^{14} &= \frac{\beta_{ij}^{34}\beta_{ij}^{43} - \beta_{ij}^{44}\beta_{ij}^{33}}{\beta_{ij}^{43} - \beta_{ij}^{33}} - \frac{\beta_{ij}^{34} - \beta_{ij}^{44}}{\beta_{ij}^{43} - \beta_{ij}^{33}}\beta_{ij}^{13} \\
\beta_{ij}^{24} &= \frac{\beta_{ij}^{34}\beta_{ij}^{43} - \beta_{ij}^{44}\beta_{ij}^{33}}{\beta_{ij}^{43} - \beta_{ij}^{33}} - \frac{\beta_{ij}^{34} - \beta_{ij}^{44}}{\beta_{ij}^{43} - \beta_{ij}^{33}}\beta_{ij}^{23} \\
\lambda_{ij} &= \frac{\beta_{ij}^{34}\beta_{ij}^{43} - \beta_{ij}^{44}\beta_{ij}^{33}}{\beta_{ij}^{34} - \beta_{ij}^{44}} \\
-\delta_{ij} &= \frac{\beta_{ij}^{34}\beta_{ij}^{43} - \beta_{ij}^{44}\beta_{ij}^{33}}{\beta_{ij}^{43} - \beta_{ij}^{33}}
\end{aligned}$$

The α_{ij}^{ll} are obtained by the following equations,

$$\begin{aligned}
\alpha_{ij}^{ee} &= \frac{\beta_{ij}^{24}-\beta_{ij}^{14}}{\beta_{ij}^{24}-\beta_{ij}^{14}\beta_{ij}^{21}} \text{ if } \beta_{ij}^{11}=1 \text{ and } \beta_{ij}^{11}\frac{\beta_{ij}^{14}-\beta_{ij}^{24}}{\beta_{ij}^{14}-\beta_{ij}^{24}\beta_{ij}^{11}} \text{ otherwise,}\\
\alpha_{ij}^{ue} &= \frac{\beta_{ij}^{14}-\beta_{ij}^{24}}{\beta_{ij}^{14}-\beta_{ij}^{24}\beta_{ij}^{11}} \text{ if } \beta_{ij}^{21}=1 \text{ and } \beta_{ij}^{21}\frac{\beta_{ij}^{24}-\beta_{ij}^{14}}{\beta_{ij}^{24}-\beta_{ij}^{14}\beta_{ij}^{21}} \text{ otherwise,}\\
\alpha_{ij}^{eu} &= \frac{\beta_{ij}^{44}-\beta_{ij}^{34}}{\beta_{ij}^{44}-\beta_{ij}^{34}\beta_{ij}^{41}} \text{ if } \beta_{ij}^{31}=1 \text{ and } \beta_{ij}^{31}\frac{\beta_{ij}^{34}-\beta_{ij}^{44}}{\beta_{ij}^{34}-\beta_{ij}^{44}\beta_{ij}^{31}} \text{ otherwise,}\\
\alpha_{ij}^{uu} &= \frac{\beta_{ij}^{34}-\beta_{ij}^{44}}{\beta_{ij}^{34}-\beta_{ij}^{44}\beta_{ij}^{31}} \text{ if } \beta_{ij}^{41}=1 \text{ and } \beta_{ij}^{41}\frac{\beta_{ij}^{44}-\beta_{ij}^{34}}{\beta_{ij}^{44}-\beta_{ij}^{34}\beta_{ij}^{41}} \text{ otherwise.}
\end{aligned}$$

Given α_{ij}^{ll} one can obtain the λ_{ij} and δ_{ij} from the third and fourth row.

The aggregate number of marriages is according to our theory (compare the outflow from singlehood in equation (3.13)) given by,

$$\widetilde{\Lambda}_t = \xi \left(\widetilde{s}_{mt}\right)^{-\frac{1}{2}} \left(\widetilde{s}_{ft}\right)^{-\frac{1}{2}} \iiiint \alpha_{ij}^{l''l} \widetilde{s}_{it}^{l''} \widetilde{s}_{jt}^{l} didjdl''dl,$$

the aggregate number of divorces (compare the inflow into singlehood in equation (3.13)) by

$$\begin{aligned}
\widetilde{\Delta}_t &= \delta \iiiint \left(1-\alpha_{ij}^{l''l}\right) \widetilde{m}_{ijt}^{l''l} didjdl''dl\\
&\quad + \iiiint \widetilde{\tau}_{it}^{l''} \max\left[1-\left(\alpha_{ij}^{l'l}/\alpha_{ij}^{l''l}\right), 0\right] \widetilde{m}_{ijt}^{l''l} didjdl''dl\\
&\quad + \iiiint \widetilde{\tau}_{jt}^{l} \max\left[1-\left(\alpha_{ij}^{l''l'}/\alpha_{ij}^{l''l}\right), 0\right] \widetilde{m}_{ijt}^{l''l} didjdl''dl.
\end{aligned}$$

Bibliography

Abowd, John M., Francis Kramarz, and David N. Margolis (1999) "High Wage Workers and High Wage Firms," *Econometrica*, Vol. 67, pp. 251 – 333.

Adjemian, Stéphane, Houtan Bastani, Frédéric Karamé, Michel Juillard, Junior Maih, Ferhat Mihoubi, George Perendia, Johannes Pfeifer, Marco Ratto, and Sébastien Villemot (2014) "Dynare: Reference Manual Version 4," Dynare Working Papers 1, CEPREMAP.

Aigner, Dennis, C.A.Knox Lovell, and Peter Schmidt (1977) "Formulation and Estimation of Stochastic Frontier Production Function Models," *Journal of Econometrics*, Vol. 6, pp. 21 – 37.

Amato, Paul R and Brett Beattie (2011) "Does the unemployment rate affect the divorce rate? An analysis of state data 1960–2005," *Social Science Research*, Vol. 40, pp. 705 – 715.

Andolfatto, David (1996) "Business Cycles and Labor-Market Search," *American Economic Review*, Vol. 86, pp. 112–132.

Andrews, M. J., L. Gill, T. Schank, and R. 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)*, Vol. 171, pp. pp. 673–697.

Andrews, M.J., L. Gill, T. Schank, and R. Upward (2012) "High Wage Workers Match with High Wage Firms: Clear Evidence of the Effects of Limited Mobility Bias," *Economics Letters*, Vol. 117, pp. 824 – 827.

Atakan, Alp E. (2006) “Assortative Matching with Explicit Search Costs,” *Econometrica*, Vol. 74, pp. pp. 667–680.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2006) “The Polarization of the U.S. Labor Market,” *American Economic Review*, Vol. 96, pp. 189–194.

Bagger, Jesper and Rasmus Lentz (2016) “An Empirical Model of Wage Dispersion with Sorting,” unpublished.

Bartolucci, Cristian, Francesco Devicienti, and Ignacio Monzón (2015) “Identifying Sorting in Practice,” unpublished.

Bauer, Anja and Benjamin Lochner (2016) “History Dependence in Wages and the Role of Match Quality: Evidence from Germany,” unpublished.

Beaudry, Paul and John DiNardo (1991) “The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data,” *Journal of Political Economy*, Vol. 99, pp. 665 – 688.

Becker, Gary S. (1973) “A Theory of Marriage: Part I,” *Journal of Political Economy*, Vol. 81, pp. 813 – 846.

——— (1974) “A Theory of Marriage: Part II,” *Journal of Political Economy*, Vol. 82, pp. S11–S26.

Blanchard, Olivier and Jordi Galí (2010) “Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment,” *American Economic Journal: Macroeconomics*, Vol. 2, pp. 1–30.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2016) “A Distributional Framework for Matched Employer Employee Data,” unpublished.

Bundesagentur für Arbeit (1988) *Klassifizierung der Berufe - Systematisches und Alphabetisches Verzeichnis der Berufsbenennungen*, Bundesagentur für Arbeit, Nürnberg.

Burdett, Ken and Melvyn G. Coles (1997) “Marriage and Class,” *The Quarterly Journal of Economics*, Vol. 112, p. 141.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin (2006) "Wage Bargaining with On-the-Job Search: Theory and Evidence," *Econometrica*, Vol. 74, pp. 323 – 364.

Calvo, Guillermo A. (1983) "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, Vol. 12, pp. 383–398.

Card, David, Jörg Heining, and Patrick Kline (2013) "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, Vol. 128, pp. 967–1015.

Chade, Hector, Jan Eeckhout, and Lones Smith (2017) "Sorting through Search and Matching Models in Economics," *Journal of Economic Literature*, Vol. 55, pp. 493–544.

Chiappori, Pierre-André, Murat Iyigun, and Yoram Weiss (2009) "Investment in Schooling and the Marriage Market," *American Economic Review*, Vol. 99, pp. 1689–1713.

Choo, Eugene and Aloysius Siow (2006) "Who Marries Whom and Why," *Journal of Political Economy*, Vol. 114, pp. 175–201.

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt (2016) "Unemployment and Business Cycles," *Econometrica*, Vol. 84, pp. 1523 – 1569.

Christoffel, Kai and Tobias Linzert (2005) "The Role of Real Wage Rigidity and Labor Market Frictions for Unemployment and Inflation Dynamics," European Central Bank Working Paper 556.

Cornwell, Christopher, Peter Schmidt, and Robin C. Sickles (1990) "Production frontiers with cross-sectional and time-series variation in efficiency levels," *Journal of Econometrics*, Vol. 46, pp. 185 – 200.

Costain, James S. and Michael Reiter (2008) "Business Cycles, Unemployment Insurance, and the Calibration of Matching Models," *Journal of Economic Dynamics and Control*, Vol. 32, pp. 1120–1155.

Diamond, Peter A. (1982) "Wage Determination and Efficiency in Search Equilibrium," *Review of Economic Studies*, Vol. 49, pp. 217–227.

Doepke, M. and M. Tertilt (2016) "Chapter 23 - Families in Macroeconomics," Vol. 2 of Handbook of Macroeconomics: Elsevier, pp. 1789 – 1891.

Dustmann, Christian, Bernd Fitzenberger, Uta Schönberg, and Alexandra Spitz-Oener (2014) "From Sick Man of Europe to Economic Superstar: Germany's Resurgent Economy," *Journal of Economic Perspectives*, Vol. 28, pp. 167–88.

Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg (2009) "Revisiting the German Wage Structure," *Quarterly Journal of Economics*, Vol. 124, pp. 843–881.

Eeckhout, Jan and Philipp Kircher (2011) "Identifying Sorting–In Theory," *Review of Economic Studies*, Vol. 78, pp. 872 – 906.

——— (2017) "Assortative Matching with Large Firms," forthcoming in *Econometrica*.

Eika, Lasse, Magne Mogstad, and Basit Zafar (2014) "Educational Assortative Mating and Household Income Inequality," Working Paper 20271, National Bureau of Economic Research.

Evers, Lisa, Christoph Spengel, and Julia Braun (2015) "Fiscal Investment Climate and the Cost of Capital in Germany and the EU," ZEW policy brief, ZEW Mannheim.

Faia, Ester (2008) "Optimal Monetary Policy Rules with Labor Market Frictions," *Journal of Economic Dynamics and Control*, Vol. 32, pp. 1600–1621.

Fernández, Raquel, Nezih Guner, and John Knowles (2005) "Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality," *The Quarterly Journal of Economics*, Vol. 120, pp. 273–344.

Fernández, Raquel and Richard Rogerson (2001) "Sorting and Long-Run Inequality," *The Quarterly Journal of Economics*, Vol. 116, pp. 1305–1341.

Fischer, Gabriele, Florian Janik, Dana Müller, and Alexandra Schmucker (2009) "The IAB Establishment Panel - Things Users Should Know," *Schmollers Jahrbuch: Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften*, Vol. 129, pp. 133–148.

Fisman, Raymond, Sheena S. Iyengar, Emir Kamenica, and Itamar Simonson (2006) “Gender Differences in Mate Selection: Evidence From a Speed Dating Experiment,” *The Quarterly Journal of Economics*, Vol. 121, pp. 673–697.

Fitzenberger, Bernd, Aderonke Osikominu, and Robert Völter (2006) “Imputation Rules to Improve the Education Variable in the IAB Employment Subsample,” *Schmollers Jahrbuch: Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften*, Vol. 126, pp. 405–436.

Gabaix, Xavier and Augustin Landier (2008) “Why has CEO Pay Increased So Much?” *The Quarterly Journal of Economics*, Vol. 123, pp. 49–100.

Ganzer, Andreas, Alexandra Schmucker, Philipp vom Berge, and Anja Wurdack (2016) “Sample of Integrated Labour Market Biographies Regional File 1975-2014 (SIAB-R 7514),” FDZ Data report 01/2017, The Research Data Centre (FDZ) at the Institute for Employment Research (IAB), Nürnberg.

Gartner, Hermann and Christian Holzner (2015) “Wage Posting as a Positive Selection Device: Theory and Empirical Evidence,” CESifo Working Paper 5494, CESifo.

Gautier, Pieter A. and Coen N. Teulings (2006) “How large are search frictions?” *Journal of the European Economic Association*, Vol. 4, pp. 1193–1225.

——— (2015) “Sorting and the Output Loss Due to Search Frictions,” *Journal of the European Economic Association*, Vol. 13, pp. 1136–1166.

Gautier, Pieter A., Coen N. Teulings, and Aico van Vuuren (2010) “On-the-job search, mismatch and efficiency,” *Review of Economic Studies*, Vol. 77, pp. 245–272.

Gertler, Mark and Antonella Trigari (2009) “Unemployment Fluctuations with Staggered Nash Wage Bargaining,” *Journal of Political Economy*, Vol. 117, pp. 38 – 86.

Goldschmidt, Deborah, Wolfram Klosterhuber, and Johannes F. Schmieder (2017) “Identifying couples in administrative data,” *Journal for Labour Market Research*, Vol. 50, pp. 29 – 43.

Goldschmidt, Deborah and Johannes F. Schmieder (2015) "The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure," Working Paper 21366, National Bureau of Economic Research.

González-Val, Rafael and Miriam Marcén (2017a) "Divorce and the business cycle: a cross-country analysis," *Review of Economics of the Household*, Vol. 15, pp. 879 – 904.

——— (2017b) "Unemployment, marriage and divorce," *Applied Economics*, pp. 1 – 14.

Goussé, Marion, Nicolas Jacquemet, and Jean-Marc Robin (2017) "Marriage, Labor Supply, and Home Production," forthcoming in *Econometrica*.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos (2016) "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation," *American Economic Journal: Macroeconomics*, Vol. 8, pp. 1–41.

Greenwood, Jeremy, Nezih Guner, and Guillaume Vandenbroucke (2017) "Family Economics Writ Large," Working Paper 23103, National Bureau of Economic Research.

Greenwood, Jeremy, Ananth Seshadri, and Guillaume Vandenbroucke (2005a) "The Baby Boom and Baby Bust," *American Economic Review*, Vol. 95, pp. 183–207.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu (2005b) "Engines of Liberation," *The Review of Economic Studies*, Vol. 72, pp. 109–133.

Guler, Bulent, Fatih Guvenen, and Giovanni L. Violante (2012) "Joint-search theory: New opportunities and new frictions," *Journal of Monetary Economics*, Vol. 59, pp. 352 – 369.

Gustafsson, Siv (1992) "Separate taxation and married women's labor supply," *Journal of Population Economics*, Vol. 5, pp. 61 – 85.

Haefke, Christian, Marcus Sonntag, and Thijs van Rens (2013) "Wage Rigidity and Job Creation," *Journal of Monetary Economics*, Vol. 60, pp. 887 – 899.

Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii (2017) "Identifying Equilibrium Models of Labor Market Sorting," *Econometrica*, Vol. 85, pp. 29 – 65, forthcoming.

Hagedorn, Marcus and Iourii Manovskii (2008) "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *The American Economic Review*, Vol. 98, pp. 1692 – 1706.

——— (2013) "Job Selection and Wages over the Business Cycle," *The American Economic Review*, Vol. 103, pp. 771 – 803.

Hall, Robert E. (2005) "Employment Fluctuations with Equilibrium Wage Stickiness," *The American Economic Review*, Vol. 95, pp. 50 – 65.

Hall, Robert E. and Paul R. Milgrom (2008) "The Limited Influence of Unemployment on the Wage Bargain.," *American Economic Review*, Vol. 98, pp. 1653 – 1674.

Hansen, Hans-Tore (2005) "Unemployment and marital dissolution: A panel data study of Norway," *European Sociological Review*, Vol. 21, pp. 135 – 148.

Heining, Jörg, Peter Jacobebbinghaus, Theresa Scholz, and Stefan Seth (2012) "Linked-Employer-Employee Data from the IAB: LIAB-Mover-Model 1993-2008 (LIAB MM 9308)," FDZ-Datenreport 01/2012, The Research Data Centre (FDZ) at the Institute for Employment Research (IAB), Nürnberg.

Hitsch, Gunter J., Ali Hortaçsu, and Dan Ariely (2010) "Matching and Sorting in Online Dating," *American Economic Review*, Vol. 100, pp. 130–63.

Hodrick, Robert J. and Edward C. Prescott (1997) "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking*, Vol. 29, pp. 1 – 16.

Hornstein, Andreas, Per Krusell, and Giovanni Violante (2005) "Unemployment and Vacancy Fluctuations in the Matching Model: Inspecting the Mechanism," *Federal Reserve of Richmond Economic Quarterly*, Vol. 91, pp. 19–51.

Hornstein, Andreas, Per Krusell, and Giovanni L. Violante (2011) "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *American Economic Review*, Vol. 101, pp. 2873–2898.

Hosios, Arthur J. (1990) "On the Efficiency of Matching and Related Models of Search and Unemployment," *Review of Economic Studies*, Vol. 57, pp. 279–298.

Jacquemet, Nicolas and Jean-Marc Robin (2013) "Assortative matching and search with labor supply and home production," CeMMAP working papers CWP07/13, Centre for Microdata Methods and Practice, Institute for Fiscal Studies.

Jensen, Peter and Nina Smith (1990) "Unemployment and marital dissolution," *Journal of Population Economics*, Vol. 3, pp. 215 – 229.

Jung, Philip and Moritz Kuhn (2014) "Labour Market Institutions and Worker Flows: Comparing Germany and the US," *The Economic Journal*, Vol. 124, pp. 1317 – 1342.

Kantenga, Kory and Tzuo Hann Law (2016) "Sorting and Wage Inequality," unpublished.

Kemeny, John G. and J. Laurie Snell (1962) *Mathematical Models in the Social Sciences*: Blaisdell, pp.145.

Kenyon-Mathieu, Claire and Warren Schudy (2007) "How to Rank with Few Errors," in *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '07, pp. 95 – 103, New York, NY, USA.

Kölling, Arnd (2000) "The IAB-Establishment Panel," *Schmollers Jahrbuch: Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften*, Vol. 120, pp. 291–300.

Krause, Michael U. and Thomas A. Lubik (2007) "The (Ir)Relevance of Real Wage Rigidity in the New Keynesian Model with Search Frictions," *Journal of Monetary Economics*, Vol. 54, pp. 706–727.

Kremer, Michael (1997) "How Much does Sorting Increase Inequality?" *The Quarterly Journal of Economics*, Vol. 112, pp. 115–139.

Lindenlaub, Ilse and Fabien Postel-Vinay (2016) "Multidimensional Sorting Under Random Search," unpublished.

Lise, Jeremy, Costas Meghir, and Jean-Marc Robin (2016) "Matching, Sorting and Wages," *Review of Economic Dynamics*, Vol. 19, pp. 63–87.

Lise, Jeremy and Fabien Postel-Vinay (2016) "Multidimensional Skills, Sorting, and Human Capital Accumulation," unpublished.

Lise, Jeremy and Jean-Marc Robin (2017) "The Macrodynamics of Sorting between Workers and Firms," *American Economic Review*, Vol. 107, pp. 1104–35.

Lopes de Melo, Rafael (2016) "Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence," *Journal of Political Economy*, forthcoming.

Marimon, Ramon and Fabrizio Zilibotti (1999) "Unemployment vs. Mismatch of Talents: Reconsidering Unemployment Benefits," *The Economic Journal*, Vol. 109, pp. 266–291.

Merz, Monika (1995) "Search in the Labor Market and the Real Business Cycle," *Journal of Monetary Economics*, Vol. 36, pp. 269–300.

Mortensen, Dale T. (1982) "The Matching Process as a Noncooperative Bargaining Game," in John McCall ed. *The Economics of Information and Uncertainty*: UMI, pp. 233–258.

Mortensen, Dale T. and Christopher A. Pissarides (1994) "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, Vol. 61, pp. 397–415.

Müller, Steffen (2008) "Capital Stock Approximation Using Firm Level Panel Data," *Jahrbücher für Nationalökonomie und Statistik*, Vol. 224, pp. 357–371.

Newville, Matthew, Till Stensitzki, Daniel B. Allen, and Antonino Ingargiola (2014) "LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python," URL: `https://doi.org/10.5281/zenodo.11813`, DOI: `http://dx.doi.org/10.5281/zenodo.11813`.

Nick, Murat and P. Randall Walsh (2007) "Building the Family Nest: Premarital Investments, Marriage Markets, and Spousal Allocations," *The Review of Economic Studies*, Vol. 74, pp. 507–535.

Nöldeke, Georg and Thomas Tröger (2009) "Matching Heterogeneous Agents with a Linear Search Technology," unpublished.

Petrongolo, Barbara and Christopher A. Pissarides (2001) "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, Vol. 39, pp. 390 – 431.

Petrosky-Nadeau, Nicolas and Lu Zhang (2017) "Solving the Diamond-Mortensen-Pissarides model accurately," *Quantitative Economics*, Vol. 8, pp. 611 – 650.

Pissarides, Christopher A. (1985) "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages," *American Economic Review*, Vol. 75, pp. 676–690.

——— (2000) *Equilibrium Unemployment Theory*: MIT Press, 2nd edition.

——— (2009) "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica*, Vol. 77, pp. 1339–1369.

Postel-Vinay, Fabien and Jean-Marc Robin (2002) "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, Vol. 70, pp. 2295 – 2350.

Pries, Michael J. (2008) "Worker Heterogeneity and Labor Market Volatility in Matching Models," *Review of Economic Dynamics*, Vol. 11, pp. 664–678.

Robin, Jean-Marc (2011) "On the Dynamics of Unemployment and Wage Distributions," *Econometrica*, Vol. 79, pp. 1327 – 1355.

Sax, Christoph (2017) *seasonal: R Interface to X-13-ARIMA-SEATS*, URL: `https://CRAN.R-project.org/package=seasonal`, R package version 1.6.1.

Schaller, Jessamyn (2013) "For richer, if not for poorer? Marriage and divorce over the business cycle," *Journal of Population Economics*, Vol. 26, pp. 1007 – 1033.

Schmidt, Peter and Robin C. Sickles (1984) "Production Frontiers and Panel Data," *Journal of Business & Economic Statistics*, Vol. 2, pp. 367–374.

Shimer, Robert (2005) "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *The American Economic Review*, Vol. 95, pp. 25 – 49.

Shimer, Robert and Lones Smith (2000) "Assortative Matching and Search," *Econometrica*, Vol. 68, pp. 343 – 369.

Sorkin, Isaac (2015) "Ranking Firms Using Revealed Preference," unpublished.

Teulings, Coen N. (2005) "Comparative Advantage, Relative Wages, and the Accumulation of Human Capital," *Journal of Political Economy*, Vol. 113, pp. 425–461.

Teulings, Coen N. and Pieter A. Gautier (2004) "The Right Man for the Job," *Review of Economic Studies*, Vol. 71, pp. 553–580.

Topkis, Donald M. (1998) *Supermodularity and Complementarity*: Princeton University Press.