What hides behind the German labor market miracle? Unemployment insurance reforms and labor market dynamics*

Benjamin Hartung† Philip Jung‡ Moritz Kuhn§

September 2020

Abstract

A key question in labor market research is how the unemployment insurance system affects unemployment rates and labor market dynamics. We provide new answers to this old question by studying one of the largest unemployment insurance reforms in recent decades, the German *Hartz reforms*. On average, lower separation rates into unemployment account for 76% of declining unemployment after the reform, a fact unexplained by existing research focusing on job-finding rates. Exploiting differences in treatment intensity by age, employment duration, and wages, we establish a causal link between the reform and changes in labor market dynamics. We demonstrate the consistency of our empirical findings with labor market theory where workers trade off wages against job stability. Counterfactual simulations show that absent the reform, unemployment rates would have been 50% higher a decade after the reform.

JEL-Classification: E24, J63, J64

Keywords: Unemployment insurance, labor market flows, endogenous separations

---

*We thank Rudi Bachmann, Christian Bayer, Marcus Hagedorn, Simon Jäger, Leo Kaas, Tom Krebs, Iourii Manovskii, Kurt Mitman, Thijs van Rens, Joost Roettgers, Benjamin Schöfer, and Harald Uhlig for their comments and remarks. We thank Philipp vom Berge for help regarding the IAB data. We also thank seminar and conference participants at various places for their helpful comments and remarks. The usual disclaimer applies. Hartung, Jung, and Kuhn thank the German Research Foundation (*Deutsche Forschungsgemeinschaft*) for financial support of this research project as part of the priority program SPP 1764 *The German Labour Market in a Globalised World: Challenges through Trade, Technology, and Demographics*. Jung also thanks the DFG SFB 823 for financial support. This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

†European Central Bank, Benjamin.Hartung@ecb.europa.eu

‡TU Dortmund and IZA, Philip.Jung@tu-dortmund.de

§University of Bonn, ECONtribute, CEPR, and IZA, mokuhn@uni-bonn.de

1 Introduction

A key question in labor market research is how the unemployment insurance (UI) system affects unemployment rates and labor market dynamics. We revisit this old question and provide new answers based on an analysis of one of the largest UI reforms in industrialized countries in recent decades: the German *Hartz reforms*. Economists have extensively studied how changes in the UI system affect job-finding rates either through their incentive effects on unemployed workers when searching for new jobs (Katz and Meyer (1990) and Schmieder and Von Wachter (2016)) or through their incentive effects on firms when posting new vacancies (Millard and Mortensen (1997), Krause and Uhlig (2012), Hagedorn et al. (2013)).[1] In this paper, we scrutinize the existing focus on job-finding rates (unemployment outflows) and draw attention to separation rates into unemployment (unemployment inflows). While the link between separation rates and the UI system is known in theory, little is known about its quantitative importance (Tuit and van Ours (2010)). The goal of this paper is to fill this void.

The Hartz reforms in Germany took place in the mid-2000s. In the decade after the reform, unemployment rates were cut in half. At the heart of the reform was an overhaul of the UI system that abolished long-term, wage-dependent unemployment assistance benefits. Using social security microdata, we highlight that three-quarters of the large decline in German unemployment rates after the reform resulted from lower separation rates into unemployment, while the increase in job-finding rates only accounts for the remainder. We document a large level of heterogeneity in the changes in separation rates across worker groups, with the largest reduction for long-term employed, high-wage workers. We establish a causal link between the UI reform and these observed changes in labor market dynamics by tracking the observed heterogeneity to differences in treatment intensities induced by the institutional features of the reform. In a second step, we demonstrate that these empirical results are also qualitatively and quantitatively consistent with a theoretical model of frictional labor markets. In the model, the response to less generous UI benefits is an increased willingness of workers to accept lower wages in exchange for lower separation rates, trading off wages against job stability to reduce the risk of unemployment. We derive analytically that the key determinant of a strong reaction of separation

[1] The existing literature on job search incentives builds on theoretical grounds in the large body of literature studying the (optimal) design of UI systems. This literature focuses on the trade-off between providing insurance and the cost of additional unemployment due to reduced search effort (Baily (1978), Shavell and Weiss (1979), Hopenhayn and Nicolini (1997), and Chetty (2006)). Recently, researchers have shown renewed interest in quantifying the incentive effects for firms' vacancy postings in relation to changes in UI benefits during the Great Recession in the United States (Hagedorn et al. (2013), Hagedorn et al. (2015), Chodorow-Reich and Karabarbounis (2019)) and Sweden (Fredriksson and Söderström, 2020).

rates to UI reforms is the level of the job-finding rate because low job-finding rates imply long average unemployment spells (see also Jung and Kuhn (2014)). A reduction in UI generosity in countries with low job-finding rates, such as Germany and most European countries, will therefore lower separation rates more strongly compared to countries with high job-finding rates and short average unemployment spells, such as the United States where effects on separation rates tend to be small, in line with Hagedorn et al. (2013).

This paper contributes to a growing literature that explores the *German labor market miracle* (Burda and Seele (2016)) by explaining two striking features of German unemployment dynamics: first, the reduction in unemployment rates by half within less than a decade starting in 2005, and second, the very small increase in unemployment rates during the Great Recession. We provide empirical and theoretical evidence that the Hartz reforms were the main driver behind this labor market miracle. What distinguishes our explanation from the existing literature is the focus on changes in separation rates into unemployment. Existing research that studies the Hartz reforms focuses on job-finding rates as the key margin of adjustment by highlighting changes in search effort (Krebs and Scheffel (2013)), changes in matching efficiency (Launov and Wälde (2013), Hertweck and Sigrist (2015), and Klinger and Weber (2016)), changes in labor supply (Carrillo-Tudela et al., 2018), changes in employer hiring standards (Hochmuth et al., 2019), or changes in vacancy posting behavior (Krause and Uhlig (2012)). We argue that the transmission mechanism of the Hartz reforms not only is of academic interest but also implies very different welfare effects across subgroups of workers in the labor force. In particular, we show that the reform has resulted in substantial welfare losses for the large group of long-term *employed* workers who have a very low risk of becoming unemployed. Our explanation therefore provides a potential explanation for the widespread discontent with the reforms in the population despite the massive reduction in the unemployment rate.

For our empirical analysis, we rely on social security microdata of individual employment histories in West Germany from the Sample of Integrated Labour Market Biographies (SIAB). We construct worker flow rates for one decade before and after the Hartz reforms and find that separation rates declined by 28% after the reform, while job-finding rates increased by only 13%. As a consequence, changes in separation rates account for 76% of the decline in unemployment rates. This stylized fact is robust to a wide range of sensitivity checks and is also found using alternative data sources. The average decline in separation rates hides a lot of heterogeneity that we exploit to establish a link from the UI reform to changes in labor market dynamics. The first dimension of heterogeneity consists of changes in maximum benefit duration that imply different treatment intensities by age and employment duration by the Hartz reforms. We find a statistically significant effect of these maximum benefit duration changes on separation rates supporting a causal effect

of the UI reform. As a second dimension of heterogeneity, we exploit that the lower bound on UI benefits provided by social assistance benefits remained unchanged by the reform, which turns the group of low-wage workers into a natural control group for the impact of the reform as their benefit level remained unaffected. We provide support for these implications in the data and find heterogeneous responses of separation rates in line with the differentiated impact of the reform on worker groups. With respect to heterogeneity by employment duration, we show that separation rates of long-term employed workers fell by up to 60%, while short-term employed workers show a comparatively modest decline of 20% in their separation rates.[2] For changes in benefit duration, we estimate an elasticity ranging between 0.51 and 0.58 on separation rates. Using these estimates, we account for the additional decline in separation rates of older, long-term employed workers. Along the wage distribution, as the second dimension of heterogeneity, we document for the control group of low-wage workers in the bottom decile of the wage distribution no change in separation rates and 20% lower separation rates for the median-wage worker. For the latter result, we restrict the sample to workers without changes in maximum benefit duration to rule out confounding effects from changes in maximum benefit duration.

To support the mechanism of UI benefit changes on separation rates, we rely on economic theory. We develop a labor market search model with worker heterogeneity, aggregate fluctuations, and endogenous separation decisions. Workers in the model differ in their employment status, skills, job duration, wages, and UI benefit eligibility. Worker skills increase with job duration, and individually efficient bargaining over wages and separation decisions implies that high-skill workers are also high-wage workers in stable jobs. Our model incorporates key institutional features of Germany's UI benefit eligibility rules with respect to the dependence on employment duration and wages, as in Krause and Uhlig (2012).[3] Our model also incorporates all three channels from the literature on how UI reforms affect labor market dynamics: workers' incentives to search and accept job offers, firms' incentives to post vacancies, and the decision of workers and firms to separate. Endogenous separation decisions lead to falling separation rates after a reduction in UI generosity (Pissarides, 2000, Ch.2). We calibrate the model to the pre-reform period and introduce the Hartz reforms by abolishing long-term wage-dependent benefits and shortening the benefit duration for long-term employed workers. After the reform, the

[2] Jäger et al. (2018) explore a staggered extension of UI benefit durations by age on older male workers in Austria. In line with our findings, they find large increases in separation rates due to increased benefit generosity.

[3] We share several modeling choices with Krause and Uhlig (2012) but differ in our focus. Their findings and calibration strategy focus on changes in job-finding rates through the effects on vacancy postings, rendering separation rates effectively exogenous in their quantitative analysis. Their model also does not include aggregate fluctuations to impose discipline on the elasticity of separation and job-finding rates, which we exploit for the calibration as described below.

model closely matches the observed time series for average separation and job-finding rates. We demonstrate that the model also closely matches the empirically documented heterogeneous responses. In the model, as in the data, the long-term employed, high-wage workers are most adversely affected, and the model-implied elasticity of separation rates with respect to benefit duration aligns well with our empirical estimates. We use the model to perform counterfactual simulations of the German labor market in the absence of the reform. Absent the reform, the model predicts skyrocketing unemployment rates during the Great Recession and unemployment rates that would have been 50% higher by 2014 than what has been observed in the data. We compare this prediction to a synthetic control estimate for German unemployment rates absent the Hartz reforms (Abadie et al., 2010). We find that the model-implied unemployment rate in the absence of the reform closely tracks the estimated counterfactual.

In the model, the UI reform affects workers' search incentives, firms' incentives to post vacancies, and separation decisions. The model structure imposes no predetermined relative importance on the different channels, so the question arises on how to discipline the relative importance of these three adjustment channels. In theory, there is a tight link between aggregate labor market fluctuations from productivity fluctuations and the responsiveness to changes in UI benefits (Costain and Reiter (2008a)). Through the lens of the model, productivity changes and benefit changes both directly affect the value of employment relative to the outside option so that pre-reform business cycle fluctuations inform the key reform elasticity of separation rates with respect to changes in UI benefits. Based on this insight, we calibrate the model to be consistent with business cycle moments for separation rates and job-finding rates before the Hartz reforms. For the responsiveness of workers' search behavior, we target existing estimates on the elasticity of the search intensity to changes in UI benefits from the empirical literature.[4] Our calibration only targets unconditional moments of worker flow rates but closely matches the time series dynamics of labor market flows before the reform, thereby providing support for the model mechanism. After the reform, the model still matches the time series of labor market flow rates very closely, lending support to the independently calibrated elasticities. Using a stylized model framework, we analytically derive the relationship between structural reforms and business cycle elasticities and identify low average job-finding rates as the key determinant of a strong reaction of separation rates to UI reforms. The latter result reconciles our findings with results on the U.S. labor market that highlight the important role of changes in job-finding rates, for example, Hagedorn et al. (2013).

[4] A broad empirical consensus has emerged suggesting that this effect is modest. Typical estimates find that granting one additional month of UI benefits leads to 0.15 more months of unemployment (Chetty (2006), Schmieder and Von Wachter (2016)).

We also provide empirical support for the model's underlying mechanism of a trade-off between wages and job stability. Using cross-sectional regressions of separation rates on wages, we find that before the reform, workers experiencing a productivity increase were indifferent between a 1% wage increase and a 1% lower separation rate. After the reform, the trade-off turns toward job stability with workers being indifferent between a 0.8% wage increase and a 1% lower separation probability. After the reform, workers' willingness to trade wages for job stability (a lower separation rate) increased.

In a final step, we use our microfounded framework to quantify the welfare effects of the reform for different labor market participants. We consider welfare effects abstracting from compensating transfers that the government could finance because of the lower spending on UI benefits after the reform. Put differently, we quantify how a transfer system needs to be designed to avoid welfare losses for all worker groups after the reform. This question is key when it comes to the political feasibility of UI reforms. We find that losses amount to 2.1% in terms of consumption-equivalent variation for the recipients of unemployment assistance benefits. Unemployment assistance benefits represent the long-term, wage-dependent benefits that have been abolished by the reform, so the large welfare losses for workers in this group ought to be expected. These losses also probably explain the widespread grandfathering rules and hardship regulation that accompanied the reform and which were targeted toward this group.[5] Among the employed, we find the largest welfare losses among the long-term employed, high-wage workers. We find that their consumption-equivalent variation to forgo the reform amounts to 0.64%. Long-term employed workers account for almost two-thirds of the German labor market, and the fact that their separation rates are the lowest among the employed might suggest that these workers are very detached from any changes in the UI system. Yet, we show that this is not the case and that in hindsight, their large welfare costs might explain the widespread discontent in the population with the reform.

Two potentially important policy implications for labor market and social security reforms arise from our findings. The first relates to UI reform proposals in other European countries taking the Hartz reforms as a role model. Regarding the political feasibility of such reforms, our findings imply that appropriate compensation schemes have to be designed to avoid discontent in large parts of the electorate, as we show that a quantitatively important role for changes in separation rates should be expected in most European countries. Second, the strong reaction of separation rates after changes in nonemployment benefits suggests that similar reactions ought to be expected and taken into account when evaluating other social security reforms such as early retirement programs or disability

[5]We document that during a transition period, supplementary benefits to cushion the impact of the reform were prevalent and only leveled off by 2008.

insurance programs that are widely discussed in Germany and elsewhere.

The remainder of the paper is structured as follows. We next provide a short description of the Hartz reforms. In Section 2, we describe our data and present the empirical results. We describe the labor market search model in Section 3. Section 4 shows the model results and discusses the counterfactual analysis. Before we conclude in Section 5, we discuss alternative explanations for the German labor market miracle in light of our empirical results.

1.1 The Hartz reforms

In 2002, the German government entrusted an expert commission consisting of various representatives from business, unions, and academia with the task of working out reforms for the German labor market. The chairman was Peter Hartz, at that time director of human resources at Volkswagen. The subsequent reforms are commonly referred to as the *Hartz reforms*.[6] The main focus of the reforms was to restructure the federal employment agency and enhance the matching process of unemployed workers to jobs. The ensuing reforms were enacted in four separate legislative packages commonly referred to as *Hartz I* to *Hartz IV* between 2003 and 2005.[7] They consisted of comprehensive measures to *promote and challenge* the unemployed — ranging from subsidies for self-employment to the restructuring of the unemployment benefit system and a tighter supervision of benefit recipients.[8] We provide further details of the reform steps in Appendix A.

We focus on the fourth step of the reform package (*Hartz IV*) and provide empirical evidence for its large effects on labor market dynamics and unemployment rates. In that step, the former three-tier system of unemployment benefits, unemployment assistance, and subsistence benefits was transformed into a two-tier system of unemployment and subsistence benefits. The reform constituted a substantial overhaul of the German UI system and implied a drastic cut in benefits for long-term employed workers who, before the reform, were eligible for long-term, wage-dependent unemployment assistance. After the reform, they only received subsistence benefits once unemployment benefits expired. We discuss the heterogeneity in the impact of the reform in detail below.

2 Data and empirical results

This section introduces the microdata for analyzing changes in unemployment rates and labor market flows. In a first step, we demonstrate that the microdata closely match

[6] The official title of the commission was the *Commission for Modern Labor Market Services*.

[7] The official title of the acts were *First, Second, Third,* and *Fourth Act for Modern Labor Market Services*.

[8] Steffen (2008) provides a detailed chronicle of the German social security system.

the macroeconomic trends, and we explain how we adjust for administrative changes that were part of the Hartz reforms and which would otherwise impede a consistent measurement over time. In the second step, we present empirical results on changes in labor market flows, document a large level of heterogeneity in these changes, and exploit this heterogeneity to establish a causal link from the UI reform to changes in labor market dynamics.

2.1 Data

Our main data source is the microdata on individual employment histories from the Sample of Integrated Labour Market Biographies (SIAB) provided by the Institute for Employment Research (IAB) for the period from 1975 to 2014.[9] The SIAB is a 2% representative sample of administrative data on all workers who are subject to social security contributions and on all unemployed workers in Germany. It excludes self-employed and civil servants, thus covering approximately 80% of Germany's labor force. Apart from its large size (1.8 million individuals) and its long panel dimension (up to 40 years), one further advantage of the administrative data is that they are virtually free of measurement error for the variables of interest in this paper. The data are taken from social security records and are merged with records on unemployment periods from the federal employment agency. The data contain the exact start and end dates of each employment and unemployment spell. In total, the data comprise almost 60 million individual spells. See Antoni et al. (2016) for further details on the data.

2.2 Sample selection, construction of worker flow rates, and inflow correction

We restrict our sample to workers in West Germany and exclude marginal employment in our benchmark sample. We drop a few individuals with missing information on employment status or missing geographic information, and all individuals who only receive social assistance benefits while in the sample. We consider the effect of including marginal employment and results for East Germany in our sensitivity analysis (Appendix C).

The data contain daily employment histories, and we follow Jung and Kuhn (2014) to aggregate daily labor market histories to histories at a monthly frequency. We assign monthly employment spells based on a reference week within each month. We report as the separation rate the share of employed workers entering into unemployment from one month to the next (unemployment inflows) and as the job-finding rate the share of unem-

[9] We use the weakly anonymous Sample of Integrated Labour Market Biographies (SIAB), 1975-2014. The data were accessed on-site at the Research Data Centre (FDZ) of the Federal Employment Agency (BA) at the Institute for Employment Research (IAB) and via remote data access at the FDZ.

ployed workers entering into employment between months (unemployment outflows). We assign the employment state in the reference week following a hierarchical ordering where employment supersedes unemployment and unemployment supersedes out of the labor force. This approach closely follows labor force surveys such as the Current Population Survey (CPS) for the United States. We count workers as employed if they are employed full- or part-time or work as apprentices. We count workers as unemployed if they are registered as unemployed at the employment agency, which requires that they are actively looking for a job. Registration is required to be eligible for unemployment benefits. The German unemployment insurance system distinguishes between unemployed workers and benefit recipients. In the microdata, reliable information on the registered unemployment status is available from 2000 onward. We use this information to assign employment states. We assign employment states for earlier periods based on records of benefit-recipient status and compute worker flow rates based on benefit-recipient status before 2000. We construct growth rates of these worker flow rates before the year 2000 and use these growth rates to extend the registration-based flow rates starting in the year 2000 backward. This leaves the dynamics of the flow rates unaffected but removes the level differences between the two definitions. We provide further details on the construction of monthly employment states and transition rates in Appendix B. For our empirical analysis, we focus on the decade from 1993 to 2002 to document worker flows before the first reform steps were implemented. We report the entire time series of worker flows for the period after the reform but take only the time period from 2008 to 2014 as the period when the transition period after the reform was completed. In Appendix D.2, we document the prevalence of supplementary benefits to cushion the consequences of the reform during the years after the reform and show that the number of supplementary benefit recipients fell strongly between 2005 to 2008 when they leveled off and were completely abolished by the end of 2010.

The goal of our empirical analysis is to study the changes in labor market dynamics that determine the evolution of the unemployment rate. We demonstrate first that the microdata match the macroeconomic trends of unemployment rates. The microdata do not include public servants (*Beamte*), and hence, for the microdata to be consistent with the reported unemployment rates by the German employment office, public servants have to be included. Figure 1(a) shows the unemployment rate for West Germany as reported by the German federal employment agency and the unemployment rate constructed from the SIAB microdata for the period between 1993 and 2014.[10] Both unemployment rates

[10]The German employment office reports two unemployment rates. The unemployment rate for dependent employment that we rely on excludes self-employed workers. The employment office also reports an unemployment rate including all employees.

track each other closely in trend and levels, so we rely on them to study the underlying changes in labor market dynamics. In Appendix B, we demonstrate that using the constructed worker flow rates in a two-state stock-flow model matches the dynamics of the unemployment rate over time very well. We also consider a three-state model of unemployment with flows in and out of the labor force but find no notable improvement in accounting for the dynamics of the unemployment rate compared to the two-state model.

Figure 1: German unemployment rates (1993-2014)

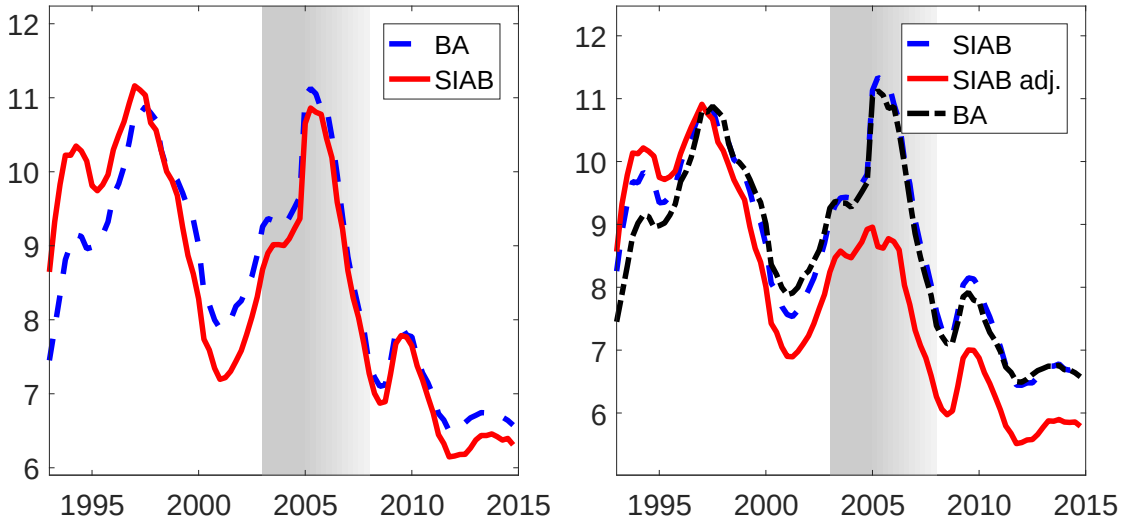

(a) BA and SIAB (b) SIAB with inflow correction

Notes: Unemployment rates for West Germany, 1993-2014, in percentages. Left panel: Unemployment rate by employment agency (BA) (blue dashed line) and unemployment rate from SIAB microdata including imputed numbers for public servants not covered by the microdata (red solid line). Right panel: Unemployment rate from SIAB microdata and employment agency as in the left panel (dashed blue and black lines) and unemployment rate from SIAB microdata after inflow correction (solid red line). See text for details. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms. Data are quarterly averages of monthly rates.

The data in Figure 1(a) show a large spike in unemployment in January 2005. The spike reflects regulatory changes in the UI system as part of the Hartz reforms that became effective in January 2005. These regulatory changes required all nonemployed who are able to work to register as unemployed to remain eligible for UI benefits. This change caused a large inflow of former social assistance recipients and spouses of unemployed into the unemployment pool and poses a challenge to obtaining a consistent measurement of worker flows over time. To account for this effect, we propose an *inflow correction* for constructing comparable and consistent transition and unemployment rates for this period.

The key challenge for this adjustment is that we cannot directly observe workers who

Table 1: Worker characteristics of entrants into unemployment

	entrants from N			other U	
	2004-01	2005-01	2005-01 (corr.)	2004-01	2005-01
female	43.3%	60.9%	45.8%	41.1%	42.1%
age	36.9	37.3	36.0	40.9	40.9
high school	23.2%	44.2%	32.5%	16.6%	18.8%
vocational training	70.4%	53.0%	62.9%	78.0%	76.0%
college	6.5%	2.9%	4.6%	5.5%	5.2%

Notes: Demographic characteristics of workers who transit to unemployment from out of the labor force (*entrants from N*) or all other states (*other U*) in January 2004 and 2005. The column for the entrants from N labeled *corr.* applies the inflow correction. See text for details. Row *female* shows the share of females in inflows, row *age* shows average age, and the bottom three rows show the shares of workers with at most high school education, vocational training, and a college education.

were forced to register as unemployed to retain their unemployment benefit eligibility. We therefore exclude persons who simultaneously satisfy three conditions: (1) entered unemployment in the first six months of 2005,[11] (2) had a nonemployment spell before registering as unemployed, and (3) did not work for at least one month until the end of 2006. We compare in Table 1 the characteristics of new entrants into unemployment from out of the labor force in January 2004 and January 2005.[12] We find large differences across the two years. Comparing columns 1 and 2 of Table 1, we observe that in January 2005, new entrants are slightly older, substantially more female (61% versus 43%), and less educated (44% versus 23% with high school or less). When looking at all other entrants into unemployment (columns *other U*), we find that worker characteristics do not differ notably for this other group of workers in January 2004 and 2005. Our inflow correction excludes entrants into the unemployment pool in early 2005 who are very detached from the labor market and are likely to have registered as unemployed solely because of the new registration requirements in 2005. Comparing the composition of the inflows, the first two columns of Table 1 suggest that a large group of entrants from out of the labor force in January 2005 falls into this category. The third column, *entrants*

[11] There is evidence that administrative problems and incomplete data records during the transition period make the records for the affected group in the first months after the reform less reliable.

[12] Out of the labor force is not directly observed in the data, and we assign out of the labor force as a residual employment state to nonemployed workers who have intermittent nonemployment spells that are not unemployment spells.

from N, reports worker characteristics for entrants after the inflow correction. We find that after the inflow correction, the worker characteristics of entrants in 2005 resemble those of the entrants in 2004 much more closely, although some differences still remain. We refer to the sample after excluding these persons as the *inflow-corrected sample*, and we will use this sample as our benchmark sample for the rest of the paper. We provide a sensitivity analysis for skipping the inflow correction in Appendix C.1.

Figure 1(b) shows the unemployment rate of the inflow-corrected sample (solid red line) and the full sample (dashed blue line). The spike in January 2005 disappears almost completely in the inflow-corrected sample. In Section 4.3, we estimate the German unemployment rate using a synthetic control group composed of OECD countries (Abadie et al., 2010). The estimation is based on observed data until 2003. The estimated unemployment rate for the control group in 2005 closely follows the inflow-corrected unemployment rate, providing additional, independent support to our inflow-correction approach. The persistently lower level of the inflow-corrected sample shows that the inflow of formerly nonemployed persons into the unemployment pool in early 2005 changed the composition toward persons who are less attached to the labor market. Given that we remove these workers completely from the sample, we also change unemployment rates before 2005, but this change is small. In 2014, unemployment rates in the inflow-corrected sample are about 0.75 percentage points lower. Looking at relative changes, we find that the inflow correction reduces the decrease in unemployment rates from roughly 40% to 30%. Still, unemployment rates declined between 2005 and 2014 by more than 30%. As we demonstrate in our sensitivity analysis (Appendix C.1), our key empirical results are reinforced even if we skip the inflow correction.

2.3 Empirical results

We consider the years 2003 and 2004 as the period of the reforms. We consider the years from 1993 to 2002 as representative of the labor market situation before the reform and use the years from 2008 to 2014 as representative of the labor market situation after the reform. Alternatively, we exclude the Great Recession from the post-reform period and consider the period from 2011 to 2014 as representative of the labor market situation after the reform. We consider the period from 2005 to 2008 as transition period when supplementary benefits to cushion the impact of the reform were still very prevalent. We provide further details and discussion on the transition period in Appendix D.2. In total, the sample period includes three recessions and, in particular, the Great Recession. One challenge based on these data alone is to disentangle the relative importance of structural changes in the labor market and changes from business cycle fluctuations when comparing the pre- and post-reform periods. We postpone this question and rely on the structural

model to provide a decomposition that disentangles structural changes and business cycle effects on worker flows and unemployment rates. We also provide an extensive sensitivity analysis of our empirical results, which we summarize at the end of this section. We relegate details to Appendix C.

2.3.1 Changes in separation and job-finding rates

Figure 2(a) shows the relative change in the separation rate for the period from 1993 to 2014. The separation rate is indexed to its average pre-reform level (1993-2002 = 100). This level is low in the German labor market over the entire time period. About 0.5% of workers transit from their employer to unemployment each month (Table 2). Looking at the relative changes, we find a substantial 28% decline in separation rates between the pre-reform average and the separation rate during the post-reform period. When we consider the post-reform average including the Great Recession, the decline is smaller but still at 22%. It is interesting to note that separation rates spiked during the Great Recession, with an increase of about 40% relative to their 2007 level. Such a large spike in separations scrutinizes the narrative of short-term work as an explanation for the modest increase in unemployment rates during the Great Recession. We will return to the experience during the Great Recession in Section 4 when discussing results of the theoretical model that allows us to provide a meaningful counterfactual for this period.

Figure 2: Separation and job-finding rates (1993-2014)

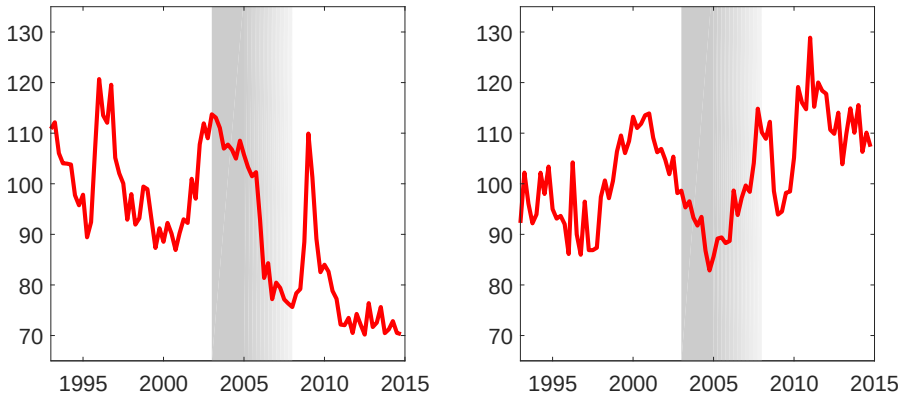

(a) Separation rate (indexed) (b) job-finding rate (indexed)

Notes: Separation and job-finding rates for West Germany, 1993-2014. Both series have been indexed to their pre-reform level (1993-2002). The grey area marks the period 2003 to 2004 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms. Data are quarterly averages of monthly rates.

Figure 2(b) shows the relative change in the job-finding rate over time, again indexed to its average pre-reform level. Job-finding rates are typically slightly above 5% before the reform period and increase to slightly below 6% after the reform. In relative terms, the

increase until 2014 constitutes a 13% increase in the job-finding rate. If we include the Great Recession in the post-reform average, the increase amounts to only 10%. During the Great Recession, job-finding rates declined by 20%, which is a modest decline given the size of the shock, and job-finding rates also recovered quickly compared to previous recessions (Jung and Kuhn, 2014). Compared to the 28% decline in the separation rates, the 13% increase in job-finding rates suggests that declining separation rates were the main driver behind the decline in unemployment rates over the decade following the Hartz reforms. The relative differences in changes remain largely unaffected when we include the Great Recession (22% versus 10%). In both cases, the decline in separation rates is twice as large as the increase in job-finding rates.

Table 2: Before- and after-reform unemployment rates, transition rates, and steady-state decomposition

				2008-2014		2011-2014	
	1993-2002	2008-2014	2011-2014	Δ	$\frac{\Delta\pi}{\Delta\bar{u}}$	Δ	$\frac{\Delta\pi}{\Delta\bar{u}}$
unemployment rate	10.5%	7.6%	7.2%	-28%		-32%	
separation rate	0.6%	0.5%	0.5%	-22%	75%	-28%	76%
job-finding rate	5.2%	5.7%	5.9%	10%	31%	13%	32%

Notes: Columns 2-4 show the level of the unemployment rate, separation rate, and job-finding rate before the Hartz reforms (1993-2002), after the Hartz reforms including the Great Recession (2008-2014), and after the Hartz reforms excluding the Great Recession (2011-2014). Columns labeled Δ report the percentage change in rates from before to after the reforms. Columns labeled $\frac{\Delta\pi}{\Delta\bar{u}}$ show the relative contribution to changes in steady-state unemployment rates from changes in separation and job-finding rates. $\Delta\bar{u}$ indicates the change in the steady-state unemployment rate from before to after the Hartz reforms based on average rates before and after the reform.

Table 2 uses a steady-state decomposition based on a two-state stock-flow model to quantify the relative contribution of separation rates and job-finding rates in explaining the 32% decline in unemployment rates until 2014.[13] We consider the period from 1993 to 2002 as the pre-reform steady state and the two periods from 2008 to 2014 and 2011 to 2014 as the post-reform steady states. The last column of Table 2 reports the relative contributions of changes in the separation rate and the job-finding rate to the unemployment rate. According to this decomposition, the declining separation rate accounts

[13]Here, we use a two-state model so that the steady-state unemployment rate is $\bar{u} = \frac{\bar{\pi}_{eu}}{\bar{\pi}_{eu}+\bar{\pi}_{ue}}$ where $\bar{\pi}_{eu}$ denotes the steady-state separation rate (unemployment inflow) and $\bar{\pi}_{ue}$ denotes the steady-state job-finding rate (unemployment outflow). In Appendix B.3, we demonstrate that two-state and three-state models deliver very similar dynamics of the unemployment rate over time.

for 76% of the decline in the unemployment rate (2011-2014 steady state). The small residual of 4% relative to the empirically observed changes demonstrates that the simple two-state stock-flow model captures the changes in the unemployment rate over time very well. Including the Great Recession in the decomposition leads to the same quantitative findings for the relative importance of separation and job-finding rates for the decline in unemployment (see column labeled "2008-2014" in Table 2).

Existing studies that explore the effect of UI reforms on the labor market focus on the effects on the job-finding rate, either from changes in search effort or from changes in contact rates for unemployed workers from more vacancy postings. The large contribution of changes in the separation rate to changes in the unemployment rate that we document here let such explanations fall short in explaining the German experience.

2.4 Heterogeneity of separation rate changes

The average decline in separation rates is the main driver of the reduction in unemployment rates in Germany after 2005. This average decline hides a lot of the heterogeneity that we trace back to the institutional features of the reform to establish a link from the reform to the observed changes. To explain the sources of this heterogeneity, Figure 3 provides a stylized description of the pre-reform and post-reform UI system in Germany and highlights the institutional features that we exploit in our empirical analysis.

Figure 3(a) sketches the three-tier UI system *before* the reform with UI benefits that are tied to the last wage, unemployment assistance benefits that long-term unemployed workers receive after their unemployment benefits expired, and as the third tier social assistance benefits that were need based at the subsistence level and independent of the last wage. Figure 3(b) sketches the UI system *after* the Hartz reforms that abolished the second tier of unemployment assistance benefits. After the reform, workers for whom UI benefits expire receive social assistance benefits at the subsistence level. Generally, this change applied to all workers, yet institutional features of the UI system led to heterogeneity in the impact of these changes. The first dimension of heterogeneity stems from the maximum benefit duration for which unemployed workers could receive UI benefits as this duration depends on previous employment duration and age. Together with the Hartz reforms, maximum benefit duration was cut differentially by age and employment duration.[14] Figure 3(c) sketches how this change led to heterogeneity in treatment intensity by the reform because especially older, long-term employed workers received, in addition to the abolition of unemployment assistance benefits, a cut in maximum UI ben-

[14]This change became effective in 2006, but as documented in Figure 23 of Appendix D.2, many unemployed still received supplementary benefits during 2005 to cushion the impact of the abolition of unemployment assistance benefits.

Figure 3: Stylized pre- and post-reform UI system and heterogeneous treatment effects

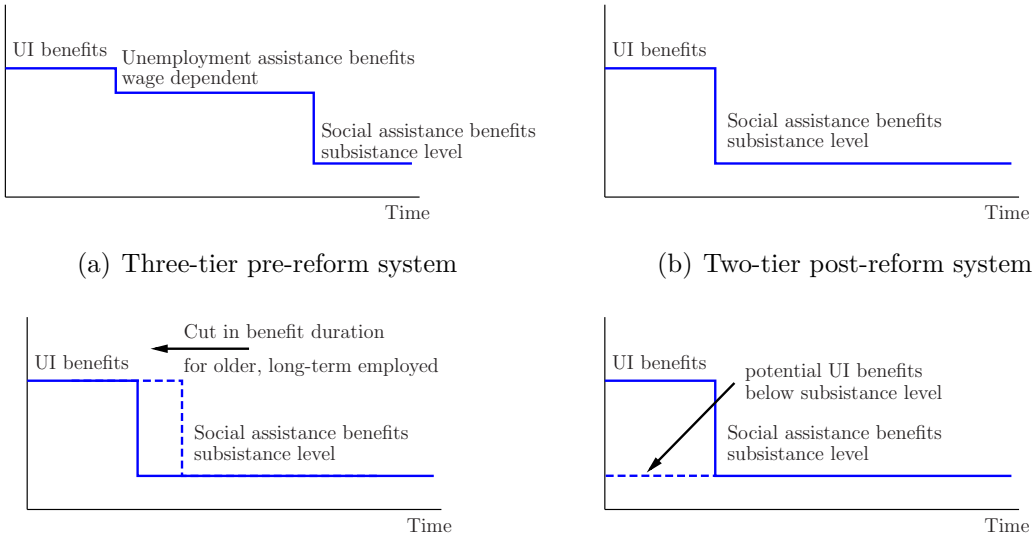

(a) Three-tier pre-reform system

(b) Two-tier post-reform system

(c) Variation in treatment intensity by age and employment duration

(d) Treatment effect and control group

Notes: Stylized pre- and post-reform UI system. The vertical axis shows the qualitative level differences in replacement rates for the average worker. The horizontal axis shows unemployment benefit duration. Top left panel shows three-tier pre-reform UI system. Top left panel shows the two-tier post-reform UI system. Bottom left panel shows the heterogeneity by age and employment duration in treatment intensity arising from a reduction in maximum benefit duration. Bottom right panel shows the case of the control group of low-wage workers with potential benefits below subsistence levels. See text for further details.

efit duration, implying a larger treatment intensity for these workers. Exploiting similar institutional variation in the German UI system by age has been used before to estimate the effect of UI generosity on search behavior (e.g., Schmieder et al. (2012)). We will rely on these estimates to calibrate our structural model in Section 3.

Figure 4 shows maximum unemployment benefit duration by employment duration and age before and after the reform. This benefit duration determines when workers lose eligibility for UI benefits and transit to unemployment assistance benefits before the reform and benefits at the subsistence level after the reform. If the abolition of the unemployment assistance benefits by the Hartz reforms is the driver of the observed changes in separation rates, we should see heterogeneity in the changes of separation rates by employment duration and age, in line with the cuts in maximum benefit duration. Looking at the pre-reform situation in Figure 4(a), we see that for workers younger than 45, the maximum benefit duration was 12 months. For older workers, we find a steep gradient in employment duration from 14 months after 30 months of previous employment to up to 30 months after 60 months of previous employment. Comparing this pattern to

Figure 4: Changes in benefit duration by age and employment duration

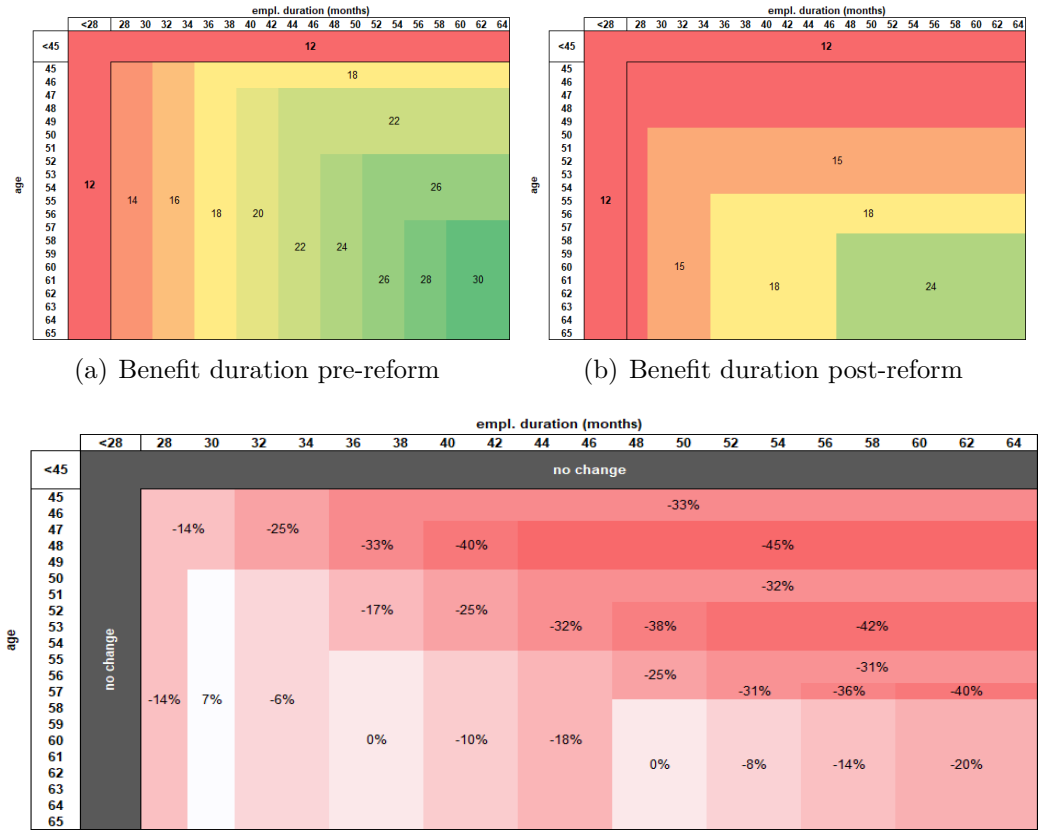

(a) Benefit duration pre-reform

(b) Benefit duration post-reform

(c) Change in benefit duration

Notes: Maximum eligibility duration for short-term unemployment benefits in months by age and employment duration. Employment duration refers to a reference period of 5-7 years prior to the unemployment spell. Panel (a) shows maximum duration before the reform. Panel (b) shows the maximum duration after the reform in 2008 (i.e., after all grandfathering rules had expired). Panel (c) shows the relative change in maximum duration in percentages for each combination of age and employment duration.

the post-reform regulation in Figure 4(b), we see that there is much less variation and that especially older, long-term employed workers see a strong decline in their benefit duration. For example, a 49-year-old worker with four years of previous employment receives, after the reform, UI benefits for up to 12 months, while before the reform she received UI benefits for up to 22 months. Figure 4(c) shows the relative changes in UI benefit durations for the different groups from before to after the reform. We see that the largest decline happened for workers with more than three years of previous employment duration between ages 45 and 55. By contrast, there have been no changes for short-term employed workers (less than 28 months) and workers younger than 45 years.

Figure 3(d) sketches the second dimension along which the institutional design of the UI system created heterogeneity in treatment effects of the Hartz reform. Workers for

whom UI benefits are below the level of need-based social assistance benefits are eligible for supplementary benefits ("Aufstocker") both before and after the reform. For these workers, abolishing wage-dependent unemployment assistance benefits had no effect on their potential UI benefit level because their potential benefit level stayed at the subsistence level and remained unaffected by the reform. In theory, this provides us with a control group for the impact of the reform. In our microdata, we cannot directly identify these workers because need-based subsistence benefits depend on household characteristics that remain unobserved in our data. We instead proxy for the control group by the lowest decile of the wage distribution. In Appendix D.1, we provide evidence supporting this choice. We document that, over time, a stable share of 10% of the unemployed receive supplementary benefits as their UI benefits are below subsistence levels. Relative to this control group, we should see a treatment effect of the reform on workers with higher wages and potential benefit levels above subsistence benefits. We exploit this heterogeneity to further tighten the causal relationship from the reform to the observed changes in separation rates and to demonstrate that no common trend brought down all separation rates in the German labor market. We will focus on workers 45 years and younger when looking at heterogeneity by wages to avoid confounding effects from the cut in maximum benefit duration.

We proceed with our empirical analysis in three steps. In a first step, we provide descriptive evidence for heterogeneous changes in separation rates by age and employment duration. This step provides direct evidence for the economic significance of the heterogeneous effects. In the second step, we rely on regression analysis to establish a statistically significant impact of the UI reform on separation rates. In a third step, we provide evidence for heterogeneity in separation rate changes along the wage distribution. We find for non-treated low-wage workers no change in the separation rates after the implementation of the reform. Combining this evidence, we decompose the total effect on separation rates in an effect from lower benefit duration and a reduction in benefit levels.

2.4.1 Heterogeneity by employment duration and age: Descriptive evidence

We first look at heterogeneous responses for workers with different employment duration. We split employed workers into two groups. The first group is short-term employed workers with at most three years of employment duration, and the second group is long-term employed workers with more than three years of employment duration. This threshold cuts the sample roughly into a first group of workers (short-term employed) who are only treated by the effect from abolishing unemployment assistance benefits and a second group of workers (long-term employed) who also experienced an additional effect from the cut in maximum benefit duration. Table 3 shows the average levels of separation

rates for the pre- and post-reform period for these two groups. Looking at the levels, we see that short-term employed workers have separation rates that are more than five times higher than those of the long-term employed workers in the period 1993 to 2002 (1.37% versus 0.26%). This difference further increases in the period 2008 to 2014 (1.15% versus 0.18%). After 2008, separation rates differ by more than a factor of six. The reason for this difference is the much stronger relative decline in the separation rate for long-term employed workers after 2008. The last column of Table 3 highlights that the decline for long-term employed workers has been twice as large as for short-term employed workers.

Table 3: Change in separation rates by employment duration and age

	1993-2002	2008-2014	Δ %
all	0.63%	0.49%	-22.0%
emp. duration $\leq$ 3 years	1.37%	1.15%	-16.2%
emp. duration $>$ 3 years	0.26%	0.18%	-33.3%

Notes: Monthly separation rates before and after the Hartz reforms by employment duration and age. Column Δ reports the percentage change in rates from the period before the Hartz reforms to the period after the Hartz reforms.

The stronger relative decline can also be seen in Figure 5. Figure 5(a) shows the indexed time series of separation rates for short-term and long-term employed workers. We see direct evidence of a strong divergence in the time series of separation rates between short-term and long-term employed workers after the Hartz reforms. The strong divergence persists so that, after the reform, separation rates of long-term employed workers have declined twice as much as those of short-term employed workers.

In addition to employment duration, we saw in Figure 4 that age determines the maximum benefit duration. In a second step, we further dissect the data in Figures 5(b) and 5(c) by looking at young (44 years and younger) and older workers with different employment durations. Young workers in Figure 5(b) are only affected by the abolition of unemployment benefit assistance but not by the effect from the cut in duration. In line with such a homogeneous treatment effect by the reform, we find no differential changes between short-term and long-term employed young workers, and separation rates decline in lockstep. By contrast, we observe differential treatment effects from changes in eligibility duration in Figure 5(c) when we consider older long-term employed and short-term employed workers. We find the strongest reduction in separation rates for long-term employed, older workers with almost 60% lower separation rates after the reform compared

to their pre-reform average. By contrast, the reduction for older short-term employed workers is only about half as large. Looking at short-term employed workers across age groups in Figure 5(d), we find a strikingly close tracking of separation rate changes for short-term employed young (age 15-44) and short-term employed old workers (age 45-64). For both age groups, separation rates decline in lockstep following the Hartz reforms. Figure 4 suggests exactly such an evolution as the two groups experience homogeneous treatment effects from the reform.

In Appendix C.3, we provide additional details of changes by age groups. One finding from this analysis is that workers closer to retirement show an even stronger decline in separation rates. Their decline in separation rates follows a longer-run trend that accelerated during the 2000s so that, over time, unemployment rates for older workers decreased more than those of younger workers. This trend was accompanied by a strongly rising labor force participation rate of workers close to retirement age (Carrillo-Tudela et al., 2018). We abstract from this fact of independent interest as it is beyond the scope of this paper.[15] These results of heterogeneity in separation rate changes already provide strong descriptive evidence for a reform effect on labor market dynamics. In the second step, we provide systematic regression evidence for such a relationship.

[15] Jäger et al. (2018) provide a detailed investigation of this topic. They study changes in separation rates of male workers toward the end of working life (50 years and older) in Austria after changes in UI benefit duration. They exploit staggered changes in UI eligibility similar to those shown in Figure 4 in combination with regional variation. Jäger et al. rely on a microeconometric analysis to characterize marginal jobs separations after changes in workers' outside options. In line with our empirical results, they document large changes in separation rates after changes in potential benefit duration.

Figure 5: Separation rates by age and employment duration (1993-2014)

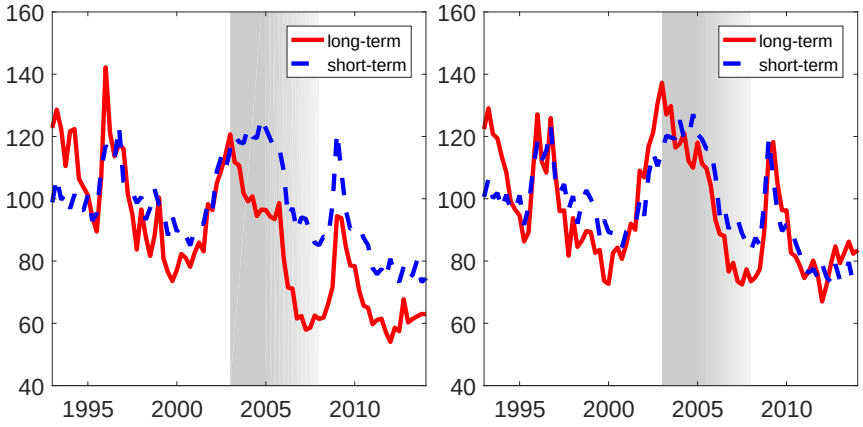

(a) all workers

(b) age 15-44

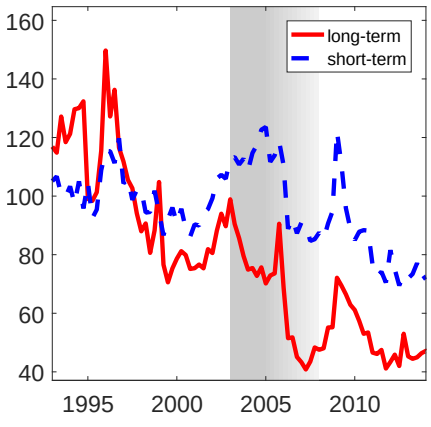

(c) age 45-64

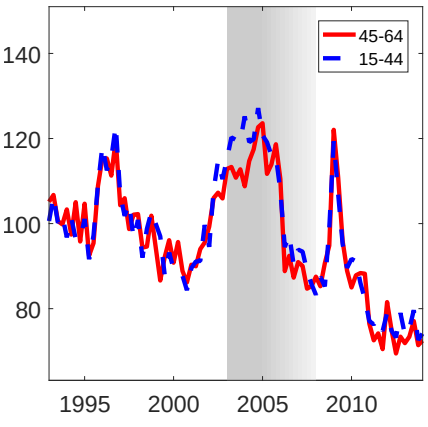

(d) short-term employed

Notes: Separation rates by employment duration and age for West Germany, 1993-2014, indexed to their pre-reform level (1993-2002). The solid red lines in panels (a)-(c) mark the separation rate for long-term employed workers who were continuously employed for three years or more. The dashed blue lines in panels (a)-(c) mark the separation rate for short-term employed workers with at most three years of continuous employment. Panel (d) shows the separation rate for short-term employed workers separately for young (age 15-44, dashed blue line) and old (age 45-64, solid red line) employees. The grey area indicates the period of the implementation of the Hartz reforms. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms. Data are quarterly averages of monthly rates.

2.4.2 Heterogeneity by employment duration and age: Regression evidence

In the next step, we regress changes in separation rates on the differential treatment effects by the reform to explore if there is a systematic relationship between treatment intensity and separation rate changes. We define treatment groups by the cells in Figure 4(c) and define the treatment intensity by the log-difference in maximum benefit eligibility (ΔD^{max}). Young and short-term employed workers with unchanged maximum benefit duration have a treatment intensity of zero ($\Delta D^{max} = 0$) and serve as the control group for this regression. In a first step, we run a flexible specification of treatment effects on (log) separation rates $\log(\pi)$,

$$\log(\pi_{i,t}) = \mu_i + \gamma_t + \sum_{s=1993}^{2014} \Delta D_i^{max} \tau_s + \varepsilon_{i,t}, \tag{1}$$

where μ_i are fixed effects for treatment groups indexed by i, γ_t are year fixed effects, and τ_s denote year-specific treatment intensity coefficients that also cover the pre-reform period (1993-2002). In this flexible specification, we interact treatment intensities ΔD^{max} with year effects to allow for unrestricted reform effects. We weight all observations by the average employment size of the treatment group over time. We report estimation results in Figure 6 where we show the predicted average separation rates for treated ($\Delta D^{max} < 0$) and untreated ($\Delta D^{max} = 0$) workers indexed to the pre-reform period.[16] In line with the descriptive evidence from Figure 5, we estimate a drop for untreated workers of approximately 20%. For treated workers, we find that the cut in benefit duration reduced separation rates by an additional 20%. We further note that the estimated pre-reform treatment effects are negligible, so predicted transition rates for treated and untreated workers before the reform do not show diverging trends but follow a common trend until entering the reform period (2003-2004). Consistent with a causal effect of the reform, we only see a positive treatment effect during the post-reform period (2005-2014) with increasing impact during the transition period (2005-2008).

The specification in equation (1) is highly flexible but has too few observations for statistical significance of treatment effects. We therefore run an alternative specification where we pool data over the pre- and post-reform period and use the log-difference in separation rates $\Delta\pi$ as our outcome variable. We include a constant in the regression that captures the baseline effect for workers with a treatment intensity of zero. Specifically, we run the regression

$$\Delta\pi_i = \beta_0 + \beta_1 \Delta D_i^{max} + \varepsilon_i, \tag{2}$$

[16] We construct the predicted separation rate for treatment group i in year t as $\hat{\pi}_{i,t} = \exp(\hat{\mu}_i + \hat{\gamma}_t + \Delta D_i^{max} \hat{\tau}_t)$.

Figure 6: Average effect of change in benefit duration

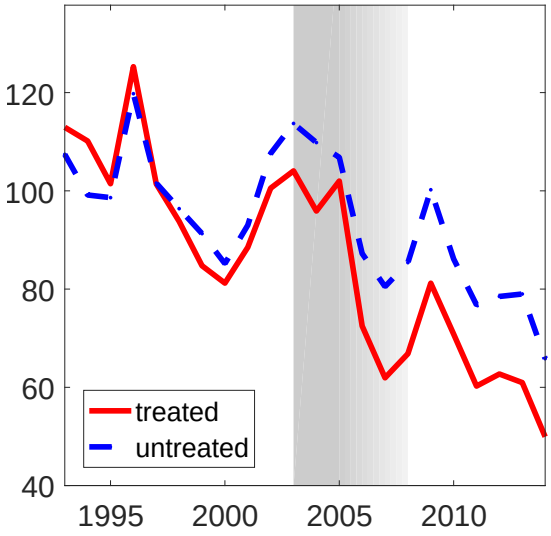

Notes: Predicted separation rates changes from change in maximum benefit duration. Predicted separation rates from equation (1) indexed to pre-reform period (1993−2002 = 100). Separation rates averaged across treated ($\Delta D^{max} < 0$) and untreated worker groups ($\Delta D^{max} = 0$). Solid red line shows average change across treated worker groups. Dashed blue line shows average change across untreated workers.

where i again identifies the different treatment groups. Note that the specification in first differences already takes out fixed characteristics across treatment cells. As before, we expect separation rates to decline on average (negative β_0) and to fall in treatment intensity (positive treatment coefficient β_1). Table 4 reports the estimated regression coefficients for four different specifications that differ with respect to the post-reform period either including or excluding the Great Recession and if treatment groups are weighted by their average employment size.

Looking at the regressions in columns 1 and 3 that include the Great Recession, we find a negative β_0 coefficient that is slightly larger than 0.2 in absolute value. This implies that, on average, separation rates of non-treated workers declined by approximately 20%, in line with the effects in Figures 5(b) and 5(d). If we exclude the Great Recession in columns 2 and 4, the coefficients decrease by 9 percentage points, consistent with the descriptive analysis. Weighting observations by employment has a negligible effect on the estimated coefficients as the unweighted regression results in columns 3 and 4 show. Across all specifications, we find β_0 to be statistically significant at a 5% level. Importantly, the estimated treatment effect β_1 has the expected positive sign and is statistically significant at the 5% level in the weighted regressions. The estimated coefficients imply an elasticity of separation rates with respect to unemployment benefit duration that is slightly larger than 0.5. Evaluated at the average cut in benefit duration of 33% ($\Delta D^{max} = -0.42$, log

Table 4: Regression of separation rate change on change in benefit duration

	(1)	(2)	(3)	(4)
$\hat{\beta}_0$	−0.21	−0.30	−0.22	−0.32
	(0.06)	(0.07)	(0.10)	(0.10)
$\hat{\beta}_1$	0.54	0.58	0.51	0.52
	(0.23)	(0.25)	(0.33)	(0.33)
post-reform	2008 - 2014	2011 - 2014	2008 - 2014	2011 - 2014
weighted	yes	yes	no	no
obs.	28	28	28	28

Notes: Regression coefficients from regression of separation rate change on change in maximum benefit duration for different specifications. Coefficient estimates $\hat{\beta}_0$ and $\hat{\beta}_1$ for constant and slope coefficient, standard errors below coefficients in parentheses. Row *post-reform* indicates the data years used for the post-reform period. Row *weighted* indicates if observations have been weighted by employment in cell. Row *obs.* indicates number of observations in regression.

difference), we get a treatment effect that lowers separation rates of all treated workers by 23% in addition to the baseline effect of 20% across all workers. These effects are qualitatively and quantitatively consistent with the results from the more flexible specification in the first regression (Figure 6).

To summarize, we estimate a statistically and economically significant treatment effect of the cut in maximum benefit duration on separation rates. We also find a large and significant baseline effect of the reform on average separation rates of all workers. To link this baseline effect to the reform, we provide in the next step further evidence based on heterogeneity along the wage distribution. As argued above, low-wage workers provide a natural control group to estimate the baseline effect from the abolition of unemployment assistance benefits because their potential benefit level did not change as a result of the reform. Consistently, we document that for low-wage workers, there has been no change in separation rates after unemployment assistance benefits have been abolished.

2.4.3 Heterogeneity by wage levels

For the following analysis, we exploit heterogeneity in separation rate changes by wage levels. We report separation rates for worker groups that have been grouped year by year into wage deciles. We pool data at an annual frequency to get precise estimates of transition rates especially for high-wage workers who have low average transition rates into unemployment. The bottom decile of the wage distribution forms our control group

for the abolition of unemployment assistance benefits by the Hartz reforms. While we cannot directly observe potential benefit levels in the data, we consider this group of workers in the lowest wage decile to have the largest overlap with the group of unemployed workers for whom potential unemployment benefits are below subsistence levels so that they are eligible to receive social assistance benefits that remained unchanged by the reform (see Appendix D.1 for further discussion). Their constant potential benefit level turns this group into a natural control group for the impact of the Hartz reforms.

Figure 7(a) compares changes in separation rates of low-wage workers and median-wage workers (fifth decile) over time. Importantly, we consider here only workers age 45 and younger so that the treatment effect only stems from abolishing unemployment assistance benefits and there is no additional treatment effect from a cut in maximum benefit duration as discussed in the previous section.[17] We find that the separation rates of the two groups evolve before 2005 in lockstep and show a common trend. Consistent with abolishing unemployment assistance benefits, we find separation rates for the treatment and control group diverging when the reform became effective in 2005. By 2014, separation rates of treated median-wage workers declined by roughly 20%, consistent with the estimated baseline level effect in equation (2). Hence, we find no change in average separation rates from the decade before the reform to the decade after the reform for our control group of low-wage workers. For these workers, potential benefit levels stayed constant at subsistence levels over the entire time period. For the median worker in the group of workers age 45 years and younger, we find a large drop in separation rates at the time the reform became effective. We associate this drop with the abolition of unemployment assistance benefits by the Hartz reforms.

Figure 7(b) provides a broader view on separation rates changes along the wage distribution. It shows for each wage decile by how much the average separation rate decreased from the decade before the reform (1993-2002) to the post-reform period (2008-2014). We now include workers of all age groups. Evidently, the higher wage deciles experienced the largest declines in separation rates. There is no notable change in separation rates in the lowest wage decile, only a small effect in the second decile, and starting in the third decile to the sixth decile, separation rates decline by between 20% to 30%. At the top, separation rates plummet by 40% to 50%. These groups include most of the older, long-term employed workers who have also been affected by the change in maximum benefit duration that, as we documented before, led on average to an additional decline in separation rates.

The results on employment duration, age, and wage heterogeneity all show a large amount

[17] Treatment intensity is $\Delta D^{max} = 0$ for all workers age 45 and younger.

Figure 7: Changes in separation rates by wages

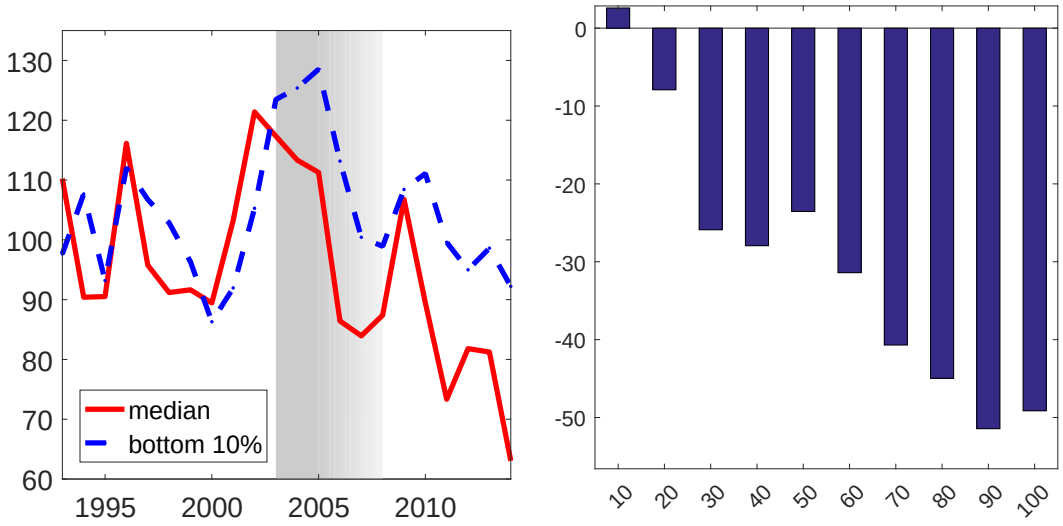

(a) Separation rate of control group (b) Relative change post-reform (in %)

Notes: The left panel shows changes in separation rates for the control group of low-wage workers and the median-wage group of workers in West Germany, 1993-2014. The solid red line shows separation rates of workers in the fifth wage decile of the wage distribution (*median*). The dashed blue line shows separation rates of the control group of workers in the bottom decile of the wage distribution. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms. Data are pooled at the annual level. The right panel shows relative declines in average separation rates for the entire wage distribution from before the reform (1993-2002) to after the reform (2008-2014) in percentages. Deciles of the wage distribution are shown on the horizontal axis.

of heterogeneity in treatment effects on separation rates across worker groups that increase with the treatment intensity of the Hartz reforms. Treatment effects are statistically and economically significant. We speak to this heterogeneity in our quantitative model below where the abolition of long-term wage-dependent benefits will lead to heterogeneous reactions in separation rates, and high-wage, long-term employed workers will see a stronger reduction in their separation rates, in line with the empirical evidence from this section. The model also informs us about the economic mechanism behind declining separation rates. Workers are willing to accept lower wages in exchange for lower separation rates into unemployment. We provide empirical evidence for such a wage-job stability trade-off after we discuss the model and its results.

2.5 Sensitivity and comparison to other data sources

Finally, we summarize the sensitivity results for our empirical analysis. We relegate details to Appendix C. First, we discuss evidence from other independent data sources

to further support our empirical evidence on the dominant role of falling separation rates in explaining the decline in German unemployment rates after 2005. The first additional data source is from the reports of the employment agency on monthly unemployment benefit claims. In Hartung et al. (2016), we construct a historical series on worker flows for the period 1967 to 2014 based on these data and demonstrate that, during the period of overlap, it closely matches worker flows from the SIAB microdata. We explain in Hartung et al. (2016) how this data series can be constructed in real time from publicly available data sources. The second data source is flow rates in and out of unemployment reported by the German employment office since 2006. These flow rates are based on registered cases of workers transiting from employment into unemployment and vice versa. These rates are based on case counts rather than worker counts. To be consistent with our structural model, we use worker counts based on reference weeks for our empirical analysis. This difference in measurement will lead to differences in the level of rates because multiple cases can occur for one worker within one month. This is the well-known time aggregation problem, as discussed, for example, in Shimer (2012).

Figure 8: Alternative measures for transition rates

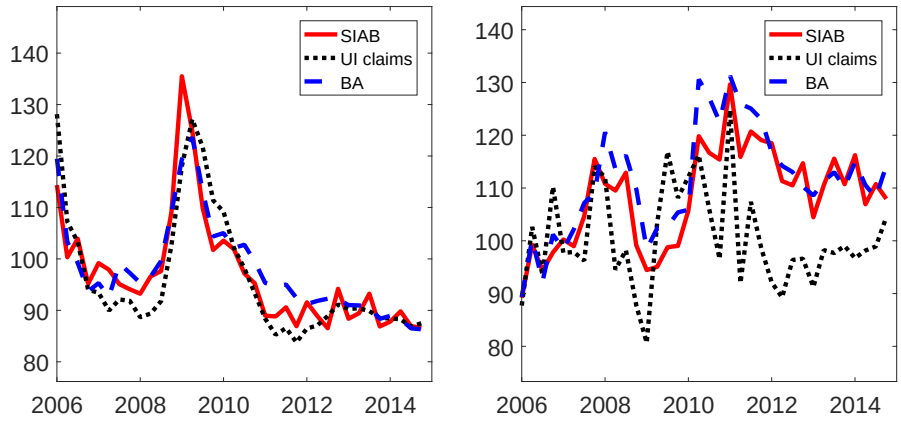

(a) Separation rates (b) job-finding rates

Notes: The figures show separation and job-finding rates for the benchmark sample from the SIAB microdata (solid red line). The dashed blue line shows flow rates reported by the German employment office. The dotted black line shows flow rates constructed in Hartung et al. (2016) based on new unemployment benefit claims. All rates are indexed to the level in the first two years displayed in the graphs (2006-2007). See text for further details.

Figure 8 shows the three alternative measures for the separation rate and job-finding rate. The first one is our benchmark measure constructed from the SIAB microdata (solid red line); the second one is constructed by the German employment office (dashed blue line), the so-called inflow hazard rate (*Zugangsrisiko*) and departure rate (*Abgangschance*); and the third one is the measure constructed from UI benefit claims in Hartung et al. (2016)

(dotted black line). We find that the two additional measures strongly support our finding of decreasing separation rates as the macroeconomic driver of falling unemployment. In particular, the flow rates reported by the German employment office (dashed blue line) track our estimated time series remarkably well.

Appendix C provides details on a battery of additional sensitivity checks that we only summarize here. In a first step of our sensitivity analysis, we demonstrate that skipping the inflow correction mainly leads to lower job-finding rates after the reform because of the larger unemployment pool (see Appendix C.1). In a second step, we control for changes in the composition of the employed in terms of worker characteristics using a linear regression model. Fixing the composition of the employed at the level in 2000, we find that compositional changes alone are too small for explaining changes in separation rates over time (see Appendix C.2). In a third step, we provide results for East Germany (see Appendix C.4), counting marginally employed workers who are registered as unemployed as employed (see Appendix C.5), and counting workers in active labor market programs among the employed (see Appendix C.6).[18] We find the documented results to be robust.

3 Model

This section applies economic theory to explore the link between changes in the unemployment insurance system and changes in labor market dynamics and unemployment rates. We develop a labor market search and matching model with aggregate fluctuations, endogenous separations, and worker heterogeneity.

In the model, time is discrete and there is a continuum of workers of measure one and a positive measure of firms. Workers and firms are risk neutral and discount the future at rate $\tilde{\beta}$. Each period there is a positive probability that a worker leaves the labor force for good. We denote this probability by ω and the product of the time discount factor and the probability of remaining in the labor market by $\beta = \tilde{\beta}(1-\omega)$. A worker who leaves the labor force is immediately replaced by a newborn worker so that there is always a constant mass of workers. Workers in the model are either employed or unemployed. We consider single-worker firms and refer to a worker-firm pair as a match.

Employed workers have one of two skill levels x_1 or x_2 with $x_1 < x_2$. We refer to workers with skill level x_1 as low-skill workers and workers with skill level x_2 as high-skill workers. Workers who enter the labor force start as low skill. While working, workers accumulate skills by learning-by-doing. An employed low-skilled worker stochastically gains skills at rate α. The accumulated skills are lost upon separation. Employed workers become

[18] After the reform, workers who participate in active labor market programs were no longer counted as unemployed.

eligible for unemployment benefits with employment duration. Since the accumulation of skills and benefit eligibility both depend on employment duration, we economize on the state space and assume that eligibility and skill level are perfectly correlated so that all high-skill workers are eligible for unemployment benefits.[19,20] As discussed below, low-skill workers are eligible for social assistance benefits or unemployment benefits if they separate and enter into unemployment. We denote the share of employed workers in the population in state x_1 by e_1 and the share of employed workers in state x_2 by e_2. Denoting the current period's state by x and the next period's state by x', the law of motion for x conditional on staying employed is

$$x' = x_2 \qquad \text{if } x = x_2,$$

and if $x = x_1$, the law of motion is

$$x' = \begin{cases} x_2 & \text{with probability } \alpha \\ x_1 & \text{with probability } 1-\alpha. \end{cases} \tag{3}$$

We denote the state of unemployed workers by b, and the state can take three values b_j with $j = 1, 2, 3$. The different states describe the current eligibility level of the unemployed: social assistance (b_1), unemployment assistance (b_2), and unemployment benefits (b_3). It holds that $b_1 \leq b_2 < b_3$. Upon entering unemployment, high-skill workers are eligible for unemployment benefits b_3. When entering unemployment, low-skill workers enter in state b_3 with probability γ, and with probability $1-\gamma$, they enter unemployment in state b_1. Stochastic eligibility for low-skill workers captures in a parsimonious way the more complex eligibility rules of the actual system.[21] During unemployment, the eligibility state stochastically changes. Workers in state b_3, receiving unemployment benefits, transit to state b_2, receiving unemployment assistance, with probability δ_3. Workers who are in state b_2 transit to state b_1, receiving social assistance, with probability δ_2. We denote the mass of workers in each state by u_j for $j = 1, 2, 3$. Denoting the current

[19] We abstract from age heterogeneity that would lead to the introduction of an additional state variable. The underlying economic mechanism would be identical to the mechanism that works along the employment duration dimension. Krause and Uhlig (2012) follow the same modeling approach.

[20] In general, experience and skill accumulation need not be perfectly correlated. The empirical evidence on wage growth for the German labor market finds strong returns to experience in the first two years (Dustmann and Meghir (2005)). This suggests that productivity gains and eligibility in the data are also highly correlated, so we are confident that our assumption to economize on the state space is of minor importance.

[21] Two main reasons account for the misalignment of employment duration and eligibility. First, employees with more than one year of employment duration are already eligible for UI benefits for a period of 6 months, which then gradually increases to 12 months the longer a person has been working. Second, employment duration in the legislation does not refer to the latest continuous employment spell but the accumulated duration in a reference period that varied between 2 and 7 years.

period's state by b and the next period's state by b', the law of motion for b conditional on staying unemployed is

$$b' = b_1 \qquad \text{if } b = b_1,$$

and if $b = b_j$ for $j = 2, 3$, the law of motion is

$$b' = \begin{cases} b_j & \text{with probability } 1 - \delta_j \\ b_{j-1} & \text{with probability } \delta_j. \end{cases} \tag{4}$$

When unemployed workers reenter employment, they enter with state x_1. The law of motion for the worker state at the transition from unemployment to employment is $x' = x_1$ independent of b. When transiting from employment into unemployment, the law of motion is

$$b' = b_3 \qquad \text{if } x = x_2, \tag{5}$$

and if $x = x_1$, the law of motion is

$$b' = \begin{cases} b_3 & \text{with probability } \gamma \\ b_1 & \text{with probability } 1 - \gamma. \end{cases} \tag{6}$$

Each period consists of two stages. The first stage is the separation stage when each match decides about separating into unemployment or entering the production stage. The second stage is the production stage for the employed and the search stage for the unemployed. Search happens simultaneously with production. We refer to this stage, respectively, as the search or production stage depending on whether the unemployed or the employed are considered. We abstract from on-the-job search. Labor market exit happens with probability ω at the end of the period. A match that does not separate enters the production stage and produces $y = \exp(a + x)$ units of output depending on skill level x and the aggregate productivity state a. The aggregate productivity state a follows an AR(1) process with autocorrelation ρ and variance σ_a^2.

The aggregate state of the economy s comprises the aggregate productivity state a and the distribution of workers over states $s = \{a, e_1, e_2, u_1, u_2\}$ where we dropped u_3 because of the identity $e_1 + e_2 + u_1 + u_2 + u_3 = 1$. The state of a match at the beginning of the period is described by the tuple (x, s) of the idiosyncratic state x and the aggregate state s. The state of an unemployed worker is (b, s), where the idiosyncratic state is the current benefit eligibility.

At the separation stage, each match draws an idiosyncratic cost shock ε and then, depending on the state of the match (x, s), decides whether to enter the production stage. For analytical tractability, we assume that the shock ε is independently and identically

distributed across matches and time and is drawn from a logistic distribution F with mean $\bar{\varepsilon}$ and variance $\sigma_\varepsilon^2 = \pi^2 \frac{\psi_\varepsilon^2}{3}$. A match that decides to separate does not pay these costs. Optimal behavior follows a threshold rule where separations happen when the idiosyncratic cost shock ε is larger than a state-specific threshold $\varepsilon^u(x,s)$. This threshold is determined as part of the bargaining process between the worker and the firm so that separation decisions will be individually efficient. The average separation rate of a match with state (x,s) is $\pi_{eu}(x,s) = Prob(\varepsilon \geq \varepsilon^u(x,s))$. Workers who separate at the separation stage enter unemployment in the current period, receive benefits, and start searching during the search stage of the current period. Aggregate output in a period is $y = \sum_i e_i(1-\pi_{eu}(x_i,s))\exp(a+x_i)$, where $e_i(1-\pi_{eu}(x_i,s))$ is the mass of employed workers of type i who produce at the production stage.[22]

We denote the value of a firm matched to a worker of skill type x before the realization of the idiosyncratic shock ε by $J(x,s)$. The value $J(x,s)$ expressed recursively is

$$J(x,s) = \int_{-\infty}^{\varepsilon^u(x,s)} \Big(\exp(a+x) - \varepsilon - w(x,s) + \beta \mathbb{E}[J(x',s')|x,s] \Big) dF(\varepsilon), \tag{7}$$

where $w(x,s)$ denotes the wage for the worker and expectations are taken over the realization of the idiosyncratic and aggregate state next period (x',s') conditional on the current state (x,s). The upper integration bound is the threshold value $\varepsilon^u(x,s)$ that determines separation. Because of free entry of firms, the continuation value of the firm after separation is zero in equilibrium. Below, we explain how $\varepsilon^u(x,s)$ and $w(x,a)$ are determined. We exploit the properties of the logistic distribution to get a closed-form solution for the integral of the idiosyncratic shocks ε that we denote by $\Psi_\varepsilon(\pi_{eu})$:

$$\Psi_\varepsilon(\pi_{eu}) = \int_{-\infty}^{\varepsilon^u} -\varepsilon dF(\varepsilon) = -(1-\pi_{eu})\bar{\varepsilon} - \psi_\varepsilon \Big((1-\pi_{eu})\log(1-\pi_{eu}) + \pi_{eu}\log(\pi_{eu}) \Big),$$

with $\pi_{eu} = 1 - F(\varepsilon^u)$ denoting the separation probability given the threshold value ε^u. The firm value simplifies to

$$J(x,s) = (1-\pi_{eu}(x,s)) \Big(\exp(a+x) - w(x,s) + \beta \mathbb{E}[J(x',s')|x,s] \Big) + \Psi_\varepsilon(\pi_{eu}(x,s)). \tag{8}$$

The state of an unemployed worker at the beginning of the period is (b,s), with the idiosyncratic state b describing the worker's current benefit level. The worker's flow

[22] The share e_i is at the beginning of the period before the separation stage. Of all employed workers in state (x,s), only a fraction $1-\pi_{eu}(x,s)$ will not separate and produce at the production stage.

utility in unemployment is $b+h$, where h is the utility value of leisure relative to working (the disutility of working is normalized to zero). Search is random, and all workers receive job offers with the same probability $\lambda(s)$ that only depends on the aggregate state of the economy. We assume that each job offer is associated with an idiosyncratic stochastic utility component ν capturing the personal valuation of workers for jobs. This stochastic non-pecuniary job component comprises, among other things, commuting time, workplace atmosphere, and working schedules of the offered job. It captures in a parsimonious way endogenous search behavior of the unemployed. Unemployed workers optimally follow a reservation utility rule and accept all job offers with ν larger than a state-dependent threshold $\nu^u(b,s)$. We assume ν is independently and identically distributed and is drawn from a logistic distribution G with state-specific mean $\overline{\nu}(b)$ and variance $\sigma_\nu^2 = \pi\frac{\psi_\nu^2}{3}$. The average acceptance probability of an unemployed worker in state (b,s) is $q(b,s) = 1-G(\nu^u(b,s))$, and the transition rate into employment is $\pi_{ue}(b,s) = \lambda(s)q(b,s)$ combining contact rate $\lambda(s)$ and acceptance rate $q(b,s)$. The recursive formulation of the value of an unemployed worker in state (b,s) is

$$\begin{aligned} V_u(b,s) &= b+h+\beta\Bigg(\lambda(s)\int\limits_{\nu^u(b,s)}^{\infty}\Big(\mathbb{E}[V_e(x',s')|b,s]-\nu\Big)dG(\nu) \\ &\quad + \lambda(s)\int\limits_{-\infty}^{\nu^u(b,s)}\mathbb{E}[V_u(b',s')|b,s]dG(\nu) + (1-\lambda(s))\mathbb{E}[V_u(b',s')|b,s]\Bigg) \\ &= b+h+\beta\Bigg(\pi_{ue}(b,s)\mathbb{E}[V_e(x',s')|b,s] + (1-\pi_{ue}(b,s))\mathbb{E}[V_u(b',s')|b,s] \\ &\quad + \lambda(s)\Psi_\nu(q(b,s))\Bigg), \end{aligned} \tag{9}$$

where $V_e(x,s)$ denotes the value of being employed in state (x,s) and the last line again exploits the properties of the logistic distribution with $\Psi_\nu(q) = -q\overline{\nu}(b) - \psi_\nu((1-q)\log(1-q) + q\log(q))$. The state-specific means $\overline{\nu}(b)$ allow us to obtain job-finding rates that are falling with unemployment duration. Such changing utilities capture, for example, decreasing motivation to apply for jobs, more effort to prepare for job interviews, and more effort to be up-to-date with job requirements.

An employed worker who does not separate at the separation stage receives her wage at the production stage. At the end of the production stage, the stochastic skill accumulation takes place. The recursive representation of the value function of employed workers is

$$V_e(x,s) = (1-\pi_{eu}(x,s))\Big(w(x,s)+\beta\mathbb{E}[V_e(x',s')|x,s]\Big) + \pi_{eu}(x,s)\mathbb{E}[V_u(b',s)|x]. \tag{10}$$

Note that in the case of separation, expectations are only over the idiosyncratic benefit state b, as the worker becomes unemployed in the current period. In an abuse of notation, we denote the stochastic benefit level by b'. The benefit level follows the laws of motion for b in equations (5) and (6).

A Cobb-Douglas matching function $m = \varkappa v^{1-\varrho} u^{\varrho}$ determines the number of matches m between vacancies v and unemployed workers $u = u_1 + u_2 + u_3$ during the search stage of each period. The contact rate from a worker's perspective is $\lambda = \frac{m}{u} = \varkappa \theta^{1-\varrho}$ and from a firm's perspective is $\lambda_v = \frac{m}{v} = \varkappa \theta^{-\varrho}$ with labor market tightness $\theta = \frac{v}{u}$. The number of vacancies at the search stage of each period is determined by a free-entry condition

$$\kappa = \lambda_v(s)\beta \sum_{j=1}^{3} q(b_j, s) \frac{u_j}{u} \mathbb{E}[J(x', s')|b_j, s], \tag{11}$$

where κ denotes the per-period cost to post a vacancy. Firms posting vacancies take into account the acceptance rates $q(b_j, s)$ of workers with different unemployment benefit eligibility. Recall that all newly hired workers start with $x' = x_1$ so there is only uncertainty regarding the aggregate state s' for the next period when posting a vacancy. Wages and threshold values for separation decisions $\varepsilon^u(x, s)$, equivalently separation probabilities $\pi_{eu}(x, s)$, are determined by a state-contingent Nash bargaining between the worker and firm over the joint surplus of the match $S(x, s) = J(x, s) + V_e(x, s) - \mathbb{E}[V_u(b', s)] \equiv J(x, s) + \Delta(x, s)$, as in Pissarides (2000, Ch. 2). We denote the bargaining power of the worker by μ. The Nash bargaining problem reads

$$\max_{\{w, \varepsilon^u\}} J(x, s)^{1-\mu} \Delta(x, s)^{\mu}.$$

The first-order condition with respect to wages delivers the standard surplus-sharing rule $\mu J(x, s) = (1 - \mu)\Delta(x, s)$. The first-order condition with respect to the separation cutoff ε^u characterizes the cutoff value in terms of the separation rate $\pi_{eu} = 1 - F(\varepsilon^u)$ as

$$\pi_{eu}(x, s) = \left(1 + \exp\left(\psi_{\varepsilon}^{-1}\left(\exp(a + x) - \bar{\varepsilon} + \tilde{S}(x, s)\right)\right)\right)^{-1}, \tag{12}$$

with $\tilde{S}(x, s) = \beta \mathbb{E}[S(x', s')|x, s] + \beta \mathbb{E}[V_u(b', s')|x, s] - \mathbb{E}[V_u(b', s)|x]$ where $\mathbb{E}[V_u(b', s)|x]$ denotes the expected value from unemployment in the current period taking into account stochastic eligibility (equation (10)). We get that the optimal separation probability $\pi_{eu}(x, s)$ is decreasing in current output $\exp(a + x)$ net of mean costs $\bar{\varepsilon}$ and in an adjusted future match surplus $\tilde{S}(x, s)$ that takes into account the option value from skill accumulation on unemployment benefit eligibility $\beta \mathbb{E}[V_u(b', s')|x, s] - \mathbb{E}[V_u(b', s)|x]$.

Table 5: Calibrated parameters

	Parameter	Value	Description
search and matching	ϱ	0.5	elasticity of the matching function
	$\varkappa$	0.164	efficiency of the matching function
	κ	1.05	vacancy posting costs
	μ	0.5	worker's bargaining power
	γ	0.6	eligibility rate of low-skill workers
	ω	0.010	labor market exit rate
preferences	$\tilde{\beta}$	0.997	time discount factor
	h	0.251	flow leisure utility
	$\bar{\nu}(b_1)$	1.34	
	$\bar{\nu}(b_2)$	1.34	means of non-pecuniary shocks
	$\bar{\nu}(b_3)$	0.10	
	ψ_ν	0.11	dispersion of non-pecuniary shocks
skills and costs	$\bar{\varepsilon}$	0.400	mean of cost shocks
	ψ_ε	0.86	dispersion of cost shocks
	α	0.028	probability of skill accumulation
	Δx	0.098	skill level difference $x_2 - x_1$
	ρ	0.94	autocorrelation of aggregate shock

3.1 Calibration

We calibrate the model to match the pre-reform labor market dynamics of the German labor market. We show all calibrated parameters in Table 5. For the calibration, we take a model period to be one month. We set a first group of parameters outside the model to standard values. The discount factor $\tilde{\beta}$ is set to match an annual interest rate of 4% so that $\tilde{\beta} = 0.997$, and the parameter ϱ of the matching function and the bargaining power of the worker μ are set to $\varrho = \mu = 0.5$.

We describe below how we set the parameters of the unemployment insurance system using independent evidence. Remaining model parameters are set within the model by targeting data moments. Dynamics in the model are only driven by aggregate productivity shocks a. To simulate the model, we linearize the model around its deterministic steady state and use a Kalman filter on GDP growth per capita to determine the time series of aggregate productivity shocks a building on Jung and Kuhn (2014) and Murtin and Robin (2016).[23] We next provide intuitive identification arguments but abstain from

[23] We use GDP per capita for Germany as data on West German GDP are not available at a quarterly frequency.

a formal proof of identification.

Each match produces output with labor and a stochastic cost component, which we interpret as payments to capital. We therefore target the mean of the cost shock $\bar{\varepsilon}$ to a capital share of 40%. Vacancy posting costs κ determine directly how many vacancies are posted and the contact rates in the search market. The contact rate determines the average job-finding rate that we take from the data ($\pi_{ue} = 0.052$). To separately identify matching efficiency $\varkappa$ from vacancy posting costs κ, we use data on the average duration to fill a vacancy from the firm's perspective. In the IAB vacancy survey, the average time to fill a vacancy is 2.2 months. For the UI eligibility parameter γ, we target a share of 60% UI benefit recipients among all inflows to unemployment. The flow utility parameter of leisure h determines the worker surplus from employment Δ, and as part of the total match surplus S, it determines the average probability of separating into unemployment (equation (12)). We match an average separation rate $\pi_{eu} = 0.006$.

Matching the observed volatility of job creation over the business cycle is a challenge for this class of models (Shimer (2005), Hagedorn and Manovskii (2008)). The variation in acceptance rates $q(b, s)$ of workers over the business cycle provides additional amplification to job creation decisions (equation (11)). To impose discipline on the level and variation in acceptance rates, we target the elasticity of average acceptance probabilities with respect to changes in unemployment benefits $\frac{\partial q}{\partial b}\frac{b}{q}$ and target the estimate of 0.53 for Germany from Schmieder and Von Wachter (2016).[24] For a given dispersion of non-pecuniary shocks, this elasticity pins down one of the means of the non-pecuniary shocks. We use it to pin down $\bar{\nu}(b_3)$. We impose the condition that recipients of unemployment assistance benefits b_2 and benefits at a subsistence level b_1 have the same mean of shocks $\bar{\nu}(b_1) = \bar{\nu}(b_2)$. This condition effectively implies different mean utilities for the short- and long-term unemployed. Hence, duration dependence in job-finding rates is informative about the difference between $\bar{\nu}(b_1)$ and $\bar{\nu}(b_3)$. For the duration dependence, we use a difference in job-finding rates between 6 and 12 months of 25%.[25] Very related is the identification of the parameter ψ_ν determining the dispersion of the non-pecuniary shock distribution. While we use the cross-sectional variation in job-finding rates to determine means of the non-pecuniary shock distribution, we leverage the time series variation in job-finding rates to identify ψ_ν. We target a volatility of job-finding rates that corre-

[24] This elasticity of search $\frac{\partial q}{\partial b}\frac{b}{q}$ in the model is the percentage change in the acceptance probability of an unemployed worker receiving unemployment benefits with respect to a percentage change in the benefit level for given contact and separation rates.

[25] Mean job-finding rates of these two benefit groups are computed from aggregate data between 1996 and 2004 on average durations in the respective group. We assume constant job-finding rates within each benefit type. To obtain the job-finding rate of short-term benefit recipients, we further assume that they transit to long-term benefits after 12 months. We can then back out the implied job-finding rate from the mean duration of the truncated distribution.

sponds to 6.4 times the volatility of output. Similarly, we use the time series volatility of separation rates to identify the dispersion of cost shocks ψ_ε. We target a volatility of separation rates that corresponds to 7.8 times the volatility of output. The volatility of separations is higher than the volatility of job-finding rates, in line with existing evidence (Jung and Kuhn (2014), Elsby et al. (2013)).

These elasticities are key when we change the unemployment insurance system. To see this, recall that a 1% change in the surplus of the match from a change in productivity works similarly to a 1% change in the surplus from a change in the outside option. Hence, the time series variation of transition rates is informative about the effects from structural changes in labor market institutions (Costain and Reiter (2008a)).

For the skill process, we use the one-to-one relation between the average duration of short-term employment that we set to 3 years and the probability of skill accumulation α. Similarly, we use the one-to-one relation between the share of long-term employed workers and the probability of labor market exit ω. Short-term and long-term employed workers differ in their productivity levels x_1 and x_2. We exploit the documented separation rate differences between the two groups to pin down the skill difference $\Delta x = x_2 - x_1$. We normalize x_1 and use the difference between the short-term employed workers' separation rate of 0.014 and the long-term employed workers' separation rate of 0.003 from Table 3 to determine the skill difference Δx.

Table 6: Parameters of the unemployment insurance system

pre-reform		post-reform	
b_1	0.251	b_1	0.251
b_2	0.339	b_2	**0.251**
b_3	0.376	b_3	0.372
δ_2	0.028	δ_2	0.028
δ_3	0.062	δ_3	**0.072**

We calibrate parameters of the unemployment insurance system to independent evidence on replacement rates from the OECD and benefit duration from Figure 4. Parameters for the period before and after the reform are shown in Table 6. According to the OECD, a single worker with the average wage before 2004 received unemployment insurance benefits corresponding to 60% of the previous wage during the first year of unemployment and 54% of the previous wage for the following four years. We use these replacement rates to pin down b_3 and b_2. We set δ_3 for the duration of UI benefits to 16.2 months, in line with

the average duration in Figure 4 when using the underlying employment distribution for the pre-reform period. We set δ_2 to match an average duration of receiving unemployment assistance of 36 months. For the subsistence level b_1, we match the average ratio of subsistence benefits to unemployment benefits over the period 1996 to 2002 based on data from the German Statistical Office (earlier data not available). The average ratio corresponds to $\frac{b_1}{b_3}$ in the model, and we fit it to be 67% as in the data ($\frac{b_1}{b_3} = 0.67$).

When exploring the effects of changes in the UI system from the Hartz reforms on labor market dynamics, we abolish long-term unemployment benefits (unemployment assistance benefits) and cut the maximum benefit duration for long-term employed workers. As in Krause and Uhlig (2012), we implement the first part of the reform by setting long-term unemployment benefits b_2 to the level of subsistence social security benefits b_1 (i.e., we set $b_1 = b_2$). The duration parameter δ_2 becomes irrelevant because transitions happen between states with the same benefit levels, and mean utility shocks $\bar{\nu}(b_1)$ and $\bar{\nu}(b_2)$ are set identical across the two states in the calibration. For the change in maximum benefit duration, we decrease the expected benefit duration of UI benefits b_3 from 16.2 months to 13.9 months by increasing the probability that they expire δ_3 (column "post-reform" in Table 6). We obtain the post-reform duration again by averaging the weighted maximum benefit duration after the reform from Figure 4.

In the model, these changes become effective in January 2006. As described above, the Hartz reforms became effective in January 2005, but the law scheduled the new benefit rules to affect workers only if they became unemployed after February 2006. In addition, a wide range of grandfathering rules and hardship clauses were provided with the law, such that it became only slowly applicable to all workers. We implement the complex and detailed legislation by gradually increasing the impact of the reform on labor market dynamics. Specifically, we use different policy functions based on a linear approximation of the steady-state systems before and after the Hartz reforms. We assume a linear weighting scheme that spreads the implementation over four years so that the reform is fully effective in January 2010.[26] When implementing the Hartz reforms in the model, we keep *all* other parameters except for the UI system constant over time.

4 Results

In the first step, we demonstrate the model's ability to match the dynamics of observed labor market flows over time. Dynamics in the model are driven by two sources: aggregate

[26] We also tried implementing the reform directly, with the only difference that the dynamics during the transition period are matched less closely. Obviously, this assumption does not affect changes in steady states but only the behavior of the model during the transition phase. Hence, our key results do not depend on the specific implementation of the transition period.

productivity fluctuations and the structural change of the UI system after the Hartz reforms. As described before, parameters are only calibrated to match selected means and volatilities of labor market flow rates before the Hartz reforms, and the Hartz reforms constitute a parsimonious change in the parameters of the UI system. Figure 9 shows simulated times series of separation and job-finding rates from the model together with the data counterparts of these series. We index all series to the pre-reform steady state that we match as part of the calibration.

Figure 9: Fit for average labor market mobility (1993-2014)

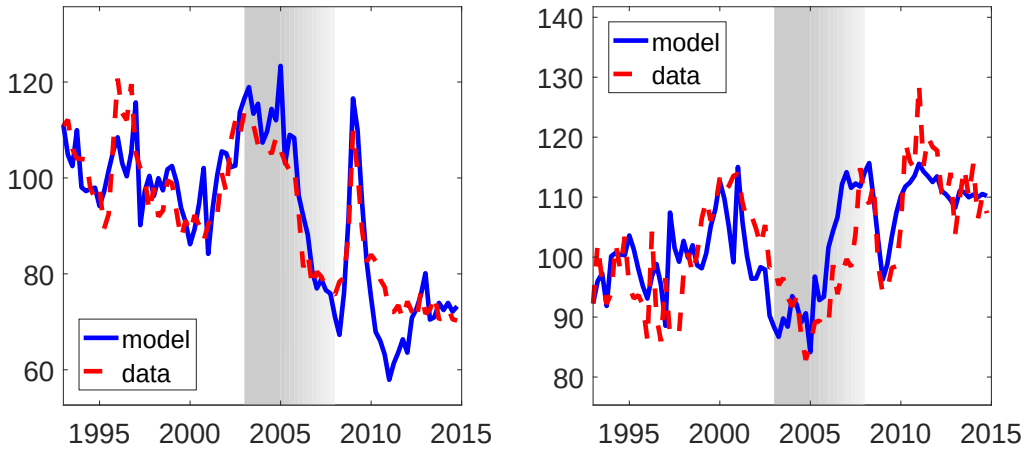

(a) separation rates (b) job finding rates

Notes: Model fit 1993-2014. The solid blue lines mark the model prediction, and the dashed red lines mark the respective flow rate in the SIAB microdata. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

Figure 9(a) shows the model's close fit to the separation rate from the data. The empirical and simulated time series largely lie on top of each other. This is true both before the reform and after the reform. Except for a short period around 2010, the model matches the dynamics of the separation rate closely, notably, also during the financial crisis of 2008.

Figure 9(b) shows the simulated job-finding rates together with the data counterparts. Job-finding rates before 2005 are again matched very closely. After the reform, the model closely matches the dynamics and level changes, with the exception of a period between 2005 and 2009 when the model predicts a more immediate increase in job-finding rates compared to the data. The divergence happens during the transition period after the Hartz reforms when our implementation of grandfathering rules and hardship regulation is very rudimentary. However, what is important is that the changes in average rates between the pre-reform period and the post-reform period are matched almost exactly

by the model outside the transition period.

Our empirical analysis emphasizes the large heterogeneity of changes in separation rates after the reform. While the heterogeneity in the model remains stylized, we demonstrate in Figure 10 the key dimension of heterogeneity by employment duration and the model's ability to match such heterogeneity in changes in separation rates. As for the average separation rate, levels and level differences between short-term and long-term employed workers before the reform have been calibrated so that they are matched by construction. Heterogeneity in changes after the reform in untargeted and provides a check of the empirically documented relationship against the prediction of economic theory. Results in Figure 10 demonstrate that the model supports that the reform led to the observed changes in labor market dynamics by demonstrating a close match of the heterogeneous responses in separation rates between the model and the data.

Figure 10: Fit for heterogeneity in labor market mobility (1993-2014)

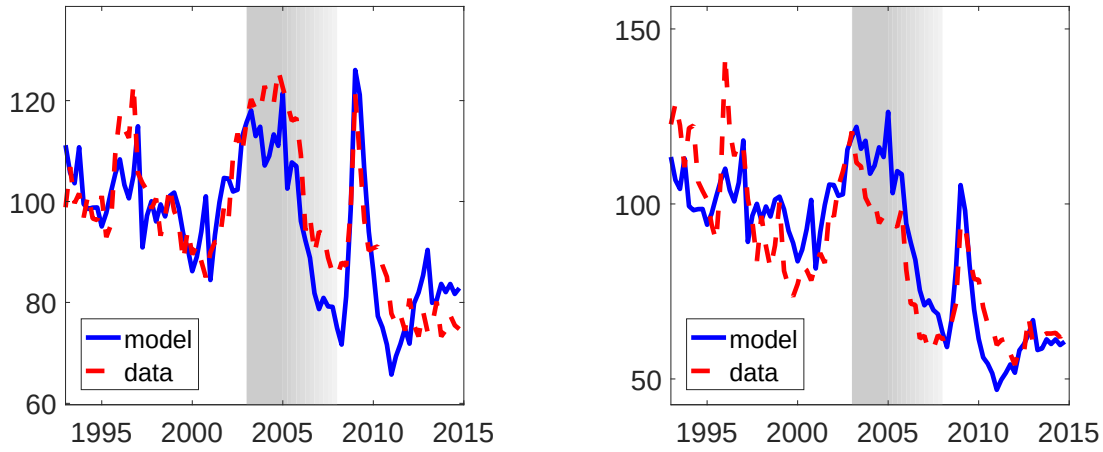

(a) separation rates ($\leq$ 3 years) (b) separation rates ($>$ 3 years)

Notes: Model fit for separation rates with low ($\leq$ 3 years, left panel) and high ($>$ 3 years, right panel) employment duration from 1993 to 2014. The solid blue lines mark the model prediction, and the dashed red lines mark the respective flow rate in the SIAB microdata. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

Figure 10(a) shows the simulated and empirical separation rates for short-term employed workers with employment durations of less than three years. The model matches the time series, including their volatilities, very closely. Unlike for the average separation rates, heterogeneous volatilities of separation rates for short-term and long-term employed workers have not been part of the calibration but are an endogenous prediction of the model. Over the long run, the model predicts a slightly lower decline in separation rates for short-term employed workers relative to the data (10% versus 20%), but importantly,

the model predicts, in line with the data, a substantially smaller decline in separation rates for short-term employed workers relative to the average in Figure 9(a) and relative to long-term employed workers in Figure 10(b).

Figure 10(b) compares the separation rates of long-term employed workers between the model and data. We find that the time series for the long-term employed workers are matched closely in both volatility and trend. We find for the long-term employed workers that the model slightly overstates the decline in separation rates, but importantly again, the separation rates of long-term employed workers decline more than the average separation rates.

Separation rates of long-term employed workers are affected by the abolition of unemployment assistance benefits b_2 but also by the reduction of the maximum benefit duration δ_3. In Table 4, we estimate elasticities between 0.51 and 0.58 of separation rates with respect to changes in maximum benefit duration. In the model, we derive the corresponding elasticity by varying δ_3 at post-reform benefit levels. The implied elasticity of separation rates with respect to changes in benefit duration is 0.62, which is just slightly outside the range of empirical point estimates but well within their confidence intervals. This close alignment of model and data for this untargeted elasticity lends further support for the underlying calibrated elasticities of our quantitative model.

Overall, our parsimonious model of labor market dynamics aligns closely with the key empirical pattern for the changes in separation rates and job-finding rates. The causal mechanism in the model is the reform of the UI system by the Hartz reforms. In the next section, we use our structural model framework to provide counterfactual simulations for labor market dynamics absent the UI reform.

4.1 Counterfactual simulations

Simulating the German labor market absent the Hartz reforms delivers labor market dynamics that are strongly at odds with the data. By contrast, the same model closely matches labor market dynamics in Germany for the two decades from 1994 to 2014 when the Hartz reforms are implemented. This finding provides further support for the Hartz reforms as the main driver of the observed changes in the German labor market dynamics after 2005. Furthermore, the counterfactual simulation provides an approach to determine the contributions of business cycle fluctuations to changes in labor market dynamics as the end of the sample period was in the middle of an ongoing economic expansion.

The construction of the counterfactual simulation in the absence of the Hartz reforms on labor market dynamics is simple and transparent. We keep all model parameters constant over time, including the parameters of the UI system, so that no structural

change takes place. We also keep the aggregate shock series identical and feed in the previously estimated productivity shocks from the Kalman filter. This counterfactual simulation provides time series of separation rates, job-finding rates, and unemployment rates in the absence of the Hartz reforms. Figure 11 shows the counterfactual simulation results for the time period from 1993 to 2014.

Figure 11: Counterfactual model simulation absent Hartz reforms (1993-2014)

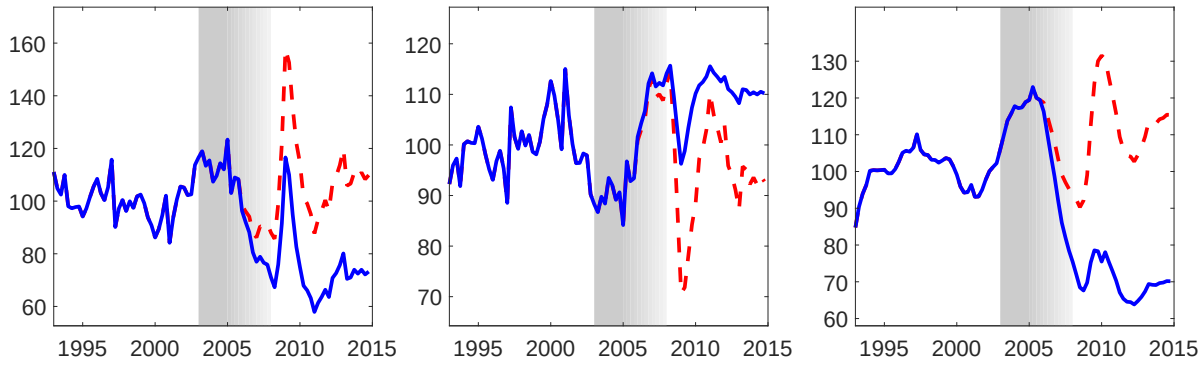

(a) Separation rates (b) Job-finding rates (c) Unemployment rates

Notes: Model simulations with and without the Hartz reforms for the period 1993 to 2014. The solid blue lines show the model with the Hartz reforms, and the dashed red lines show the counterfactual rate without policy change. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

By construction, the time series from the baseline and the counterfactual in the period before the implementation of the Hartz reforms lie exactly on top of each other as we rule out any anticipation effects.[27] After the implementation of the reform, the two simulated time series strongly diverge. Separation rates of the counterfactual remain high and fluctuate around their pre-reform level, as shown by the dashed red line in Figure 11(a). Separation rates of the counterfactual simulation strongly spike during the financial crisis of 2008, to almost 160% of their steady-state level. In the case of the reform, the separation rate still spikes but increases only to slightly more than 120% of the old steady-state level.

Job-finding rates in Figure 11(b) also evolve identically between baseline and counterfactual up to the implementation of the reform, when the two series start to diverge. In the new steady state after the reform, the job-finding rates increase permanently by 10%, whereas by construction, they fluctuate around the old steady-state level in the absence of the reform. Over time, the divergence is strongest during the financial crisis. In the

[27] Anticipation effects are likely small as the implementation of the reform happened on short notice. Parliament approved the law that became effective in January 2005 only in June 2004. See Hochmuth et al. (2019) for additional discussion supporting the assumption of no or very small anticipation effects.

counterfactual scenario without the Hartz reforms, job-finding rates plummet to around 70% of their steady-state level. In the case with the Hartz reforms, the job-finding rate still decreases, but only to a level slightly below its old steady-state level. The divergence of the separation and job-finding rates manifests itself in very different dynamics of the unemployment rate. While unemployment in the baseline simulation with the Hartz reforms declines by 30% relative to the pre-reform steady state, the unemployment rate, by construction, stays put at its pre-reform level absent the reform. A second part of the German labor market miracle was the small increase in unemployment rates during the financial crisis, a fact that our model accounts for. We also find a marked difference in the evolution of unemployment rates between the simulations with and without the Hartz reforms during the financial crisis. The counterfactual simulation shows an increase in the unemployment rate of almost 30% over its long-run average. Such sharply and strongly rising unemployment rates are reminiscent of the typical European country and the United States during these years. In the case of the implementation of the Hartz reforms, the rise in unemployment rates is substantially smaller compared to the largest labor market crisis in decades that most other countries experienced. Unemployment rates increased about 10% over their new steady-state level, which itself is 30% below the pre-reform level. The reason for the modest increase in unemployment after the reform is that while separation rates spike in both simulations, the relative decline in the job-finding rate is much smaller in the case of the Hartz reforms.[28] These strikingly different dynamics based on our theoretical labor market model provide our final argument for the relationship between the Hartz reforms and the German labor market miracle.

4.2 Decomposing cyclical and structural changes

We use the counterfactual model simulations to further quantify the contribution of the business cycle to the decline in unemployment since 2004. Figure 11 shows that in 2004 all data series are away from their respective steady states. Our motivating evidence in Figure 1, like most of the public debate, focuses on 2004 as a year of reference to assess the effect of the Hartz reforms on unemployment rates. Taking 2004 as the reference, the decline in the unemployment rate between 2004 and 2014 contains some part that arises from the business cycle variation and not from a structural change. We rely on the

[28] Germany's reliance on short-time work is oftentimes suggested as an explanation for the low rise in unemployment rates during the Great Recession. Balleer et al. (2016) find that short-term work reduces the increase in unemployment rates by around 20%, so unemployment rates would have gone up by 36% rather than 30% without short-term work. Balleer et al. (2016) also find that the smaller reaction results mostly from lower separation rates, whereas we explain the small reaction by a smaller decline in job-finding rates in comparison with the counterfactual simulation. We abstract from a detailed investigation of short-time work, but we acknowledge that such an investigation is important but still beyond the scope of the current paper.

counterfactual simulation to isolate this business cycle component. Our decomposition approach is straightforward: We attribute all changes in the counterfactual simulation absent the Hartz reforms to the business cycle, and by subtracting these changes from the baseline model, we isolate the structural component of the changes in separation rates, job-finding rates, and unemployment rates. Table 7 shows average unemployment, separation, and job-finding rates in 2004 and 2014 from model simulations with and without the Hartz reforms. The columns labeled *change* show the percentage change in the respective rates between 2004 and 2014. The change in the baseline case with the Hartz reforms compounds the business cycle effects with the effects from the structural reform, whereas the change in the case when the reform is not implemented results only from business cycle variation. We report the relative contribution of the business cycle in the last column of Table 7.

Table 7: Business cycle contribution

	with reform			absent reform			Business cycle
	2004	2014	change	2004	2014	change	contribution
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
unemployment rate	10.38	6.16	-40.6%	10.38	10.13	-2.5%	6.1%
job-finding rate	4.80	5.80	20.8%	4.80	4.90	1.9%	9.2%
separation rate	0.69	0.45	-34.2%	0.69	0.68	-1.3%	3.7%
separation rate (short-term)	1.60	1.21	-24.6%	1.60	1.57	-1.9%	7.6%
separation rate (long-term)	0.33	0.18	-45.8%	0.33	0.32	-2.3%	5.0%

Notes: This table shows the unemployment and flow rates in the model before the reform (2004) and in the most recent year (2014). Columns 1-3 show the rates implied by a model with the Hartz reforms, and columns 4-6 show the rates without the reform but with the same business cycle shocks. The last column shows the relative contribution of the business cycle to the overall change in the respective variable.

First, we consider the baseline case when the reform is implemented. The key driver of the lower unemployment rates is the decline in the separation rate by 34%; the job-finding rate increased by 21%. Comparing these effects to the case absent the reform in the middle columns of Table 7 isolates the business cycle effect and shows that the business

cycle effect is small. The last column shows the constructed business cycle contribution to changes in the unemployment, separation, and job-finding rates; we find that they never exceed 10%. Looking at the unemployment rate, we find that of the 41% decline, only 6% stems from the business cycle, whereas most of the decline after 2004 is structural. Based on these results, we conclude that the business cycle effect is small and of minor importance for the German labor market miracle.

4.3 Germany without the Hartz reforms

The strong deviation of the counterfactual simulation in the absence of the Hartz reforms supports the main hypothesis of this paper that the Hartz reforms are the key driver of the German labor market miracle. An important question is whether the quantitative size of the model-predicted effects is realistic, that is, whether the counterfactual provides a good description of what would have happened had the reforms not been implemented. To estimate an empirical counterfactual for the German unemployment rate, we estimate a synthetic control for Germany (Abadie et al., 2010). The synthetic control estimation constructs a weighted average from a pool of candidate countries. We take a large set of OECD countries for which sufficient data are available. The synthetic control approach determines weights for the different countries to match the pre-reform trend in Germany. We provide further details on the approach in Appendix E. Figure 12(a) shows the close fit between the estimated synthetic control group (dashed blue line) and the German unemployment rate (solid red line) over the pre-reform period (1993-2002). The countries in the synthetic control group did not implement the Hartz reforms and provide an empirical estimate for what would have happened to the German unemployment rate had the Hartz reforms not been implemented. The dashed green line shows our counterfactual model prediction in case the Hartz reforms are not implemented. Comparing the model counterfactual and synthetic control estimates, we find a close comovement between the model prediction and empirical estimate for the evolution of the unemployment rate absent the Hartz reforms. Both the synthetic control and model predict a large rise in unemployment rates during the financial crisis, but the model predicts a slightly stronger recovery. Most importantly, both counterfactuals show a strong divergence from the data after 2005 when the Hartz reforms changed the UI system. The model predicts unemployment rates for 2014 that are 50% higher than what we observe in the data. This estimate is more conservative compared to that of the synthetic control group, which predicts 60% higher unemployment rates in Germany absent the reform. This analysis provides two important conclusions. First, the model is consistent with empirically observed elasticities of the UI reform, and second, unemployment rates today would be at least 50% higher than observed in the data had the Hartz reforms not been

implemented.

Figure 12: Synthetic control, model prediction, and inflow correction

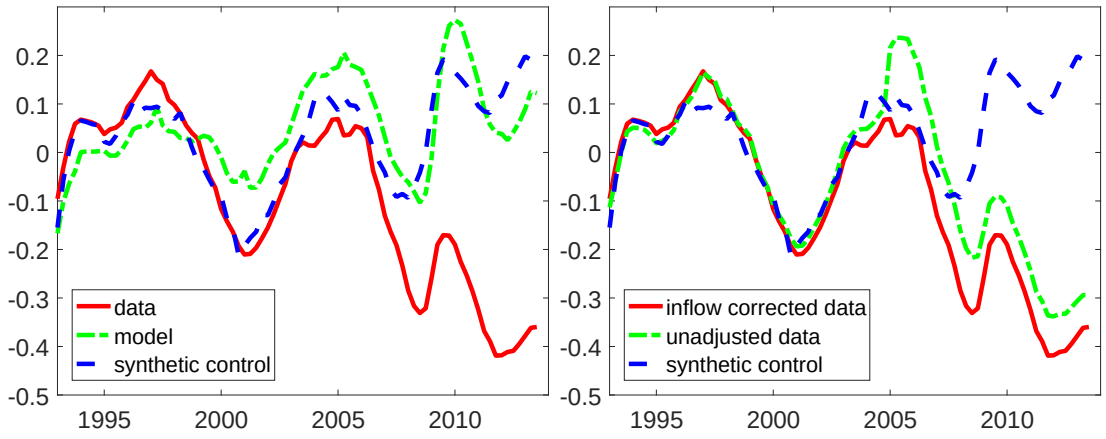

(a) Counterfactual (b) Inflow correction

Notes: The figure shows the model prediction of the unemployment rate, the empirical unemployment rate, and the synthetic control estimate for the unemployment rate. All unemployment rates are expressed as percentage deviations from their pre-reform mean. Left panel: Solid red line shows the (inflow-adjusted) unemployment rate. The dashed-dotted green line shows the model counterfactual for the unemployment rate if the Hartz reform is not implemented. The dashed blue line shows the synthetic control estimate for the unemployment rate. Right panel: The solid red line shows the (inflow-adjusted) unemployment rate, the dashed-dotted green line shows the unadjusted unemployment rate, and dashed blue line shows the synthetic control estimate for the unemployment rate.

The synthetic control estimate also provides the opportunity for a validity check of our inflow correction. The synthetic control estimate uses only pre-reform data until 2003, and it provides an estimate for the evolution of the German unemployment rates absent the Hartz reform and other structural changes such as the changes in eligibility rules in 2005 that we correct for using our inflow correction approach. One way to assess the validity of our inflow correction is to compare the synthetic control estimate for the unemployment rate to the inflow-adjusted unemployment rate. Both are supposed to represent the unemployment rate absent changes in eligibility rules. While this is not a formal statistical test, it still provides a hint as to whether the inflow correction is a reasonable estimate of the unemployment rates at the beginning of 2005. Figure 12(b) compares the inflow corrected unemployment rate (solid red line) to the unadjusted unemployment rate (dashed-dotted green line) and the synthetic control estimate (dashed blue line). We find that the synthetic control estimate and the inflow-adjusted estimate align closely in January 2005 when the unadjusted unemployment rate spikes after the large inflow resulting from changes in eligibility rules. We take this as supporting evidence for the validity of our suggested inflow correction approach to account for the changes in

eligibility rules in 2005.

4.4 Relative importance of the separation rate channel

We have documented empirical evidence and economic theory that falling separation rates are the key driver of declining unemployment rates after the Hartz reforms in Germany. By contrast, most existing labor market research has focused on exploring changes in job-finding rates after changes in the UI system; see, for example, the recent work by Hagedorn et al. (2013) and Chodorow-Reich and Karabarbounis (2019). At first glance, these results suggest a tension between our findings and the focus of existing research. What we will show is that the divergence of our results from the existing focus of the literature not only is consistent with economic theory but also ought to be expected. For our discussion, we focus on a stylized static version of the model from Section 3 that allows us to derive the reaction of the separation rate to a UI reform in closed form.

The economic environment is as follows. All workers are employed at the beginning of the period. Each worker-firm match has stochastic output y that is composed of an aggregate (deterministic) productivity component A and an idiosyncratic cost shock ε, so that output is $y = A - \varepsilon$. Idiosyncratic cost shocks ε have a distribution function $F(\varepsilon)$ and density function $f(\varepsilon)$. Workers and firms make separation decisions at the beginning of the period after having received their idiosyncratic cost shock. If they stay matched, the worker receives the bargained wage w. If they separate, the worker becomes unemployed. Unemployed workers receive UI benefits b for a fraction of time $1 - \pi_{ue}$, and for the fraction of time π_{ue}, they work in a job with average productivity A and $\varepsilon = 0$. The resulting value of being unemployed V_u is $V_u = \pi_{ue}A + (1 - \pi_{ue})b$, the value of employment is $V_e = w$, and the value of a filled job to the firm is $J = y - w$. The worker surplus is $\Delta = w - V_u$, and the total surplus is $S = J + \Delta = y - V_u$. Nash bargaining over wages and separation decisions delivers $w = \mu y + (1 - \mu)V_u$ with μ denoting the worker's bargaining power. The bargained separation cutoff for cost shocks ε^u is $(1 - \pi_{ue})(A - b)$ so that separation decisions are individually efficient and occur if $S < 0$. The implied (ex ante) separation probability (separation rate) is $\pi_{eu} = Prob(\varepsilon > \varepsilon^u) = 1 - F((1 - \pi_{ue})(A - b))$. Using this result, the elasticity of separations with respect to a change in UI generosity (b) is

$$\frac{\partial \pi_{eu}}{\partial b}\frac{b}{\pi_{eu}} = \frac{f(\varepsilon^u)}{1 - F(\varepsilon^u)}(1 - \pi_{ue})b. \tag{13}$$

The elasticity depends negatively on the job-finding rate π_{ue} and positively on the generosity of benefits b.[29] A high separation rate elasticity is therefore a direct consequence of long average unemployment durations $1 - \pi_{ue}$. Such long unemployment durations are

[29] If shocks are distributed logistically with mean zero and scale parameter ψ, the elasticity simplifies

a characteristic of the German labor market and many European labor markets. If unemployment durations are short as in the United States (π_{ue} is high), the formula implies a low separation rate elasticity and justifies a focus on job-finding rates. Intuitively, separation decisions do not react strongly to changes in the UI system if unemployment is very transient. This effect is further amplified if UI generosity (b) is higher in Germany than in the United States. Generally, if UI benefits are high or low, this will affect currently employed workers in a country such as the United States only marginally as they expect to be unemployed for only a short period of time. If unemployment is persistent, a cut in benefits has a large, long-lasting effect on workers' welfare, leading to a strong reaction of separation rates after UI reforms. This intuition easily reconciles what at first glance appeared to be a tension in the cross-country analysis. It also corroborates the focus on job-finding rates when analyzing the United States, and it provides the argument for a focus on the separation rate response in Germany.

The elasticity formula in equation (13) also connects our results to insights in Costain and Reiter (2008b). The elasticity can be reformulated as an elasticity with respect to changes in aggregate productivity A (business cycle shocks):

$$\frac{\partial \pi_{eu}}{\partial b}\frac{b}{\pi_{eu}} = - \underbrace{\frac{\partial \pi_{eu}}{\partial A}\frac{A}{\pi_{eu}}}_{\text{business cycle elasticity}} \frac{b}{A}. \tag{14}$$

This result highlights the tight connection between the business cycle volatility of separation rates and their reaction to changes in the UI system. We exploit this relationship for our model calibration. Using data from 1980 to 2004, Jung and Kuhn (2014) document that the business cycle volatility of separation rates is three times higher in Germany than in the United States. Adding further data from Elsby et al. (2013), they show that such higher volatility in separation rates is a common feature across European labor markets and correlates strongly with lower job-finding rates (see their Figure 2). Indeed, Jung and Kuhn (2014) document that the United States has the lowest separation rate volatility across all OECD countries, scrutinizing the transferability of results on the consequences of UI reforms on labor market dynamics from the United States to the typical European country.

---

to

$$\frac{\partial \pi_{eu}}{\partial b}\frac{b}{\pi_{eu}} = (1-\pi_{ue})\frac{(1-\pi_{eu})b}{\psi}.$$

4.5 Wages and job stability

One prediction of economic theory is that a less generous UI system leads to lower wages. In this case, declining wages and increasing profits of firms lead to more vacancy postings, and, as a consequence, job-finding rates will increase so that unemployment rates will decline. This mechanism plays the key role in Hagedorn et al. (2013) and Krause and Uhlig (2012). While this mechanism is also at work in our model, we find it to be of minor importance quantitatively. Consistently, Jäger et al. (2020) find very small wage changes after a change in the UI system in Austria. In our calibrated model, wages of long-term employed workers fall by only 1% in the new steady state. For short-term employed workers, our model delivers a wage increase of 0.8%. Wages increase because the treatment of short-term employed workers by the reform was smaller (Section 2.4), and increasing job-finding rates improve their outside option. These small wage changes suggest that the effect of lower wages on vacancy creation is of minor importance in accounting for changes in labor market dynamics after the Hartz reforms. This is further corroborated by the small overall increase in job-finding rates after the reform. Our interpretation of the structural change hinges on a second theoretical prediction. If the UI system becomes less generous and separation decisions are part of the bargaining process between the worker and the firm, then workers want to trade off job stability in the form of lower separation rates against lower wages (Jung and Kuhn (2018)). If separation rates are low, as is the case in the German labor market, then a small cut in wages can have large relative effects on separation rates. The ability to adjust wages and separation decisions jointly in the bargaining process implies that we should observe in the data a negatively sloped locus of bargained separation rates and wages across productivity levels. Put simply, high-wage workers are in more stable jobs. Jung and Kuhn (2018) provide evidence for such a wage-job stability relationship in U.S. data. We adopt their approach to provide evidence for such a relationship in the German data. Specifically, we regress the probability of separating into unemployment for individual i over the next six months $\pi^6_{eu,i}$ on the contemporaneous (log-)wage $\log(w_{i,0})$ controlling for worker observables X_i,

$$\pi^6_{eu,i} = \alpha + \beta \log(w_{i,0}) + \gamma X_i, \tag{15}$$

where the vector X_i contains dummies for gender, 10-year age brackets, education levels, and time and industry dummies; $\log(w_{i,t})$ refers to average daily earnings.[30] The dependent variable $\pi^6_{eu,i}$ is a binary variable that is equal to one if the worker separates into unemployment at least once over the next six months. The coefficient β is the key

[30] We focus only on full-time workers for this regression.

coefficient of interest as it determines how strongly unexplained wage variation is related to separation rates. To support the theory, we expect a negative coefficient. When we divide the coefficient β by the average separation rate $\bar{\pi}^6_{eu}$, it corresponds to the elasticity of separation rates with respect to wages.

Table 8: Wages and separation rates

period	$\bar{\pi}^6_{eu}$	β	elasticity
pre-reform	0.020	-0.020***	-1.0
post-reform	0.013	-0.017***	-1.3

Notes: Regression results for the relationship between wages and separation rates before and after the Hartz reforms. The column labeled $\bar{\pi}^6_{eu}$ shows the average 6-month separation rate. The column labeled β shows the regression coefficient from equation (15). The last column reports the implied elasticity of separation rates on wages. We indicate by *, **, *** significance at the 10%, 5%, and 1% level. See text for further details.

We estimate the regression coefficient β separately for the period before the Hartz reforms (1993-2002) and after the reforms (2008-2014). Table 8 reports the coefficient estimates. The last column reports the implied elasticity of separation rates with respect to wages. For both time periods, the regression coefficient is negative and highly statistically significant. The key model prediction of a negatively sloped wage-job stability locus is therefore not rejected by the data. Before the reform, we find an elasticity of -1.0 so that a (residual) productivity increase associated with a 1% wage increase reduces the separation rate by 1%. After the reform, this elasticity increases by almost one-third to -1.3. This means that after the reform, the trade-off shifted toward job stability. Before the reform, workers were indifferent between a 1% wage increase and a 1% lower separation probability, whereas after the reform, workers are indifferent between a 0.77% wage increase ($= 1/1.3$) and a 1% lower separation rate. In other words, job stability became more important as workers are, after the reform, willing to accept a 0.23% smaller wage increase in order to obtain a 1% increase in job stability.[31]

4.6 Welfare effects

Our empirical and theoretical analysis demonstrates that changes in separation rates have been the driver of the German labor market miracle starting in the mid-2000s.

[31] We also ran a logit regression that directly estimates the elasticities. The estimated elasticities in terms of level and change are similar to the case of the linear regression. We estimate an elasticity of -0.87 before the reform and -1.06 after the reform.

We document and explain why the decline in separation rates has not been uniform in the population and that long-term employed, high-wage workers saw the strongest decline in their separation rates in reaction to the reform. Job-finding rates increased and thereby increased the probability that both short- and long-term unemployed can find jobs and enter into employment more quickly. Our structural model allows us to investigate the welfare consequences of these changes for the different worker groups. We derive welfare consequences as the consumption-equivalent variation in steady-state consumption for a worker (i.e., we quantify a worker's willingness to pay to avoid the reform). We compute welfare consequences by relying on a steady-state comparison for all worker types: short- and long-term employed workers and workers in each of the three tiers of the unemployment insurance system.[32] Note that this equivalent variation is uncompensated in the sense that because of lower unemployment after the reform, the government could redistribute gains from the reform. Our equivalent variation is before any redistribution and indicates the compensation necessary to make workers of each group indifferent between implementing the reform and not implementing it.

Table 9: Welfare effects from the unemployment insurance reform

	employed		unemployed		
worker group	short-term employed	long-term employed	social assistance	unemployment assistance	unemployment benefits
equivalent variation	0.11%	0.64%	0.03%	2.11%	1.18%

Notes: Welfare effects of the reform expressed as consumption-equivalent variation for avoiding the implementation of the unemployment insurance reform.

Table 9 shows the welfare effects for the different groups of workers. We find the largest welfare losses for former recipients of unemployment assistance benefits, with a consumption-equivalent variation larger than 2%. Unemployment assistance benefits were abolished by the reform and such a large welfare loss likely provides an explanation for the existence of the grandfathering and hardship regulation that accompanied the reform. Note that here we compare steady states so that, even in our model with the staggered implementation, the welfare effects including the transition would be lower. The group with the second-largest welfare losses has been the unemployed, with an equivalent variation of 1.2%. Unemployed workers receiving social assistance benefits experience hardly any welfare effect because their benefits remain unchanged by the reform. The

[32]The assumption of risk neutrality leads to simple formulas for the consumption-equivalent variation. Denoting the value function before the reform by V_0 and after the reform by V_1, the consumption-equivalent variation is $\Delta = \frac{V_0 - V_1}{V_1}$.

non-zero effect results from an indirect effect from lower wages after skill accumulation in the case of reemployment. Turning to the employed, we find much larger effects for the long-term employed compared to short-term employed workers. The group of long-term employed workers with low separation rates experiences a welfare loss corresponding to a consumption-equivalent variation of 0.6%. This group corresponds to more than 60% of all employed workers in the German labor market. Their low separation rates might suggest that this group is the least affected by the reform, yet we find large welfare losses for them. The reason is highly intuitive and closely connected to the causal mechanism of this paper. Welfare effects are large because the outside option for these workers deteriorates most strongly with the abolition of long-term, wage-dependent unemployment assistance benefits. Hence, a group of almost two-thirds of the German labor market experienced large welfare losses from the reform. These losses remained largely uncompensated in the aftermath of the reform and might therefore explain the discontent with the reform by large parts of the German population.

These results might have important implications beyond the specific case of the German Hartz reforms for reform proposals in other European countries. The results suggest that the political feasibility of UI reforms might critically depend on the compensation of the large group of long-term employed workers with secure jobs who, at first glance, might appear very detached from the topic of unemployment benefit reforms. Yet, we show that they are at the center of the adjustment process in countries that on average feature long periods of unemployment after a job loss.

4.7 Alternative explanations

We provide extensive empirical evidence in connection with economic theory to argue that the cause of the German labor market miracle has been the unemployment insurance reform that was part of the Hartz reforms in the mid-2000s. The German labor market miracle (Burda and Seele (2016)) has recently been widely studied, and various narratives have been proposed in addition to the ones that highlight changes in job-finding rates as a key driver. We provide a short summary of our investigation regarding such alternative explanations and relegate details to Appendix F. The most prominent narrative may be from Dustmann et al. (2014), who argue that Germany's unit labor costs and wages were declining relative to other European countries even before the Hartz reforms. They point to declining union power as a possible source for the decline. From the viewpoint of economic theory, wage trends alone are hard to interpret and need to be discussed relative to productivity trends and trends in the outside option. We show in Appendix F.1 that the declining trend in unit labor costs in Germany is hard to reconcile with the relative evolution of the unemployment rate in Germany relative to other European countries.

In particular, this explanation struggles to account for the increase in unemployment rates in Germany during the 1990s and the sudden reversal after 2005. Another related narrative focuses on globalization and an export-demand-driven boom in Germany (see Dauth et al. (2016) for some evidence on globalization effects in export- and import-exposed industries). Looking at industries by export exposure following the industry classification used by Dauth et al. (2016), we show in Appendix F.2 that separation rate changes in industries classified as export-exposed behave similarly to the ones classified as non-exposed, suggesting that export exposure is likely not the main driver of the decline in separation rates. Generally, explanations that, through different channels, affect aggregate GDP growth will already be captured by our analysis as we include aggregate GDP changes in our analysis. Section 4.2 provides an upper bound of the contribution of these effects. Finally, we study whether the Hartz reforms have affected long-term unemployed workers in particular (see Klinger and Rothe (2012) for a more extensive empirical analysis). Reducing long-term unemployment was one of the explicit goals of the reform. We show in Appendix F.3 that the share of long-term unemployed remained largely constant between the pre- and post-reform periods. Together with the evidence on job-finding rates, this finding suggests that the effects via a reduction in long-term unemployment are likely to be very modest.

5 Conclusions

A key question in labor market research is how UI reforms affect unemployment rates and labor market dynamics. We revisit this old question by exploring the German labor market miracle that followed after the German Hartz reforms, one of the largest UI reforms in industrialized countries in recent decades. By combining an empirical analysis of worker flows with economic theory on labor market dynamics, we trace the German labor market miracle back to the Hartz reforms' overhaul of the UI system. Our analysis highlights changes in separation rates after the reform as the quantitatively important channel through which the UI reform brought down unemployment rates.

Specifically, we provide evidence that a decrease in separation rates after the reform accounts for 76% of declining unemployment. The reduction in separation rates is heterogeneous, with long-term employed, high-wage workers being most affected. We exploit this heterogeneity in combination with differences in treatment intensities by the reform to establish a causal link from the reform to changes in separation rates. We use economic theory to support this empirical relationship qualitatively and quantitatively. Using the calibrated labor market model, we provide a counterfactual simulation of the German labor market absent the reform and find that unemployment rates would be 50% higher one decade after the reform. We derive theoretically that such a strong response of separation

rates after the UI reform ought to be expected because of long average unemployment durations in Germany that are also characteristic of most European labor markets and in contrast to the U.S. labor market where unemployment is typically transient. Exploring the welfare consequences of the Hartz reforms, we find that long-term employed, high-wage workers suffered substantial welfare losses in the absence of compensating transfers. This worker group accounts for almost two-thirds of Germany's workforce and has the lowest probability of becoming unemployed. In hindsight, their welfare losses might explain the discontent of a large part of the German population with the reforms despite strongly falling unemployment rates.

Our results have at least two implications for labor market reforms. The first is related to future labor market reforms in other European countries that have been discussed after observing Germany's labor market miracle. For these reforms, our welfare analysis suggests that appropriate compensation schemes must be designed to avoid discontent in the large parts of the population that at first glance seem very detached from the consequence of any UI reform because of low unemployment risk. Second, the strong reaction of separation rates after changes in non-work benefits highlights the importance of this channel for other labor market reforms such as early retirement programs or disability insurance programs.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American statistical Association*, 2010, *105* (490), 493–505.

Antoni, Manfred, Andreas Ganzer, and Philipp vom Berge, "Sample of Integrated Labour Market Biographies (SIAB) 1975-2014," Technical Report, Data report 04/2016, Institute for Labor Market Research (IAB), Nuremberg 2016.

Autor, David, David Dorn, and Gordon Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103*, 2121–68.

Baily, Martin, "Some aspects of optimal unemployment insurance," *Journal of Public Economics*, 1978, *10*, 379–402.

Balleer, Almut, Britta Gehrke, Wolfgang Lechthaler, and Christian Merkl, "Does short-time work save jobs? A business cycle analysis," *European Economic Review*, 2016, *84*, 99 – 122. European Labor Market Issues.

Born, Benjamin, Gernot Müller, Moritz Schularick, and Petr Sedlacek, "The Costs of Economic Nationalism: Evidence from the Brexit Experiment," *The Economic Journal*, 2019, *129* (623), 2722–2744.

Burda, Michael and Stefanie Seele, "No Role for the Hartz Reforms? Demand and Supply Factors in the German Labor Market 1993-2014," *SFB 649 Discussion Paper*, 2016.

Carrillo-Tudela, Carlos, Andrey Launov, and Jean-Marc Robin, "The fall in German unemployment: A flow analysis," 2018.

Chetty, Raj, "A general formula for the optimal level of social insurance," *Journal of Public Economics*, 2006, *90*, 1879–1901.

Chodorow-Reich, Gabriel and Loukas Karabarbounis, "The macro effects of unemployment benefit extensions: a measurement error approach," *The Quarterly Journal of Economics*, 2019, *134* (1), 227–279.

Costain, James and Michael Reiter, "Business Cycles, Unemployment Insurance, and the Calibration of Matching Models," *Journal of Economic Dynamics and Control*, 2008, *32* (4), 1120 – 1155.

Costain, James S and Michael Reiter, "Business cycles, unemployment insurance, and the calibration of matching models," *Journal of Economic Dynamics and control*, 2008, *32* (4), 1120–1155.

Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum, "Adjusting to Globalization - Evidence from Worker-Establishment Matches in Germany," *mimeo, Duesseldorf Institute for Competition Economics*, 2016.

Dustmann, Christian and Costas Meghir, "Wages, Experience and Seniority," *Review of Economic Studies*, 2005, *72* (1), 77–108.

_ , Bernd Fitzenberger, Uta Schönberg, and Alexandra Spitz-Oener, "From sick man of Europe to economic superstar: Germany's resurgent economy," *The Journal of Economic Perspectives*, 2014, *28* (1), 167–188.

Elsby, Michael WL, Bart Hobijn, and Ayşegül Şahin, "Unemployment dynamics in the OECD," *Review of Economics and Statistics*, 2013, *95* (2), 530–548.

Fredriksson, Peter and Martin Söderström, "The equilibrium impact of unemployment insurance on unemployment: Evidence from a non-linear policy rule," *Journal of Public Economics*, 2020, *187*, 104199.

Hagedorn, Marcus and Iourii Manovskii, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, 2008, *98* (4), 1692–1706.

_ , Fatih Karahan, Iourii Manovskii, and Kurt Mitman, "Unemployment benefits and unemployment in the great recession: the role of macro effects," Technical Report, National Bureau of Economic Research 2013.

_ , Iourii Manovskii, and Kurt Mitman, "The impact of unemployment benefit extensions on employment: the 2014 employment miracle?," Technical Report, National Bureau of Economic Research 2015.

Hartung, Benjamin, Philip Jung, and Moritz Kuhn, "Etiopathology of Europe's Sick Man: Worker Flows in Germany, 1959-2016," 2016, (10341).

Hertweck, M. and O. Sigrist, "The Ins and Outs of German Unemployment: A Transatlantic Perspective," *Oxford Economic Papers*, 2015, *67* (4), 1078–1095.

Hochmuth, Brigitte, Britta Kohlbrecher, Christian Merkl, and Hermann Gartner, "Hartz IV and the decline of German unemployment: A macroeconomic evaluation," 2019.

Hopenhayn, Hugo and Juan Pablo Nicolini, "Optimal Unemployment Insurance," *Journal of Political Economy*, 1997, *105* (2).

Jäger, Simon, Benjamin Schoefer, and Josef Zweimüller, "Marginal jobs and job surplus: Evidence from separations and unemployment insurance," Technical Report, Working Paper 2018.

_ , _ , Samuel Young, and Josef Zweimüller, "Wages and the Value of Nonemployment," *The Quarterly Journal of Economics*, 2020, *135* (4), 1905–1963.

Jung, Philip and Moritz Kuhn, "Labour market institutions and worker flows: comparing Germany and the US," *The Economic Journal*, 2014, *124* (581), 1317–1342.

_ and _ , "Earnings Losses and Labor Mobility Over the Life Cycle," *The Journal of the European Economic Association*, 2018.

Katz, Lawrence F and Bruce D Meyer, "The impact of the potential duration of unemployment benefits on the duration of unemployment," *Journal of Public Economics*, 1990, *41* (1), 45–72.

Klinger, Sabine and Enzo Weber, "Decomposing Beveridge Curve Dynamics by Correlated Unobserved Components," *Oxford Bulletin of Economics and Statistics*, 2016, *78* (6), 877–894.

_ and Thomas Rothe, "Der Rueckgang der Langzeitarbeitslosigkeit in Deutschland: Ein Erfolg der Hartz-Reformen oder konjunktureller Effekt," *Schmollers Jahrbuch*, 2012, *132* (1), 89–121.

Krause, Michael U and Harald Uhlig, "Transitions in the German labor market: Structure and crisis," *Journal of Monetary Economics*, 2012, *59* (1), 64–79.

Krebs, Tom and Martin Scheffel, "Macroeconomic evaluation of labor market reform in Germany," *IMF Economic Review*, 2013, *61* (4), 664–701.

Launov, Andrey and Klaus Wälde, "Estimating incentive and welfare effects of nonstationary unemployment benefits," *International Economic Review*, 2013, *54* (4), 1159–1198.

Millard, Stephen P. and Dale T. Mortensen, "The unemployment and welfare effects of labour market policy: A comparison of the USA and the UK," in Dennis J. Snower and Guillermo de la Dehesa, eds., *Unemployment Policy: Government Options for the Labour Market*, Cambridge University Press, 1997, pp. 545–572.

Murtin, Fabrice and Jean-Marc Robin, "Labor market reforms and unemployment dynamics," *Labour Economics*, 2016.

Pissarides, Christopher A, *Equilibrium Unemployment Theory*, MIT press, 2000.

Schmieder, Johannes F and Till Von Wachter, "The effects of unemployment insurance benefits: New evidence and interpretation," *Annual Review of Economics*, 2016, *8*, 547–581.

_ , _ , and Stefan Bender, "The effects of extended unemployment insurance over the business cycle: Evidence from regression discontinuity estimates over 20 years," *The Quarterly Journal of Economics*, 2012, *127* (2), 701–752.

Shavell, Steven and Laurence Weiss, "The Optimal Payment of Unemployment Insurance Benefits over Time," *Journal of Political Economy*, 1979, *87* (6), 1347–1362.

Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment, Vacancies, and Wages: Evidence and Theory," *American Economic Review*, 2005, *95* (1), 25–49.

_ , "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 2012, *15* (2), 127 – 148.

Steffen, Johannes, "Sozialpolitische Chronik," *Arbeitslosenversicherung (seit 1969). Arbeitnehmerkammer Bremen. Bremen*, 2008.

Tuit, Sanders and Jan van Ours, "How changes in unemployment benefit duration affect the inflow into unemployment," *Economic Letters*, 2010, *109* (2), 105–107.

A The Hartz reforms

The *Hartz reforms* in Germany consisted of four legislative packages (*Hartz I* to *Hartz IV*) that became effective between 2003 and 2005. The first two parts of the reform were enacted in 2003 and contained several steps. *Hartz I* changed the legal framework for temporary work, making it more attractive for firms to hire temporary workers by lifting restrictions. *Hartz II* changed the regulations for marginal employment and introduced an additional form of social security tax-favored employment (*midi-jobs*) and subsidies for unemployed workers starting their own business.

Hartz III was enacted in 2004 and restructured the federal employment agency. In particular, placement agencies (*Arbeitsämter*) and social security offices (*Sozialämter*) were combined into single institutions (*Arbeitsagenturen*). Newly created *job centers* were set up, and case managers supported the job search of unemployed workers.

Hartz IV was enacted in 2005. This part of the reform constituted the large overhaul of the German UI system that is the focus of our investigation. It is also publicly the most debated and controversial part of the reforms because it substantially reduced unemployment benefits for several groups of workers by abolishing the system of unemployment assistance benefits (*Arbeitslosenhilfe*). Before the reform, unemployment assistance could be received for several years after unemployment benefits expired, depending only on some weak eligibility criteria. Workers who were not eligible for unemployment assistance received a minimum subsistence level (*Sozialhilfe*) that included rent payments but was not linked to previous wages. *Hartz IV* abolished the wage-dependent benefits for the long-term unemployed so that after the reform they would receive the minimum subsistence level (*Arbeitslosengeld II*). Unemployment benefits (*Arbeitslosengeld I*) remained largely unchanged.

The duration of eligibility for unemployment benefits depends on past employment under social security legislation and changed simultaneously with the Hartz reforms. The changes became effective in February 2006. Before the change, workers were eligible for age-specific maximum benefit durations ranging from 12 months for workers younger than 45 years up to 32 months for workers 57 years and older (see Figure 4). The general rule was that two months of employment resulted in one month of benefit eligibility up to the maximum eligibility threshold. Hence, for most workers, two years of employment guaranteed maximum eligibility. After the reform, the maximum benefit duration was set at one year, and three months of employment were necessary for one additional month of eligibility. For workers 55 and older, the maximum duration was cut to 18 months.[33] We

[33] In 2009 this change was partly reversed again. Workers of age 50, 55, and 58 could then receive benefits for up to 15, 18, and 24 months again.

exploit this variation in our empirical analysis. The fact that this change only became effective in 2006 and the fact that additional grandfathering and hardship regulations were introduced motivate our approach of introducing the reform with a transition phase.

To summarize, the Hartz IV reform transformed the former three-tier system of unemployment benefits, unemployment assistance, and subsistence benefits into a two-tier system of unemployment benefits and subsistence benefits.

B Data details

B.1 Sample selection

In our baseline analysis, we focus on the West German labor market from 1993 to 2014 in order to reduce the impact of the German reunification on unemployment and transition rates. We restrict our sample to persons who had employment or unemployment spells only in West Germany. We also drop persons for whom the SIAB does not contain any information on their geographic location or employment status. We provide results for East Germany as part of our sensitivity analysis in Section C.4 of this appendix.

B.2 Definition of labor market states

We define a worker as employed if the worker is full- or part-time employed or employed as an apprentice. We require current wages to be non-zero to exclude dormant employment relationships, for example, workers on maternity leave. We also exclude marginally employed workers in our baseline definition of employment and define them as being unemployed if they have a parallel unemployment spell and as not in the labor force if there is no parallel spell. The SIAB microdata are derived from social security records with information on dependent employment under social security legislation, so we do not cover self-employed workers and public servants (*Beamte*) in our employment definition.

We define a worker as unemployed if the person is registered as unemployed at the federal employment agency.[34] The SIAB microdata provide comprehensive information on unemployment registrations from 2000 onward. For the period 1993 to 2000, we rely on information on benefit-recipient status to define workers as unemployed. This includes all workers who receive unemployment benefits and unemployment assistance. To construct worker flow rates for the entire period 1993 to 2014, we extend the registration-based worker flow rates backward starting in 2000 using the growth rates of benefit-based worker flow rates for the period 1993 to 2000. Extending the time series using growth rates avoids level breaks in the series but preserves the cyclical properties of worker flow rates.

[34] Workers can remain registered as unemployed as long as they work less than 15 hours per week.

Figure 13: Unemployment rates, 2000-2014

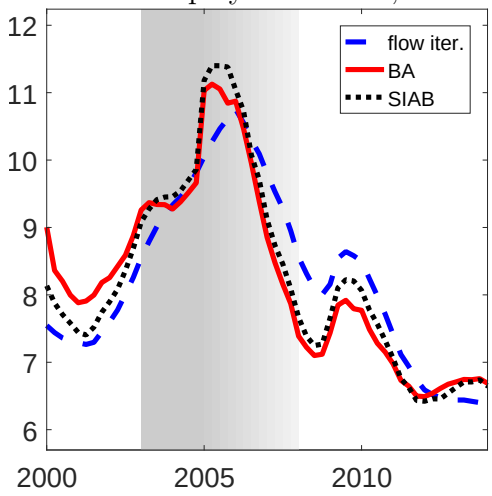

Notes: The solid red line displays the registered unemployment rate in the SIAB data for West Germany, 2000-2014. The dashed blue line displays the unemployment rate obtained by iterating forward the SIAB unemployment rate in January 2000 using monthly separation and job-finding rates: $u_{t+1} = u_t(1 - \pi_{ue,t}) + (1 - u_t)\pi_{eu,t}$. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

In our empirical analysis, we study the evolution of worker flow rates to uncover changes in the underlying dynamics of the inflows and outflows to unemployment. Hence, what is most important for our analysis is that the constructed worker flow rates account for the changes in the unemployment rate over time. Figure 13 shows the unemployment rate from the SIAB microdata (dotted black line) and the unemployment rate from the federal employment agency (solid red line), as in Figure 1. In addition, we construct a flow-based unemployment rate using the law of motion of a two-state approximation of unemployment dynamics,

$$u_{t+1} = u_t(1 - \pi_{ue,t}) + e_t\pi_{eu,t}, \tag{16}$$

where e_t denotes the employment rate of workers covered by social security legislation (see Section 2.2). Such a two-state approximation of unemployment dynamics also underlies our labor market model in Section 3. We use this law of motion to iterate forward the unemployment rate over time. Changes in the unemployment rate using this flow-based approach are only determined by changes in separation rates $\pi_{eu,t}$ and job-finding rates $\pi_{ue,t}$. The unemployment rate from this flow-based approach is shown as the dashed blue line in Figure 13. We find that this unemployment rate closely tracks the dynamics of the aggregate unemployment rate. Hence, changes in the transition rates based on these definitions and construction account for the observed changes in the unemployment rate

over time and are therefore informative about the drivers of declining unemployment.

B.3 Unemployment rates and out of the labor force

For the analysis Section 2 in the main part of the paper, we rely on a two-state representation of labor market dynamics abstracting from flows in and out of the labor force. Here, we demonstrate that the approximation error from the two-state model is small and that the resulting unemployment dynamics align very closely between the two-state approximation and the full three-state model that accounts for flows in and out of the labor force. Recently, Carrillo-Tudela et al. (2018) point out the importance of the flows from out of the labor force to account for changes in the employment stock and its composition (part-time and full-time workers) over time. Figure 14 shows, indexed to the pre-reform period (1993-2003 = 100), the steady-state approximation of the unemployment rate of the two-state model u_t^2 abstracting from flows in and out of the labor force and the three-state steady-state unemployment rate u_t^3,

$$u_t^2 = \frac{\pi_{eu,t}}{\pi_{eu,t} + \pi_{ue,t}} \qquad u_t^3 = \frac{\pi_{eu,t} + \frac{\pi_{nu,t}}{\pi_{nu,t}+\pi_{ne,t}}\pi_{en,t}}{\pi_{eu,t} + \frac{\pi_{nu,t}}{\pi_{nu,t}+\pi_{ne,t}}\pi_{en,t} + \pi_{ue,t} + \frac{\pi_{ne,t}}{\pi_{nu,t}+\pi_{ne,t}}\pi_{un,t}},$$

where e denotes employment, u unemployment, and n out of the labor force and $\pi_{ij,t}$ denotes the respective flow rate from labor market state i to labor market state j in period t.[35]

Figure 14: Steady-state unemployment rates from a two- and three-state model

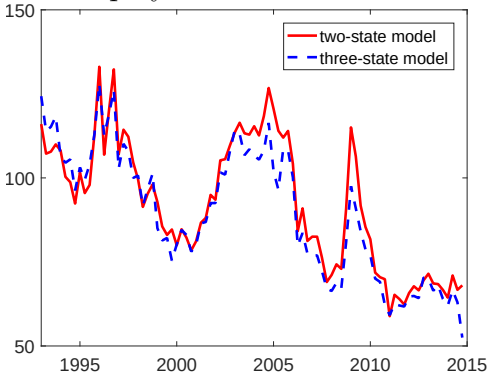

Notes: The solid red line displays the steady-state unemployment rate based on a two-state approximation. The dashed blue line displays the steady-state unemployment rate for the three-state model. Underlying transition rates are quarterly averages of monthly rates.

Figure 14 highlights that the two steady-state unemployment rates track each other

[35]Note that the stock of workers out of the labor force remains unobserved in the SIAB data. The steady state of the unemployment rate in the three-state model can be computed with the level of worker flows as the stock cancels out from the respective ratios.

closely over the entire time period from 1993 to 2014 and that the changes in separation rates π_{eu} and job-finding rates π_{ue} in the two-state approximation alone already track the key dynamics of the unemployment rate over time. As already demonstrated in Figure 13, a two-state stock-flow model also closely matches observed unemployment rates.

C Sensitivity analysis

This section provides a sensitivity analysis of the empirical analysis of Section 2. We consider in Section C.1 a sample in which we do not apply the *inflow correction* described in Section 2.2. Related to skipping the inflow correction, we explore in Section C.2 how much changes in the composition of the employed have contributed to the changes in the separation rates over time. We provide a further detailed discussion of heterogeneity in separation rate changes by age and employment duration in Section C.3. In Section C.4, we compare East and West German worker flow rates. In the main part of the paper, we restrict attention to West Germany. Section C.5 includes marginally employed workers in the definition of employment. In the main part, we do not include marginally employed workers in the definition of the employment state. Section C.6 looks at the effect on job-finding rates from changes in how workers in active labor market programs are counted before and after the reform.

C.1 Worker flows without inflow correction

Figure 15 shows separation and job-finding rates for the baseline sample with the inflow correction and for a sensitivity sample in which we skip the inflow correction.

Figure 15: Separation and job-finding rates (1993-2014)

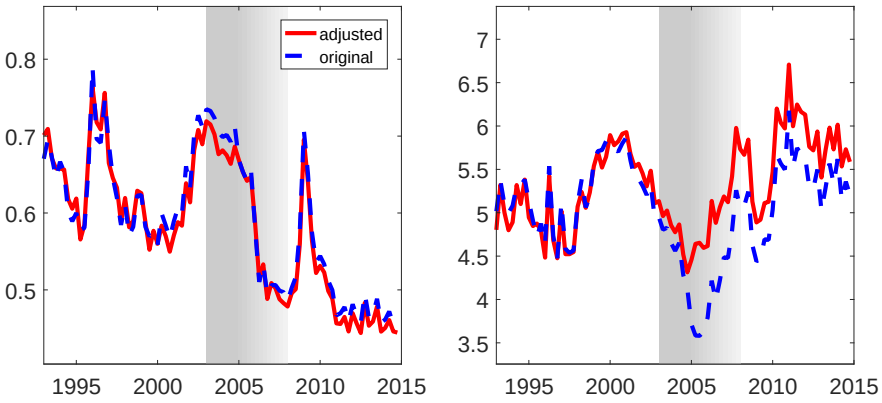

(a) Separation rate (b) Job-finding rate

Notes: the left panel shows separation rates in percentage points. The right panel shows the job-finding rate in percentage points. The solid red lines exclude non-employed entering the unemployment pool in the first half of 2005 who did not become employed until the end of 2006 (inflow correction). The dashed blue lines show the separation rates without inflow correction. The horizontal axis shows the years from 1993 to 2014. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

Looking at separation rates in Figure 15(a), we see that the inflow correction hardly affects separation rates because those workers whom we exclude with our inflow correction are only weakly attached to the labor market. In the case in which they become employed, they constitute only a negligible fraction of total employment so that separation rates remain almost unaffected. This is not true for the job-finding rates in Figure 15(b). Job-finding rates are almost 20% lower in January 2005 in the full sample compared to the inflow-corrected sample. This difference decreases over time but remains sizable even at the end of our sample in 2014. Job-finding rates before 2005 remain largely unaffected, in line with the idea that these workers are only weakly attached to the labor force. Hence, if we do not apply the inflow correction, the increase in job-finding rates would be smaller, and the contribution of the decreasing separation rate to the decrease in the unemployment rate would be even larger.

C.2 Controlling for composition

Our empirical analysis in Section 2 and Section C.3 of this appendix documents large heterogeneity in separation rates across worker groups. One potential reason for decreasing separation rates that would be unrelated to the UI reform could be changes in the composition of worker groups with different separation rates over time. To assess the quantitative importance of composition effects on separation rates, we run a linear prob-

ability model of separation rates on a large set of observable worker characteristics. We run the following regression:

$$\mathbf{1}_{eu,i,t} = X_{i,t}\beta_t + \varepsilon_{i,t},$$

where $\mathbf{1}_{eu,i,t}$ denotes an indicator function that is one if in year t we observe a transition from employment into unemployment of individual i, and where $X_{i,t}$ denotes a vector with dummies for individual characteristics of individual i in year t, β_t denotes the coefficient vector that we allow to vary across years, and $\varepsilon_{i,t}$ denotes the error term. We include dummies for gender, age, education, employment duration, temporary work, and wage percentiles. We pool all transitions of one year in the regression so that one worker can have multiple transitions within one year. Predicted average transition rates are then average population characteristics that we denote by $\bar{X}_t$ times the coefficient vector $\hat{\pi}_{eu,t} = \bar{X}_t\beta_t$. The predicted average separation rate corresponds by construction to the observed average rate.[36]

Figure 16: Separation rates controlling for worker characteristics

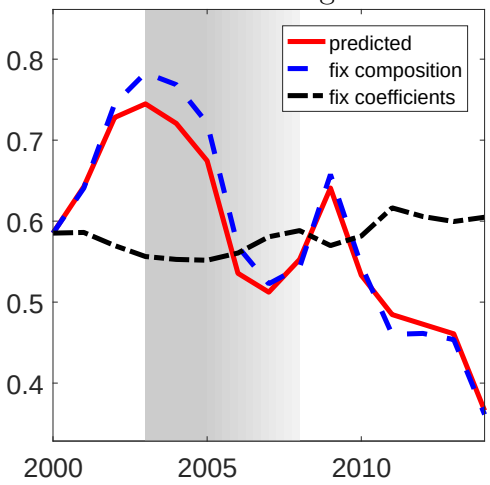

Notes: Yearly averages of monthly separation rates 2000-2014. The solid red line marks the predicted (actual) separation rate. The dashed blue line marks the separation rate keeping the composition of all observables fixed at their level in 2000. The dashed-dotted black line marks the separation rate keeping the coefficients of all observables fixed at their level in 2000. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

We then construct two counterfactual transition rates. For the first counterfactual transition rate, we keep population shares at their level in 2000 and only vary coefficients over time $\tilde{\pi}^{2000}_{eu,t} = \bar{X}_{2000}\beta_t$. This captures changes in separation rates for a fixed population of workers. Through the lens of our structural model in Section 3, these are changes in

[36] We pool all transitions within a year to compute the transition rates. This approach can lead to small deviations in comparison to an average of monthly rates, but in our case, the difference is negligible.

behavior, for example, resulting from changes in the UI system. For the second counterfactual transition rate, we keep coefficients at their level in 2000 and only vary population shares over time $\check{\pi}_{eu,t}^{2000} = \bar{X}_t \beta_{2000}$. This captures the effects from changes in the composition of worker groups. Figure 16 shows the predicted separation rate $\hat{\pi}_{eu,t}$ (solid red line), the counterfactual transition rate with fixed population shares $\tilde{\pi}_{eu,t}^{2000}$ (dashed blue line), and the counterfactual transition rate with fixed coefficients $\check{\pi}_{eu,t}^{2000}$ (dashed-dotted black line). We find that the counterfactual transition rate with changes in coefficients β_t tracks the drop in separation rates over time very closely. The counterfactual transition rate that keeps all coefficients fixed at their level in 2000 and where we only vary population shares over time hardly changes. This evidence strongly supports the idea that it was behavioral changes resulting from changes in the macroeconomic environment that explain the decline in the separation rate over time rather than changes in the composition of the workforce.

C.3 Heterogeneity in transition rates by age groups

This section provides further details on the heterogeneity in the changes in separation rates by age discussed in Section 2.4. Table 10 provides detailed information on separation rate changes by age and employment duration. The upper part of the table shows results for all workers and for three different age groups. Workers age 15-44 show the smallest decline in separation rates (-14.2%), and workers in the age group from 45 to 64 years show the strongest decline in separation rates (-25.2%). These age differences still hide important heterogeneity arising from employment duration. The lower part of Table 10 distinguishes workers by age and employment duration. Here, we find that changes in separation rates mirror the relative differences in changes in benefit eligibility from Figure 4. Short-term employed workers show across age groups a rather uniform decline in separation rates varying between 14.6% and 17.7%. The decline is always less than the average decline over this time period of 22.0%. We also find a much stronger decline for long-term employed workers age 45 and older. Their separation rates decline by 32.5% and 48.8%. For younger long-term employed workers, we find a smaller decline. This is in line with the relative cut in benefits shown in Figure 4 that does not show any variation in the cuts in benefit eligibility among young workers. The larger decline among the oldest age group of long-term employed workers cannot be explained by the cut in benefit eligibility from Figure 4 alone. The longer-run trend in Figure 17(a) suggests that the likely explanation predates the Hartz reforms. The separation rates for the oldest group of workers seem to follow a longer-run downward trend starting in the mid-1990s. A detailed investigation of this trend is of independent interest but beyond the scope of this paper. We leave a detailed investigation of the reasons behind this trend to future

research.

Table 10: Change in separation rates by employment duration and age

	1993-2002	2008-2014	Δ %
age: 15-44	0.72%	0.61%	-14.2%
age: 45-54	0.43%	0.35%	-18.3%
age: 45-64	0.46%	0.35%	-25.2%
age: 15-44, emp. duration $\leq$ 3 years	1.36%	1.13%	-16.8%
age: 15-44, emp. duration $>$ 3 years	0.26%	0.22%	-15.4%
age: 45-54, emp. duration $\leq$ 3 years	1.47%	1.25%	-14.6%
age: 45-54, emp. duration $>$ 3 years	0.18%	0.12%	-32.5%
age: 45-64, emp. duration $\leq$ 3 years	1.48%	1.22%	-17.7%
age: 45-64, emp. duration $>$ 3 years	0.27%	0.14%	-48.8%

Notes: Monthly separation rates before and after the Hartz reforms by employment duration and age. Column Δ reports the percentage change in rates from the period before the Hartz reforms to the period after the Hartz reforms.

Figure 17: Separation and job finding rates by age

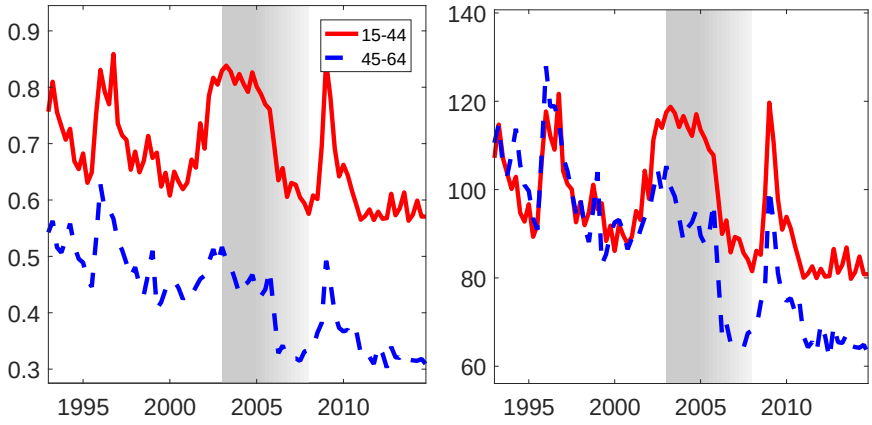

(a) Separation rate (b) Indexed separation rate

Notes: Separation rates for age groups 15-44 years (solid red lines) and 45-64 (dashed blue lines). The left panel shows the level of the separation rate. The right panel shows the change in the separation rate relative to its pre-reform level (1993-2002). The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

C.4 East Germany

For our empirical analysis in Section 2.3, we exclude workers who have employment or unemployment spells in East Germany. We do this to abstract from any effects of a transition of the East German labor market in the decade after reunification. In this section, we explore separation and job-finding rates for East Germany starting in 1995. Figure 18 shows the time series for separation rates and job-finding rates for East German workers and applies the inflow correction described in Section 2.2. The corresponding results for the West German labor market are in Figure 2.

Separation rates in East Germany are higher than in our baseline West German sample. Before the reform, the monthly separation rate is slightly higher than 1.4%. Figure 18(a) shows that separation rates in East Germany plummet in 2006 to 70% of their pre-reform level and in 2014 stand at 50% of their pre-reform trend. The data suggest an ongoing falling trend in the separation rate. Hence, the decline in the separation rate is stronger in the East than in the West German labor market. Regarding job-finding rates, the results are even more striking. Relative to their pre-reform level of 5.4%, the job-finding rate in the East German labor market stands in 2014 at its pre-reform level. All changes in East German unemployment therefore result from a decline in separation rates, thereby further reinforcing our findings from the West German labor market.

Figure 19 provides results on the heterogeneity in the changes in separation rates for the

Figure 18: Changes in separation and job-finding rates East Germany (1995-2014)

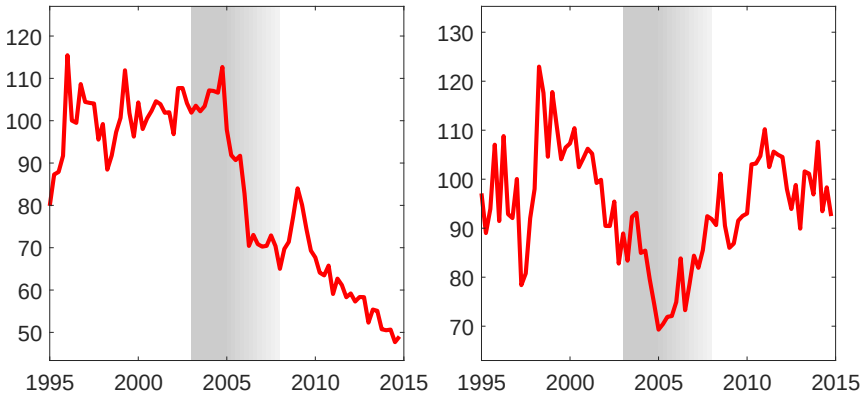

(a) Separation rate (indexed) (b) Job-finding rate (indexed)

Notes: Separation and job-finding rates for East Germany, 1995-2014. Both series have been indexed to their pre-reform level (1995-2002). Both series exclude non-employed entering the unemployment pool in the first half of 2005 who did not become employed until the end of 2006 (inflow correction). The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms. Data are quarterly averages of monthly rates.

East German labor market over time. The corresponding results for the West German labor market are shown in Figure 5.

The changes in separation rates by age and employment duration in the East German labor market corroborate the findings for the West German labor market. We find that long-term employed workers show a much stronger decline than short-term employed workers (Figure 19(a)). Looking at workers in the age range from 15 to 44 years in Figure 19(b), we find a roughly equal decline by 50% from the pre-reform period to 2014. The short-term employed typically show a slightly smaller decline than the long-term employed but also started from a higher level in 2005. For workers in the age group 45-64 years, we find a much stronger decline for the long-term employed, in line with our results for the West German labor market (Figure 19(c)). Separation rates for the long-term employed workers decline roughly 20% more than those for the short-term employed workers. The average decline in East Germany is larger. Finally, when comparing short-term employed workers in the age group 15-44 years to workers in the age group 45-64, we again find, as in the case of the West German labor market, that their separation rates lie virtually on top of each other and decline in lockstep between 2005 and 2014 (Figure 19(d)).

Figure 19: Separation rates by age and employment duration (1995-2014) for East Germany

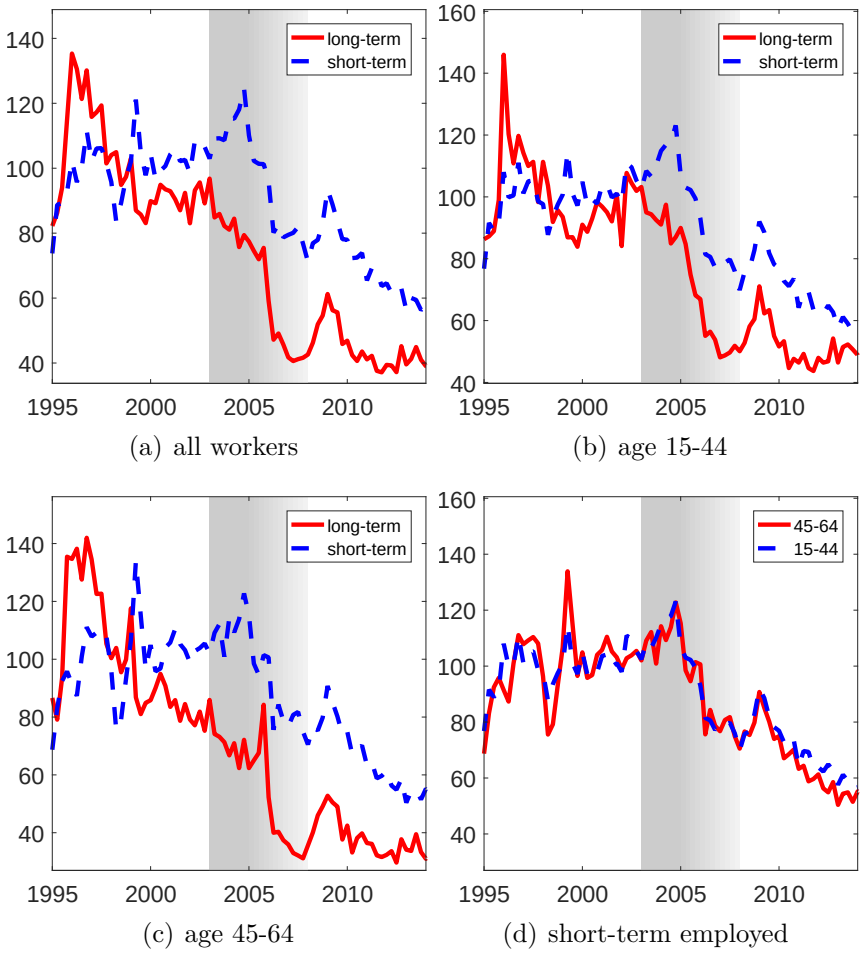

(a) all workers (b) age 15-44

(c) age 45-64 (d) short-term employed

Notes: Separation rates by employment duration and age for East Germany, 1995-2014, indexed to their pre-reform level (1995-2002). The solid red lines in panels (a)-(c) mark the separation rate for long-term employed workers who were continuously employed for three years or more. The dashed blue lines in panels (a)-(c) mark the separation rate for short-term employed workers with at most three years of continuous employment. Panel (d) shows the separation rate for short-term employed workers separately for young (age 15-44, dashed blue line) and old (age 45-64, solid red line) employees. The gray area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms. Data are quarterly averages of monthly rates.

C.5 Including marginally employed

For our baseline sample, we do not define workers as employed if their only employment relationship is under marginal employment regulation. As described in Section B.2, we count these persons as either unemployed or out of the labor force depending on whether or not they have a parallel unemployment spell in that month. A main reason for excluding marginal employment in our baseline sample is to derive consistent time series of worker flow rates. Information on marginal employment becomes comprehensive in the microdata after 1999, so we cannot construct a consistent time series going back to 1993. Before 1999, information on marginal employment is typically not recorded. As a sensitivity analysis, we include all available information on marginal employment when defining employment states. Figure 20 shows the separation rates and job-finding rates including marginal employment information in comparison to the rates from the baseline sample.

Figure 20: Separation and job-finding rates including marginal employment

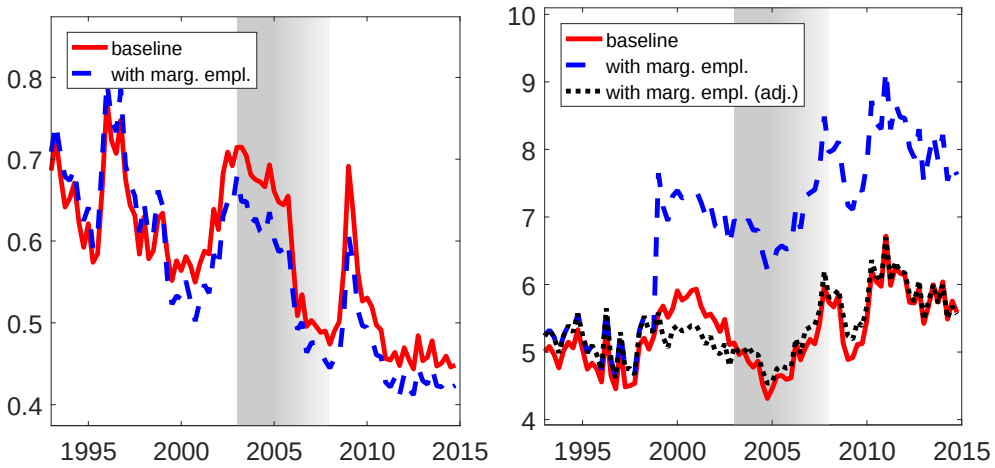

(a) Separation rates (b) Job-finding rates

Notes: Separation rates and job-finding rates in West Germany, 1993-2014, for the baseline sample (solid red line) and for a sample in which marginal employment is included in the employment definition (dashed blue line). The dotted black line in the right panel shows the job-finding rates including the marginally employed adjusted for the structural break in 1999. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

Figure 20(a) shows separation rates for the baseline sample (solid red line) and the sensitivity sample including marginal employment information (dashed blue line). Marginal employment accounts for only a small fraction of total employment so that the change in aggregate separation rates is small. The decline in separation rates becomes slightly more pronounced in the sensitivity sample, and including marginal employment would lead to

a larger decline in separation rates compared to the baseline sample. Figure 20(b) shows job-finding rates from the baseline sample (solid red line) and sensitivity sample (dashed blue line). The job finding rate in the sensitivity sample shows a structural break in 1999 when complete information on marginal employment becomes available. We provide an additional estimate for the sensitivity sample, where we remove the structural break by removing the level shift (dotted black line).[37] We find that after we remove the structural break in 1999, the job-finding rates from the baseline and sensitivity sample track each other closely. If anything, the job-finding rate in the adjusted sensitivity sample is slightly higher before 1999, implying a slightly smaller increase in job-finding rates after the reform. We conclude that our empirical findings on the importance of the decline in separation rates are robust to a change in the employment definition to include marginal employment information.

C.6 Effect of active labor market policy

Section 2.2 discusses changes in regulation for unemployment registration and the inflow correction to adjust for this change. A second change that affects the microdata records that was enacted as part of the Hartz reforms was the treatment of active labor market programs. Starting in 2005, unemployed persons who participate in training programs, internships, or other measures that are part of active labor market policy are no longer recorded as unemployed in the microdata while they are taking part in such programs. Our baseline definition of employment states assigns workers in active labor market programs as out of the labor force. If these workers go from a program to regular employment, the baseline sample would not count this as a transition from unemployment to employment; as a consequence, the job-finding rate would be lower. To explore the quantitative effect of this change in recording, we exploit the information from the unemployment records that list a reason for why the worker is no longer registered as unemployed. We exploit this information to identify workers who participate in active labor market programs and explore how our estimates for job-finding rates are affected if we include workers as unemployed while they are in active labor market programs. Figure 21 shows the unemployment rate and the job-finding rate for the baseline sample and for the sensitivity sample that still counts all participants in measures of active labor market programs after 2005 as unemployed if they were unemployed before the program started.[38]

Looking at the unemployment rate in Figure 21(a), we find a very small increase in

[37] The level shift at the structural break corresponds to a 37% increase in the job-finding rate in the sensitivity sample.

[38] Because of the inflow correction, the samples differ slightly before 2005.

Figure 21: Unemployment and job-finding rates including active labor market programs

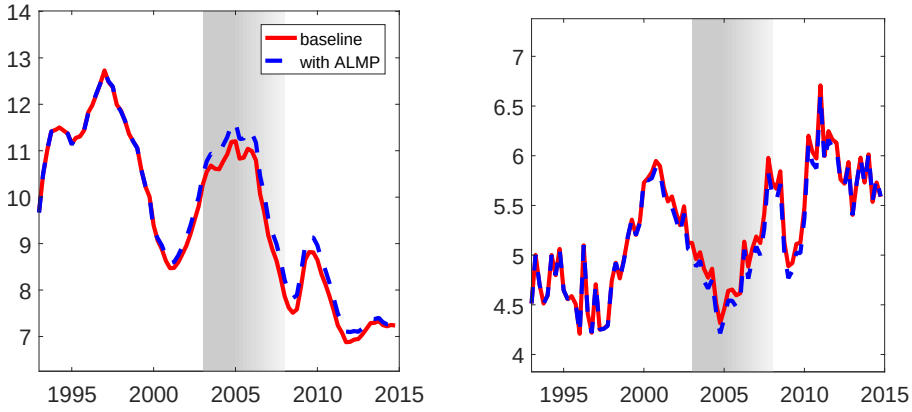

(a) Unemployment rate (b) Job finding rate

Notes: Unemployment and job-finding rates from the baseline sample and a sensitivity sample that includes workers in active labor market policy (ALMP) programs in the group of the unemployed. See text for further details. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

unemployment, yet the effect is negligible. Job-finding rates in Figure 21(b) are hardly affected. We conclude that the change in the recording of active labor market programs in the microdata has a quantitatively negligible effect on our results.

D Unemployment benefit statistics

This section reports supplementary statistics on unemployment benefit recipients. Section D.1 reports the share of unemployed workers who receive supplementary benefits as their benefit level for unemployment insurance benefits is below subsistence level. Section D.2 reports the number of recipients of supplementary benefits during a transition period after the reform to cushion the impact of the reform.

D.1 Low-wage workers

Section 2.4.3 relies on low-wage workers as a control group for the treatment effect of the UI reform. The reason low-wage workers can serve as a control group stems from the institutional setup of the German social security and unemployment insurance system. Unemployment benefits are provided as replacement rates to the last wage of workers. For some low-wage workers, these UI benefits are below the subsistence benefit level (including rents). In case of unemployment, these workers receive the higher social security benefits. These benefits remained largely unaffected by the reform, and the (potential) recipients of these benefits can therefore serve as a control group for the impact of the reform on

separation rates. A direct identification in the microdata is not possible as the level of subsistence benefits also depends on the family situation of workers, which remains unobserved in the social security data for employed workers. We therefore capture the group of workers by the lowest decile of the wage distribution. The lowest decile is in line with a share of 10% of the unemployed receiving social security benefits despite being eligible to receive UI benefits. Figure 22 shows the share of unemployed workers starting in 2007. No data are available for the period before 2007, and we rely on the available 12 years of data to estimate the share of workers who constitute the group of low-wage workers who remained unaffected by the reform.

Figure 22: Share of unemployed with benefit entitlement below subsistence level (*Aufstocker*)

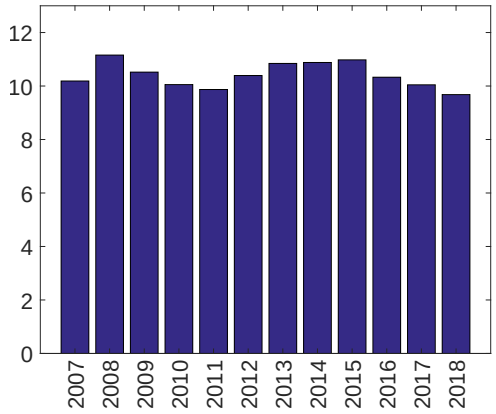

Notes: Share of unemployed workers with benefit entitlement below subsistence level who receive supplementary social assistance benefits. Shares computed as average annual stocks and shown as percentage points. Data from German employment office (Bundesagentur für Arbeit). Data only available starting in 2007.

D.2 Transition period

In our analysis, we consider the years from 2005 to 2008 as the transition period after the reform. The reason is that between 2005 and 2010, unemployed workers received benefits in addition to their subsistence benefits according to regulation in §24 SGB II. The regulation aimed at cushioning the transition from the pre-reform to the post-reform system. The regulation was abolished by the end of 2010. Figure 23 that shows the number of recipients of these supplementary benefits declined strongly between 2005 and 2008 when it leveled off. Data start in June 2006, and the dashed line indicates that these data have been extrapolated based on the available data. We take this as evidence that the period from 2005 to 2008 was a transition period after the reform. Given that the number of recipients declines almost linearly, we also linearly implement the impact of reform in our model analysis over this time period.

Figure 23: Number of workers with additional benefits according to §24 SGB II

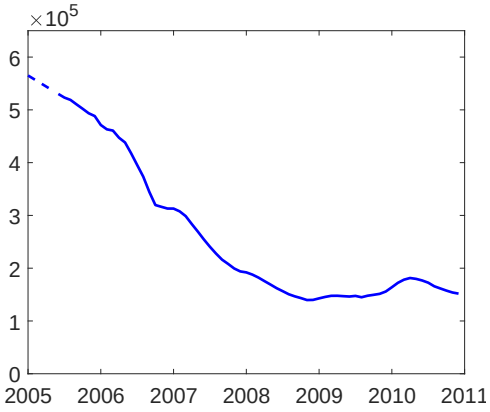

Notes: Number of workers with additional benefits according to §24 SGB II. Data start in June 2006, and the dashed line indicates that these data have been extrapolated based on the available data.

E Estimation of synthetic control counterfactual

We sketch here how we implement the synthetic control estimation for the counterfactual evolution of the unemployment rate in Germany absent the Hartz reforms (Figure 12). For details of the original method, we refer to Abadie et al. (2010). For our implementation of the estimation, we refer to Born et al. (2019). We construct the synthetic control from a pool of 17 OECD countries (donor pool). We restrict the sample to countries that provide quarterly unemployment rate data for the entire time period from 1993Q1 to 2014Q1. For Germany, we use unemployment rates as reported by the German employment office and which we discuss in Section 2. Importantly, we assume that the effect of the Hartz reforms materializes after 2003Q1, so we use data until 2002Q4 for the construction of the control group. The estimation of the control group determines a set of weights for the countries from the donor pool to minimize the mean squared distance between the German unemployment rate (treatment group) and the weighted unemployment rate of the countries from the donor pool (control group) for the pre-Hartz reform decade from 1993Q1 to 2003Q1. We abstain from additional weighting and use a diagonal weighting matrix that assigns equal relevance to the countries in the donor pool. The estimated weights to form the control group are positive for six countries where we consider weights positive if they are above the threshold of 0.01. The synthetic control group is composed of Belgium (0.05), Italy (0.14), Luxembourg (0.17), Austria (0.14), Sweden (0.35), and Japan (0.15), with the estimated weights in parentheses. Countries from the donor pool with weights of less than 0.01 are the Czech Republic, Denmark, Ireland, Spain, France, the Netherlands, Portugal, Finland, the United Kingdom, Norway, and the United States. We demonstrate that the weighted average unemployment rates of the synthetic control

group are able to match almost exactly the evolution of the German unemployment rate for the decade before the Hartz reforms. For the estimated counterfactual evolution of the unemployment rate for Germany in the absence of the Hartz reforms, we fix the estimated weights and construct the unemployment rate for Germany as the weighted average as before.

F Alternative explanations

The German labor market miracle has received a lot of attention in the public debate as many European countries struggle with high unemployment rates and the question about whether Germany can serve as a role model for these countries has been discussed. Our paper provides new empirical evidence that the UI reforms that were part of the Hartz reforms in the mid-2000s are responsible for the German labor market miracle. Other explanations for the German labor market miracle have been proposed, and we discuss three of the most prominent alternative explanations in this section. A prominent idea put forward by Dustmann et al. (2014) is falling unit labor costs in Germany relative to its European neighbors. In Section F.1, we look at changes in unemployment rates and unit labor costs in a cross section of European countries to explore this idea. A closely related explanation is an export-driven boom in the labor market arising from demand from China and other European countries. In Section F.2, we compare worker flows in export industries to those that produce for the domestic market. Finally, Section F.3 looks at the composition of unemployment and the share of long-term unemployed. One goal of the Hartz reforms was to reduce long-term unemployment, and that might have contributed to a decline in unemployment after the reforms.

F.1 Unit labor costs

One prominent narrative of the German labor market miracle comes from Dustmann et al. (2014), who link the German labor market miracle to a decline in Germany's unit labor cost and wages relative to other European countries. Figure 24 shows the ratio of German unit labor costs (ULC) (solid red line) and unemployment rate (dashed blue line) relative to other European countries.[39] We find that the relative decline in Germany's ULC had already started in the 1990s, long before the Hartz reforms were enacted. The idea behind how changes in ULC are related to changes in unemployment rates is that with falling ULC, production became relatively cheaper in Germany, thereby increasing labor demand in Germany at the cost of falling labor demand in other European countries.

[39] EU-18 is the employment-weighted average of 18 EU countries: Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Hungary, Ireland, Italy, Luxembourg, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, and the United Kingdom.

As a result, stronger labor demand leads to declining unemployment rates in Germany, and in comparison, unemployment rates in other European countries with weakened labor demand rise. The relative unemployment trends in Figure 24 paint a different picture. During the period when German ULC were falling relative to other European countries, German unemployment was rising in comparison to these countries, whereas the fall in unemployment rates between 2005 and 2014 was accompanied by stagnating or even increasing relative ULC in Germany. Without drawing causal conclusions, these negatively correlated time series are challenging to reconcile with the hypothesis that declining ULC were the main driver of the German labor market miracle.

Figure 24: Unit labor costs and unemployment rates in the EU

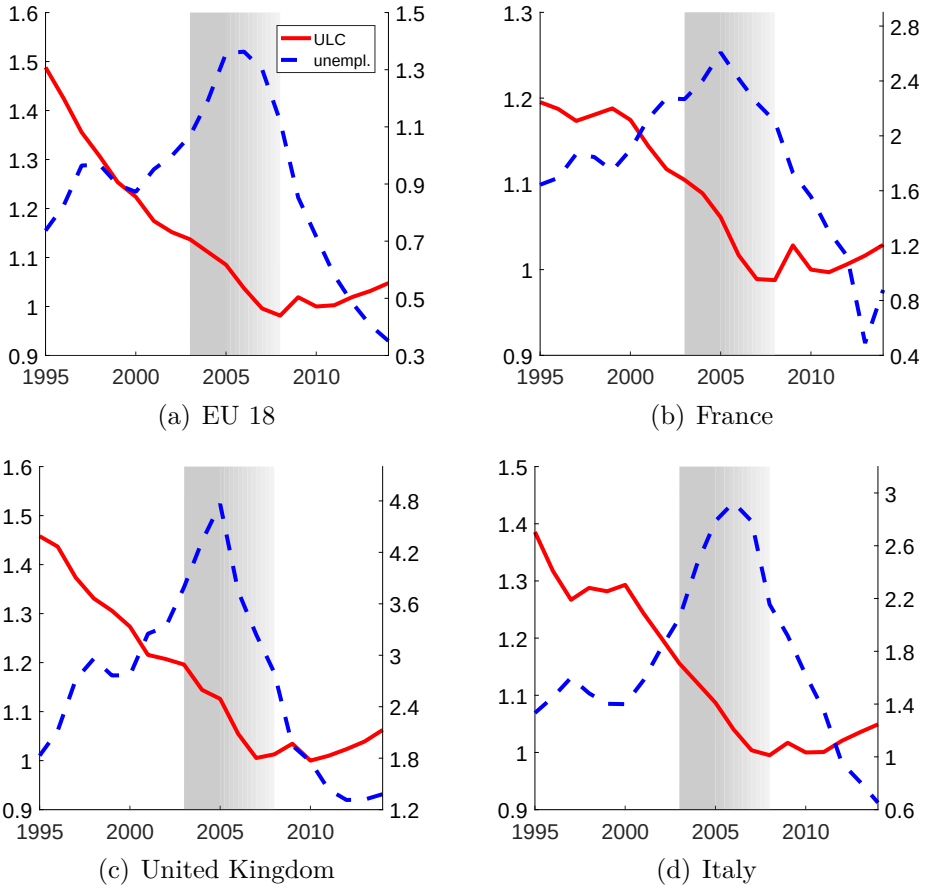

(a) EU 18 (b) France

(c) United Kingdom (d) Italy

Notes: The solid red lines (left axis) show ULC in Germany divided by ULC in other European countries. The dashed blue lines (right axis) show the German unemployment rate relative to the unemployment rate in these countries. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

F.2 Export demand

Related to the idea of falling unit labor costs in Germany relative to its European competitors is the idea that rising export demand for German goods has spurred labor demand in the German labor market and has led to a decline in separation rates contemporaneously with the labor market reforms. To investigate this idea, we follow Dauth et al. (2016) and classify industries by their export or import exposure. If labor demand from abroad is a key driver of the labor market miracle, we expect separation rates to decline in the export-exposed industries relative to import-exposed or unexposed industries. We adopt the classification of import- and export-exposed industries from Dauth et al. (2016). Their classification follows the methodology used in the seminal paper by Autor et al. (2013) on the impact of Chinese import competition on industries in the United States. Import exposure is defined as the absolute value of trade flows into an industry from a particular region relative to the trade flows into that industry stemming from all countries; export exposure is defined equivalently. Dauth et al. (2016) classify 93 industries at the 3-digit level in the SIAB sample. Out of these industries, they report the 25 most import-exposed and the 25 most export-exposed industries in Table A.2 of their appendix. We take the classification from this table to define the import-exposed and export-exposed industries in our sample.

Table 11: Change in separation rates by trade exposure

industries	1993-2002	2008-2014	Δ %
all	0.63%	0.49%	-22.0%
no exposure	0.63%	0.50%	-20.0%
export exposure	0.44%	0.26%	-40.8%
import exposure	0.62%	0.34%	-44.4%

Notes: Monthly separation rates before and after the Hartz reforms by trade exposure. Column Δ reports the percentage change in rates from the period before the Hartz reforms to the period after the Hartz reforms.

Table 11 shows the decline in separation rates across industries with no trade exposure, with export exposure, and with import exposure. Industries with neither import nor export exposure show a decline in separation rates from the pre-reform to the post-reform period of 20%, close to the average decline of 22%. We find that both import-exposed and export-exposed industries show much stronger declines relative to other industries. Separation rates in trade-exposed industries decline by more than 40%. The effect on the

overall separation rates remains modest, however, because employment shares of these industries are small and account for less than 10% of employment. Furthermore, the declines in the import-exposed and export-exposed industries are about the same size so export-induced demand as a driver of the decline in separation rates seems at odds with the observed pattern. The evidence therefore does not support a prominent role for export-driven labor demand as an explanation for the German labor market miracle.

F.3 Long-term unemployment

One goal of the Hartz reforms was to reduce long-term unemployment. To explore whether the reduction in long-term unemployment was an important driver behind the reduction in overall unemployment, we examine the composition of the unemployment pool over time. If the Hartz reforms increased, in particular, the job-finding rates of long-term unemployed, this should have shifted the composition of the unemployment pool toward short-term unemployment after the reforms. Figure 25 shows that apart from a spike during the Great Recession between 2006 and 2008, the share of unemployed who have been out of work for more than one year did not change relative to the pre-Hartz period.

Figure 25: Share of long-term unemployed in Germany (1998-2015)

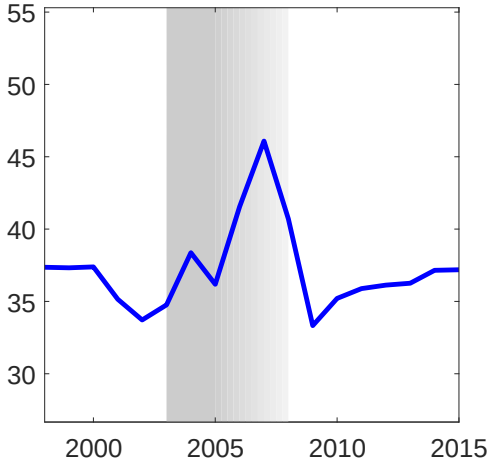

Notes: Share of long-term unemployed (more than 12 months) in total unemployment. The grey area marks the period 2003 to 2005 when the Hartz reforms were enacted. The fading out indicates the first transition years 2006 to 2008 after the reforms.

The analysis in Section 4 provides evidence that higher job-finding rates are not the main reason for the decline in the unemployment rate after 2005. The fact that we find a largely constant share of long-term unemployed after the Hartz reforms implies that there have also been no differential effects among the unemployed and a stronger increase in job-finding rates among the long-term unemployed.