DISCUSSION PAPER SERIES

Worker Churn and Employment Growth at the Establishment Level

Rüdiger Bachmann
Christian Bayer
Christian Merkl
Stefan Seth
Heiko Stüber
Felix Wellschmied

DISCUSSION PAPER SERIES

Worker Churn and Employment Growth at the Establishment Level

Rüdiger Bachmann
University of Notre Dame, CEPR, ifo, CESifo

Christian Bayer
University of Bonn and CEPR

Christian Merkl
*University of Erlangen-Nuremberg (FAU)
and IZA*

Stefan Seth
Institute for Employment Research (IAB)

Heiko Stüber
FAU and IAB

Felix Wellschmied
*Universidad Carlos III de Madrid
and IZA*

ABSTRACT

Worker Churn and Employment Growth at the Establishment Level*

We study the relationship between employment growth and worker flows in excess of job flows (churn) at the establishment level using the new German AWFP dataset spanning from 1975–2014. Churn is above 5 percent of employment along the entire employment growth distribution and most pronounced at rapidly-adjusting establishments. We find that the patterns of churn along the employment growth distribution can be explained by separation rate shocks and time-to-hire frictions. These shocks become larger on average during boom periods leading to procyclical worker churn. Distinguishing between separations into non-employment and to other establishments, we find that separations to other establishments drive all procyclical churn. In a secondary contribution, we compare German worker and job flows with their US counterparts and recent US findings.

JEL Classification:	E20, E24, E32, J23, J63
Keywords:	job flows, worker flows, churn, job-to-job transitions, aggregate fluctuations

Corresponding author:
Rüdiger Bachmann
Department of Economics
University of Notre Dame
3026 Nanovic Hall
Notre Dame, IN 46556
USA

E-mail: rbachman@nd.edu

1 Introduction

Many establishments both hire and separate from workers within relatively narrow windows of time. This leads to worker turnover in the economy that is larger than what would be necessary to accommodate observed job creation and destruction (see Burgess et al. (2000) and Davis et al. (2006, 2012)). These worker flows in excess of job flows, in short worker churn, are quantitatively large (on average much larger than job flows) and increase by about 40 percent during booms relative to recessions.

In this paper, we use micro data to study the relationship between establishment growth (job creation and destruction) and worker churn. Our analysis offers new insights into the sources of worker reallocation, the shocks and frictions establishments face when adjusting employment, and the way business cycles propagate through endogenous worker reallocation.

For our analysis, we use the new Administrative Wage and Labor Market Flow Panel (*AWFP*) for Germany. The data comprises the entire universe of German establishments from 1975–2014 at the quarterly frequency. It allows us to link establishment growth to hiring decisions (from other establishments and non-employment) and separation decisions (to other establishments and non-employment). In comparison to the US data, aggregate job and worker flows are about half the size in Germany. In both countries, worker turnover is almost twice as large as job turnover. What is more, flow rates have similar cyclical properties in the two countries.[1] The aggregate separation rate is procyclical, the job destruction rate is countercyclical, and the hiring rate is more procyclical than the job creation rate. Moreover, worker flows are more volatile and persistent than job flows leading to persistent and volatile procyclical worker churn in the aggregate.

In the cross-section, the churning rate is lowest for establishments that do not change, or change little, their number of workers. It grows in absolute employment growth, i.e., rapidly-shrinking and rapidly-growing establishments have the highest churning rate. In other words, establishments that decrease the number of workers often also hire. On average, they hire more than establishments with a constant work force. Analogously, establishments that increase the number of workers often also separate from some workers. On average, they separate from more workers than establishments with a stable number of workers.

This observation cannot be explained with simple models of employment dynamics, where establishments face a constant separation rate and make employment adjustments in reaction to idiosyncratic productivity shocks. In such a framework, rapidly-shrinking establishments are at or above their employment target and thus

---

[1]Davis et al. (2012) provide a comprehensive overview for the US data. AWFP has advantages over US data. One major obstacle for studying links between job and worker flows in the US is the availability of data sets that provide information on establishment characteristics, worker flows, and job flows. The most commonly used US data source is the Job Openings and Labor Turnover Survey (JOLTS), used by Davis et al. (2006), sampling 16,000 establishments in the US on a monthly basis. However, JOLTS only started in 2001, providing data on at most two full business cycles. By contrast, the German AWFP, similar to the LEHD data analyzed by Abowd and Vilhuber (2011), contains quarterly information on job and worker flows of all full-time employees working for all German establishments from $1975 - 2014$. This allows us to systematically study the cyclical behavior of job and worker flows and their interaction.

have no incentives to hire any workers. In fact, in such a setup, churn is highest for establishments with a constant number of workers, where every hire is a replacement of an exogenous and predictable separation.

The fact that rapidly-shrinking establishments also hire substantial numbers of workers implies that these establishments separate from more workers than they had planned or had foreseen. We interpret this as stochastic separation rate shocks to the establishment. Stochastic separations from the point of view of the establishment may reflect that workers find better employment opportunities outside the given establishment (a job ladder view), or that the establishment learns that some of its workers are not a good match (a mismatch view). In addition, what matters is that these separation shocks cannot be undone immediately; and as a result, establishments make ex post planning mistakes with respect to their employment stocks because separation rate shocks drive a wedge between desired and actual employment levels. We then argue that these separation rate shocks, in addition to productivity shocks, are an important source of uncertainty for establishments and drive short-term employment fluctuations. If it takes time to hire, establishments will try to rehire the separations they expect in excess of their desired employment changes. When separations realize below this value, the establishment grows. If more separations than anticipated happen, the establishment will shrink. Since this expectation error is by definition unrelated to the desired establishment growth, it can produce large average churning rates in fast-growing or shrinking establishments.

Next, we study the cyclical properties of worker churn. During booms, relative to recessions, the churning rate becomes larger along the entire employment growth distribution, and the distribution of employment growth shifts to the right. However, from a statistical perspective, the latter is negligible for the cyclical movements in the aggregate churning rate. Rather, the aggregate churning rate is driven by changes in the churning rates conditional on employment growth. This property of churn is remarkably different from the underlying worker flows, where both cyclical shifts in the employment growth distribution and shifts in worker flows conditional on the employment growth distribution contribute to aggregate worker flow rates. When we look at the data through the lens of our stylized model of employment dynamics, separation shocks are on average larger, but less dispersed, during a boom.

Our data allows us to decompose separations (and hires) into those going to other establishments and those going to non-employment. We find that separations (and hires) to other establishments shift up along the employment growth distribution in a parallel fashion during a boom (relative to a recession). Worker transition rates through the non-employment pool show no such cyclical behavior. We show that, as a result, cyclical aggregate worker churn is almost identical to the procyclical job-to-job transition rate. What is more, after subtracting job-to-job transitions from hires and separations, respectively, worker flows have the same cyclical properties as job flows. The hiring rate from non-employment is almost identical to the job creation rate. Similarly, the separation rate into non-employment is almost identical to the job destruction rate which means that the separation rate becomes countercyclical once we subtract the job-to-job transition rate. Put differently, booms are times of high job creation and high churn (not high job destruction), which means that churn-induced separations, that is, job-to-job transitions, get ultimately replaced by some

establishments through hiring from non-employment; and vice versa for recessions. What is more, in terms of timing, job creation and churn both start early in a boom, but the latter is more persistent and continues to increase into the maturing boom.

These findings also contribute to a recent literature that highlights the link of observable establishment characteristics with cyclical hiring and separation decisions and the resulting establishment growth. We show that cyclical churn (and thus cyclical job-to-job transitions) are not systematically linked to establishment *growth*. Moscarini and Postel-Vinay (2012, 2013) develop a framework that links establishment *size* to cyclical job-to-job transitions. Large firms grow during booms on expense of small firms by poaching workers from small firms in a procyclical way. Haltiwanger et al. (2015) question such poaching behavior and show that establishment *pay* is a better predictor for cyclical employment growth and poaching patterns.

Our findings are in line with theories that stress job-to-job transitions as means of procyclical worker reallocation (for example, Moscarini and Postel-Vinay (2013), Schaal (2015), and Fujita and Nakajima (2016)). These theories suggest that during a boom, workers systematically reallocate from low to high productivity establishments. At the same time, our results suggest that we require more heterogeneity than implied by a common ranking (productivity) of establishments by all workers. In common-rank models, highly ranked establishments have low separation rates (and hence churning rates) on average. As higher ranked establishments grow more during a boom than during recessions (see Moscarini and Postel-Vinay (2012, 2013)), the composition of high-growth establishments shifts towards high-rank and hence low-churn establishments. Therefore, in common-rank models, one observes during booms that churn increases at the low end of the employment growth distribution by more than at the upper end. Our data by contrast shows that separations to other establishments, and hence churn, shift up equally across the employment growth distribution during a boom. Therefore, while our results suggest that workers climb the job ladder faster in booms, they also suggest that the ranking of establishments across the ladder is worker specific.

The remainder of this paper is organized as follows. Section 2 introduces the new *AWFP* dataset and explains the main concepts that we use to analyze the data. Section 3 analyzes aggregate job and workers flow dynamics. Section 4 links churn to establishment growth in the cross-section. Section 5 studies the cyclical dynamics in the churning rate. Section 6 connects our empirical finding to models of procyclical labor reallocation and Section 7 concludes.

2 Dataset and Variable Definitions

2.1 The Administrative Wage and Labor Market Flow Panel

The new Administrative Wage and Labor Market Flow Panel (*AWFP*) measures employment, labor flows, and wage data[2] for the universe of German establishments (*Betriebe*) for the years 1975–2014. The *AWFP*s main data source is the Employment

---

[2]Merkl and Stüber (2016) use the *AWFP* to analyze the effects of different wage dynamics on labor flows.

History (*Beschäftigten Historik*, BeH) of the German Institute for Employment Research (IAB). The BeH is an individual-level dataset covering all workers in Germany subject to social security.[3] The information in the BeH originates from the notification procedure for social security. Essentially, this procedure requires employers to keep the social security agencies informed about their employees by reporting any start or end of employment and by annually confirming existing employment relationships.

From the BeH, the *AWFP* aggregates the worker and job flow information to the establishment level, rendering an establishment the observational unit.[4] To ensure consistency over time, most variables in the *AWFP* — and all variables used in the paper — are calculated on a 'regular worker' basis. In the *AWFP* a person is defined as a 'regular worker' when she is employed full-time and belongs to one of the following person groups: 'employees subject to social security without special features', 'seamen' or 'maritime pilots'. Therefore (marginal) part-time employees, employees in partial retirement, interns, etc., are not counted as regular workers.

The *AWFP* covers the time period 1975–2014 (West-Germany until 1992 and the re-unified Germany thereafter). It is available at an annual and a quarterly frequency. For our analysis, we use the *AWFP* at the quarterly frequency and drop all establishments that are on the territory of former East-Germany and Berlin to avoid a break in the series. For further information on the dataset we refer the reader to the *AWFP* data report (Seth and Stüber (2017)).

2.2 Variable Definitions

In the *AWFP*, a worker is considered to be working for a given establishment (henceforth plant) in a given quarter when she is employed at this plant at the end of the quarter.[5] From this definition follows the number of jobs at a plant i at the end of a quarter (J_{it}), the number of hires[6] (H_{it}), as well as the number of separations[7] (S_{it}). These are the time series from the *AWFP* from which almost all data series in our paper are constructed.[8]

Using this basic data, we compute the net job flow at a plant as $JF_{it} = J_{it} - J_{it-1}$. When a plant decreases employment ($JF_{it} < 0$) within a quarter, we count this as job destruction, JD_{it}. When employment increases ($JF_{it} > 0$), we count this as job creation, JC_{it}. A plant may hire and separate from workers within the same quarter, that is, we have $H_{it} \geq JC_{it} \geq 0$ and $S_{it} \geq JD_{it} \geq 0$ for each plant in each quarter.

Part of our analysis deals with differences in plant-level behavior given the amount of employment growth at the plant. For this purpose, we first aggregate the plant-

---

[3]Marginal part-time workers (*geringfügig Beschäftigte*) are included since 1999. The main types of employees not covered by the BeH are civil servants (*Beamte*), military personnel, and the self-employed.

[4]Before this aggregation, the data on individuals are subjected to numerous validation procedures. Further details on the dataset are described in Seth and Stüber (2017). Conceptual differences between the *AWFP* and US Data are discussed in appendix A.2.

[5]It turns out that most workers leave or join a plant at the end/the beginning of a quarter.

[6]A worker that has not been working for that plant at the end of the previous quarter.

[7]A worker that has been working for the plant at the end of the previous quarter.

[8]For some analyses we use data packages of the *AWFP* that provide a decomposition of inflows and outflows according to their source and destination.

Figure 1: Aggregate Job and Worker Flows and the Churning Rate
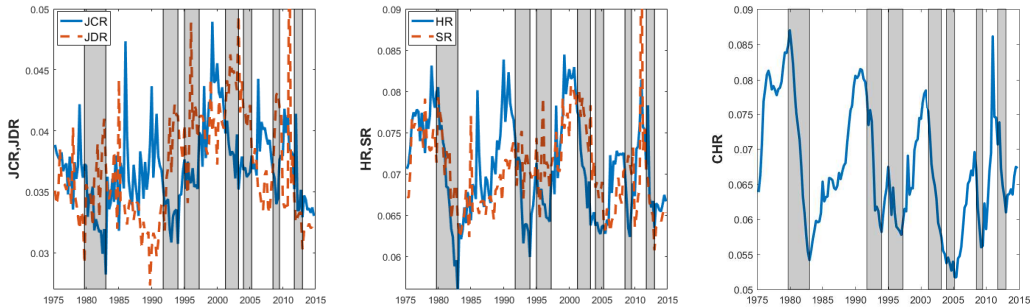
Note: The left panel displays aggregate job flows. JCR: job creation rate, JDR: job destruction rate. The center panel displays aggregate worker flows. HR: hiring rate, SR: separation rate. The right panel displays the aggregate churning rate, CHR. All rates are seasonally adjusted. West German plants only. The gray shaded areas represent periods of at least 5 consecutive quarters of unemployment growth.

level data to 21 employment growth categories/bins. Table A1 in Appendix A.1 provides these growth bins, and Figure A1 provides the time averaged employment share in each of these categories.

We allow each employment growth category to have its own seasonal component and compute seasonally adjusted series, using the X-12 ARIMA CENSUS procedure.[9] To derive the aggregate series for West Germany, we finally sum over the seasonally adjusted series for all employment growth categories.

Given either the aggregated stock/flow data or the stock/flow data by employment growth category, we define aggregate flow rates. We use the average of contemporaneous and lagged end-of-quarter employment as the denominator:

$$N_t = [J_t + J_{t-1}]/2.$$

For example, the hiring rate is given by:

$$HR_t = \frac{H_t}{N_t}. \tag{1}$$

The separation rate (SR), the job-creation rate (JCR), and the job-destruction rate (JDR) are defined analogously. Using the numerator N_t, as defined above, implies that all rates are bound in the interval $[-2, 2]$ with endpoints corresponding to the death and birth of plants.[10]

Most of our analysis deals with fluctuations at the business cycle frequency. To measure the stage of the business cycle, we use the filtered aggregate unemployment rate for West-Germany.[11] If not otherwise stated, we compute the cyclical component for the aggregate or disaggregate-by-employment-growth-rate employing an HP-filter for the series with a smoothing parameter of $100,000$ (following Shimer (2005)). The

---

[9]By allowing for series-specific seasonality, we want to ensure consistency for each variable for the sum of all individual categories and the aggregate series of West Germany.

[10]See Davis et al. (1996) for a more thorough discussion regarding the properties of this measure.

[11]Cyclical unemployment has a strong negative correlation with GDP (-0.71).

cyclical components have, thus, the interpretation of a deviation from a slowly moving non-linear trend. Given that unemployment and job and worker flows are already expressed as rates, we define the cyclical components as absolute deviations from the trend, i.e., they have to be interpreted as percentage point deviations.

3   Aggregate Job and Worker Flows

Table 1: Job and Worker Flows and the Churning Rate

	Mean	*SD*	*AC*(1)	Correlation with U_{t+j}				
				$j = -2$	-1	0	$+1$	$+2$
JCR	3.69%	0.29%	0.52	0.19*	0.08	−0.04	−0.17*	−0.28*
JDR	3.69%	0.36%	0.40	−0.02	0.05	0.15	0.23*	0.29*
HR	7.06%	0.57%	0.82	−0.26*	−0.40*	−0.53*	−0.64*	−0.72*
SR	7.06%	0.47%	0.47	−0.46*	−0.50*	−0.51*	−0.50*	−0.48*
CHR	6.74%	0.76%	0.92	−0.55*	−0.67*	−0.77*	−0.84*	−0.87*

Note: The table displays the properties of the HP(100,000)-filtered aggregate flow rates. *JCR*: job creation rate, *JDR*: job destruction rate, *HR*: hiring rate, *SR*: separation rate, *CHR*: churning rate. *SD*: standard deviation, *AC*(1): first-order auto correlation. A * indicates significance at the 5% level obtained by non-parametric block-bootstrapping with a block length of 20.

In this section, we discuss aggregate job and workers flows in Germany, as well as the aggregate churning rate, and their dynamics. The first two panels in Figure 1 displays the (unfiltered, but seasonally adjusted) job and worker flows over time. The gray shaded areas represent periods of at least 5 consecutive quarters of unemployment growth. The time average quarterly job creation and destruction rate are both around 3.7% (see also column one in Table 1).[12] Worker flows are substantially larger. The time average quarterly hiring and separation rate are both around 7.1%.

Thus, worker turnover in Germany is about twice as high as is required for the observed job turnover. The US shows a similar picture, where time average quarterly job flows are around 7.1% and time average worker flow rates are around 11.8%.[13] Burgess et al. (2000) introduce a measure that quantifies the amount of worker flows in excess of job flows at the plant level, called worker churn:[14]

$$CH_t = (H_t - JC_t) + (S_t - JD_t). \tag{2}$$

[12]It is by chance that the time average flow rates are almost equal.
[13]See tables A2 and A3 in Appendix A.2 for a comparison to US data.
[14]Lazear and Spletzer (2012) and Lazear and McCue (2017) also study worker churn in the US.

Table 2: Correlations of Job and Worker Flows

|	JCR	JDR	HR	SR	CHR
JCR	1.00				
JDR	−0.32*	1.00			
HR	0.81*	−0.29*	1.00		
SR	0.12	0.61*	0.49*	1.00	
CHR	0.45*	−0.19	0.89*	0.66*	1.00

Note: The table displays correlation coefficients of HP(100,000)-filtered flow rates. JCR: job creation rate, JDR: job destruction rate, HR: hiring rate, SR: separation rate, CHR: churning rate. A * indicates significance at the 5% level obtained by non-parametric block-bootstrapping with a block length of 20.

The right panel in Figure 1 displays the aggregate churning rate. On average, churn is around 6.7% of employment each quarter.

The upper panel in Table 1 displays summary statistics of the cyclical component of the job flow rates. The job creation rate is somewhat more persistent but fluctuates less than the job destruction rate. The job creation rate moves counter to the unemployment rate, particularly at leads of unemployment. In contrast, the job destruction rate moves together with the unemployment rate. The second panel in Table 1 displays summary statistics of the hiring and the separation rate. Worker flows are more persistent than job flows and more volatile. Moreover, both rates are procyclical. Taken together, early in a boom (recession), job creation is high (low) and job destruction is low (high). However, worker flows stay high (low) throughout the boom (recession). The fact that the hiring rate rises more than the job creation rate during a boom is made possible by a procyclical separation rate. The last panel in Table 1 shows that worker churn is also procyclical; its contemporaneous correlation with the unemployment rate is −0.77. What is more, it is more persistent and volatile than either job or worker flows. During times of low unemployment, it is about 3 percentage points higher than during times of high unemployment.

Table 2 shows that these relationships lead to the following somewhat complex correlation structure between job and worker flows: The job creation and destruction rate are negatively correlated. Job creation rate and hiring rate, and the job destruction rate and the separation rate are positively correlated. Nonetheless, the hiring and separation rate are also positively correlated, and both are positively correlated with the churning rate.

Figure 2: Churning Rates and Employment Growth
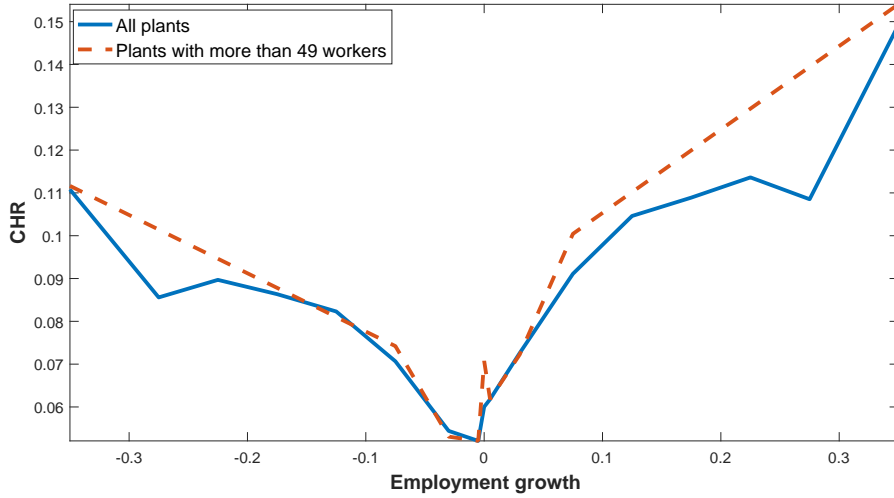
Note: The figure displays the churning rate as a function of the plant specific employment growth rate. Plants are grouped in 17 employment growth categories. We represent the employment growth category by its midpoint as an estimate of the average growth in that category. West German plants only. Pooled data, seasonally adjusted by growth category, quarterly frequency, 1975Q1 - 2014Q4. The red dashed line displays the churning rate for plants with at least 50 employees.

4 Understanding Worker Churn

Intuitively, churn occurs because non-growing plants hire workers, and growing plants separate from workers. Figure 2 displays the rate of churn across the employment growth distribution. It shows basically a U-shaped pattern in employment growth. The larger the absolute rate of employment change, the more a plant churns workers. Importantly, as Figure 2 also shows, this pattern is not exclusively driven by small plants, where small numbers of worker flows imply large flow rates.[15] In other words, growing plants not only hire a large fraction of their workforce, but they also separate from a significant number of workers and they separate from more workers than plants with a constant workforce. Vice versa, plants that shrink hire workers, and they hire more than plants with a constant work force.

This is hard to explain with simple models of plant-level employment adjustment. Instead, it requires that plants do not have full control over the number of workers they employ as we will illustrate next, making use of simple and stylized models of employment adjustment, which are not chosen for detailed realism but serve us as accounting devices to identify the intensity of shocks and frictions needed to generate the observed patterns of churning.

---

[15]We abstract from plants shrinking more than -0.4 or growing more than 0.4, which deviate from the U-shaped pattern (see Table A1 in Appendix A.1). Figure A1 in Appendix A.1, however, shows that these plants contribute little to overall employment. To understand their importance for the aggregate churning rate, we compute the churning rate resulting from the churn of plants growing in the interval $[-0.4, 0.4]$. Figure A2 in Appendix A.1 shows that the resulting churning rate is basically identical to the aggregate churning rate (interval $[-2, 2]$). In other words, this is not a paper about exiting, near-exiting, or entering plants.

4.1 Quadratic Employment Adjustment Costs

We start off with a basic model of employment dynamics at the plant level, where plants have full control of the number of workers they employ. Plants have a decreasing returns to scale production function in employment, and face shocks to idiosyncratic productivity, a constant exogenous separation rate, and quadratic costs of hiring. Let plant i produce output Y_{it} at time t according to the following decreasing returns to scale production function:

$$Y_{it} = z_{it} E_{it}^{\alpha}, \tag{3}$$

where E_{it} is the employment level, z_{it} is idiosyncratic productivity and α (with $0 < \alpha < 1$) is the curvature of the production function. Productivity follows an $AR(1)$ process in logs:

$$\log z_{it} = (1 - \rho)\mu_z + \rho \log z_{it-1} + \epsilon_{it}, \quad \epsilon_{it} \sim N(0, \sigma_\epsilon^2). \tag{4}$$

At the beginning of a period, workers separate from the plant at a constant rate s. The plant actively adjusts its workforce by $\Delta_{E_{it}}^a \in \mathbb{R}$ workers such that the number of workers at plant i evolves according to

$$E_{it} = (1 - s)E_{it-1} + \Delta_{E_{it}}^a. \tag{5}$$

If $\Delta_{E_{it}}^a > 0$, then this active adjustment is counted as hires in the model, i.e., $H_{it} = \left(\Delta_{E_{it}}^a\right)^+$. If $\Delta_{E_{it}}^a < 0$, we count the active adjustment as additional separations in the model, i.e., $S_{it} = sE_{i,t-1} + \left(\Delta_{E_{it}}^a\right)^-$. The plant decides on $\Delta_{E_{it}}^a$ after observing its productivity, i.e., it has full command over the number of workers used in production and no planning lag. Actively adjusting the number of workers is subject to quadratic adjustment costs: $c_{it} = \psi\left(\Delta_{E_{it}}^a\right)^2$. Plants choose their active employment adjustment to maximize the sum of expected profits which they discount at rate r given a wage rate w.

It is straightforward to see that for negative employment growth rates smaller than $-s$, there is no hiring and thus churn is zero – different from the data. For employment growth rates larger than $-s$, plants rehire for the workers lost through separations. Yet, as separations are a fixed fraction of employment, the model cannot produce the fact that fast-growing plants not only hire more, but also separate more from workers. Since we use the same definitions of rates in the model as in the data, i.e., based on the average employment between two adjacent periods, the separation rate, $s\frac{E_{it-1}}{(E_{it-1}+E_{it})/2}$, and analogously the churn rate even slightly decline in plant growth.

To obtain a quantitative impression of the differences between model and data, we calibrate the model and display the churn rates by employment growth in Figure 3. The parameters of this simple model are the wage, w, the returns to scale, α, the quarterly interest rate r, the mean of the log productivity process, μ_z, the autocorrelation, ρ, the standard deviation of productivity shocks, σ_ϵ, the separation rate, s, and the adjustment cost parameter, ψ.

We assume a quarterly interest rate of 0.01, set $\alpha = 0.6$, and normalize the wage to $w = 1$. We set ρ to 0.9675 as estimated by Bachmann and Bayer (2014) and

Figure 3: Churning Rates in a Model with Productivity Shocks
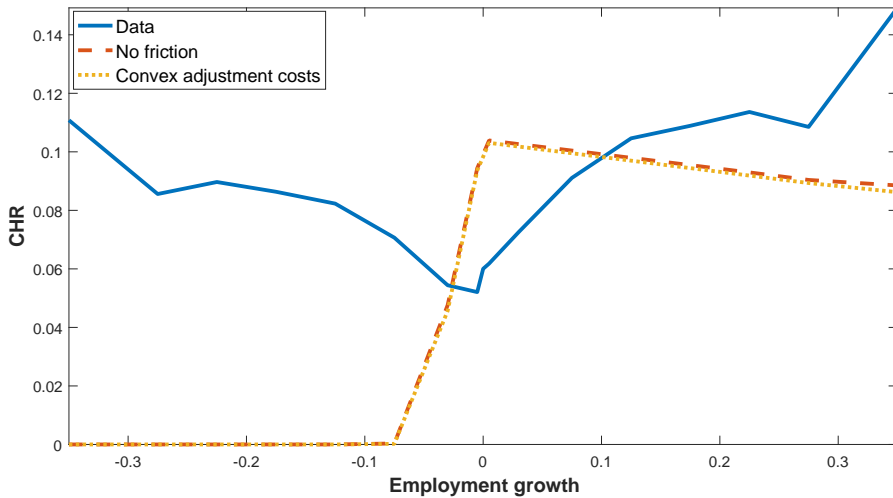
Note: The figure shows churning rates as a function of the plant-specific employment growth rates. We represent the employment growth category by its midpoint as an estimate of the average growth in that category. The blue solid line is to the empirical churning rates for the West-German sample 1975-2014. The yellow dotted line, *convex adjustment costs*, is the churning rates from the optimal active employment adjustment policy of plants in a model with productivity shocks and convex adjustment costs. The red dashed line, *No friction*, is the churning rates in the same model but adjustment costs are set to zero.

use μ_z to match the average plant size in our data of 12.6. We obtain the other three parameters, σ_ϵ, s, ψ, from an equally weighted simulated minimum distance estimator. Our moments are the aggregate separation rate and the churning rate at the sixteen employment growth categories.[16] Column (2) in Table 3 displays the estimated parameters.

Figure 3 compares the churning rate over the employment growth distribution in the model and the data. The model fails to generate any churn at rapidly-shrinking plants. These plants experience negative productivity shocks and desire to shrink; thus, they do not hire any workers. Plants with positive productivity shocks desire to grow. The churn at these plants is basically given by the exogenous separation rate s. Convex adjustment costs turn out to be of little importance to understand churn in the present framework, as Figure III also demonstrates.

4.2   Separation Shocks

Large churn at rapidly-shrinking plants suggests that more workers separate from these plants than these plants desire - hence they rehire. Conversely, it suggests that some plants shrink *because* workers and plants separate as opposed to plants separating from workers in order to shrink. These separations may result from workers finding a better employment opportunity, or from plants firing workers after new in-

---

[16]We have in total 17 employment growth categories. However, given our assumption of a continuous shock distribution and convex adjustment costs, constant employment is a zero-probability event.

Table 3: Parameter Estimates

	Convex costs		Time-to-hire		
	Productivity shocks	Separation shocks	Average	Recession	Boom
μ_z %	1.48	1.67	1.52	1.50	1.53
σ_ϵ %	0.06	–	–	–	–
s/μ_s	0.05	-3.44	-3.09	-3.26	-2.99
σ_s	–	0.70	1.08	1.20	1.00
ψ %	0.78	300.10	–	–	–

Note: The table shows the parameter estimates for the different models of churn and employment growth. μ: mean of log-productivity. σ_ϵ: standard deviation of log-productivity shocks. s: separation rate in the model with productivity shocks. μ_s: mean of log-separation rate shocks. σ_s: standard deviation of log-separation rate shocks. ψ: scaling parameter of the quadratic active adjustment cost function.

formation about these workers has arrived (e.g., a lower match quality). Given data limitations, we are silent on distinguishing between these two explanations.

To understand how important these stochastic separations are for churning, we extend our model of employment dynamics to feature stochastic separations. For clarity, we assume that all plants have a common productivity level μ_z. Instead, similar to the structure of productivity shocks, we assume that the separation rate follows a (truncated) log-normal distribution:

$$\log s \sim \mathcal{N}(\mu_s, \sigma_s^2).$$

It is straightforward to see how churn arises at shrinking plants in this framework. Plants lose workers and they rehire. Without any adjustment costs, plants would have zero employment growth. Adjustment costs lead to plants rehiring only a part of their lost workforce. Since the marginal benefit from rehiring increases more than linearly in the distance from optimal employment, plants that lose many workers rehire a larger fraction of these worker losses. Hence, churn is larger for rapidly-shrinking plants. These plants had large separation shocks and they lean more strongly against the wind than plants with small shocks.

Conversely, positive employment growth arises because a plant having had large separation shocks in the past has a too low employment stock, rehires and thus grows. Fast-growing plants are now either plants with little separation in the current period, or plants with a particularly small employment stock to start with. This implies that the larger employment growth, the larger the fraction of plants with both a lot of hiring and a lot of separations.

11

Figure 4: Churning Rates in a Model with Separation Shocks
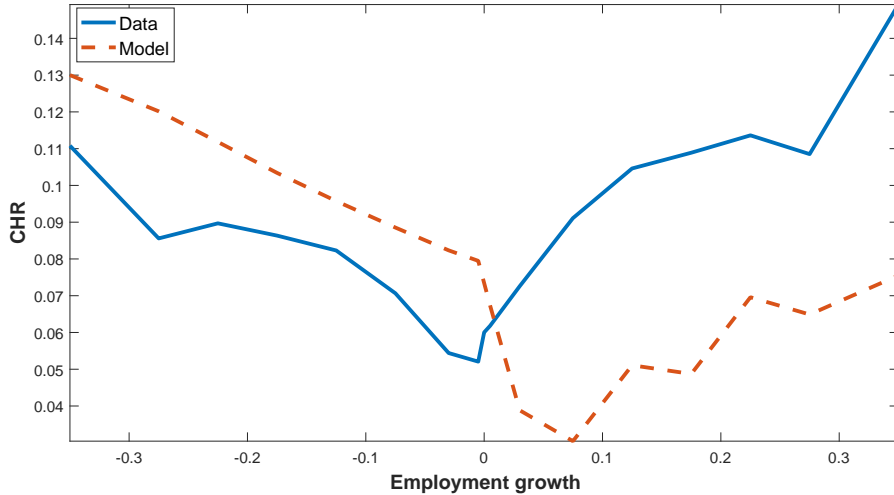
Note: The figure shows churning rates as a function of the plant-specific employment growth rates. We represent the employment growth category by its midpoint as an estimate of the average growth in that category. The blue solid line is the empirical churning rates for the West-German sample 1975-2014. The red dashed line, "model", is the churning rates from to the optimal active employment adjustment policy of plants in a model with quadratic adjustment costs and stochastic separations.

A calibrated version of the model can, thus, replicate the basic U-shape of churning rates across plant growth in the data. We set again the quarterly interest rate to 0.01, $\alpha = 0.6$, and normalize the wage to $w = 1$. The remaining parameters of the model are the level of log productivity, μ_z, the adjustment costs, ψ, the location parameter of the separation rate shocks, μ_s, and the dispersion parameter, σ_s. As before, we choose μ_z to match the average plant size in the data and obtain the other parameters by a minimum distance estimation.

Figure 4 compares the model-implied churning rate to the data. After a large separation rate shock, plants want to rehire their work force. Owing to the quadratic adjustment costs, they stretch out this rehiring process over several periods, with the most rehiring taking place in the first period. As a result, churn is larger at rapidly-growing plants, though, in contrast to the data, not as large as in rapidly-shrinking plants.

4.3 Time-to-Hire

While the separation shocks taken together with costs of employment adjustment can generate the U-shape of churning rates in the data, this model fails quantitatively for growing plants. Moreover, the estimated adjustment costs are unrealistically large (see third column in Table 3). In the model, adjusting employment by one unit costs 12 percent of average quarterly plant output. Muehlemann and Pfeifer (2016) find that average hiring costs in Germany are around two monthly wages, which translates to about 3 percent of average quarterly plant output in our model.

Figure 5: Churning Rates in a Model with Time-to-Hire and Stochastic Separations
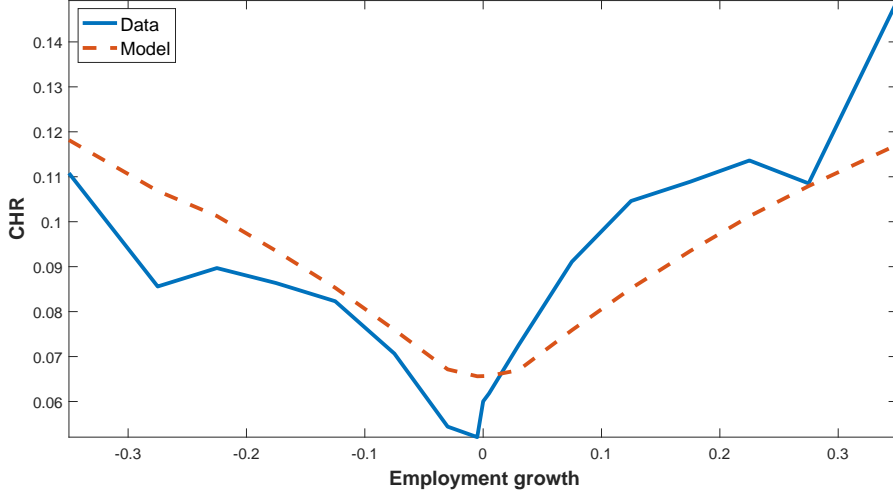
Note: The figure shows churning rates as a function of the plant-specific employment growth rates. We represent the employment growth category by its midpoint as an estimate of the average growth in that category. The blue solid line is the empirical churning rates for the West-German sample 1975-2014. The red dashed line , "model", is the churning rates from the optimal active employment adjustment policy of plants as described by model (6).

The model needs to bring down contemporaneous rehiring rates in order to increase churning at growing plants. Effectively, the model seeks to make current hiring and current separations independent with large adjustment costs.

A similar, but in our view better suited friction to achieve this decoupling of current hiring and current separations is when hiring decisions are taken based on information of the preceding period.

To be more specific, let us assume that hiring decisions take place before the separation rate shock occurs, to which plants cannot react anymore intra-period ("time-to-hire"). Plants, therefore, make mistakes in planning their employment stock. The plant chooses to actively adjust employment $\Delta^a_{E_{it}}$ to maximize:

$$\max_{\Delta^a_{E_{it}}} \Big\{ \mathbf{E}_{t-1} \{ z E_{it}^\alpha - w E_{it} \} \Big\} \tag{6}$$

$$E_{it} = (1 - s_{it})(E_{it-1} + \Delta^a_{E_{it}}), \quad \log s_{it} \sim \mathcal{N}(\mu_s, \sigma_s^2).$$

Crucially, without adjustment costs, optimal employment choices are now independent of last period's realized employment level, E_{it-1}. Therefore, plants with the largest employment growth are those who experienced large separation rates in the past. Yet, this does not mean that they necessarily have low separation rates today. On the contrary, plants with many hires tend to have many separations, $s(E_{it-1} + \Delta^a_{E_{it}})$, and thus churn.

We calibrate the model under our maintained assumptions of $r = 0.01$, $\alpha = 0.6$, and $w = 1$, calibrating the log plant productivity, μ_z, the location parameter of the separation rate shocks, μ_s, and the dispersion parameter, σ_s. Again, we use μ_z to

match the average plant size in the data and obtain the remaining parameters by a simulated minimum distance estimator. Figure 5 shows that the model is able to replicate the U-shaped pattern of the churning rate very well. Particularly, churn is largest at rapidly-growing and shrinking plants. Table 3 shows the resulting parameters. The implied uncertainty about separations is substantial. Within the 90% confidence interval, the separation rate ranges from 1 to 18 percent on a quarterly basis.[17]

To be clear, we do not mean with our analysis that productivity shocks plus adjustment costs are not important ingredients to understand plant-level labor data. Nevertheless, our analysis does suggest that another shock, stochastic separations, and another friction, time-to-hire, appear to be important to understand churning data.

5   Understanding Cyclical Churn

So far, our analysis has focused on the time average churning rate. Yet, as we have shown in Section 3, churn is particularly large during boom periods. Figure 6 shows the cyclical dynamics of the churning rate across the employment growth rate distribution. We pool the ten quarters with the lowest cyclical unemployment rate (boom) and the highest cyclical unemployment rate (recession). Table A1 in Appendix A.1 displays additional summary statistics of the cyclical dynamics of the churning rate for each individual employment growth category. Both the table and Figure 6 in Appendix A.1 show that across the employment growth distribution, churn moves counter the unemployment rate. Moreover, in absolute value, the rise during booms is similar across the distribution. The only exception are very rapidly-growing plants, but the employment share at these plants is close to acyclical.

We use our model of time-to-hire with separation shocks to estimate how these shocks must be varying over the business cycle. Table 3 shows that separation rate shocks are on average larger during booms, but their dispersion is somewhat larger during recessions. In fact, losing more than 20 percent of the workforce is more likely during recessions than during booms; but this event occurs in less than 9 percent of all cases. Put differently, the typical plant faces more separations during a boom, but shocks in the very right tail are larger during recessions. One example of such an event would be organizational restructuring that changes the desired mix of employees.

The importance of separation rate shocks for short-run plant-level employment dynamics that we estimate is particularly interesting in light of the recent debate about the role of time-varying uncertainty in business cycles (see Bloom (2014) for an overview of this literature). There, typically it is assumed that productivity shocks are more dispersed in recessions. Here, we find that dispersed separation rate shocks imply that large separation events, mass layoffs, are more likely in recessions, as should be expected.

---

[17]Key to having large churn at rapidly-growing and shrinking plants are large separation rate shocks. the log-normal distribution assumption is not critical for our results. Results are similar when we replace the log-normal with an exponential distribution.

14

Figure 6: Plant Growth and Cyclical Churn
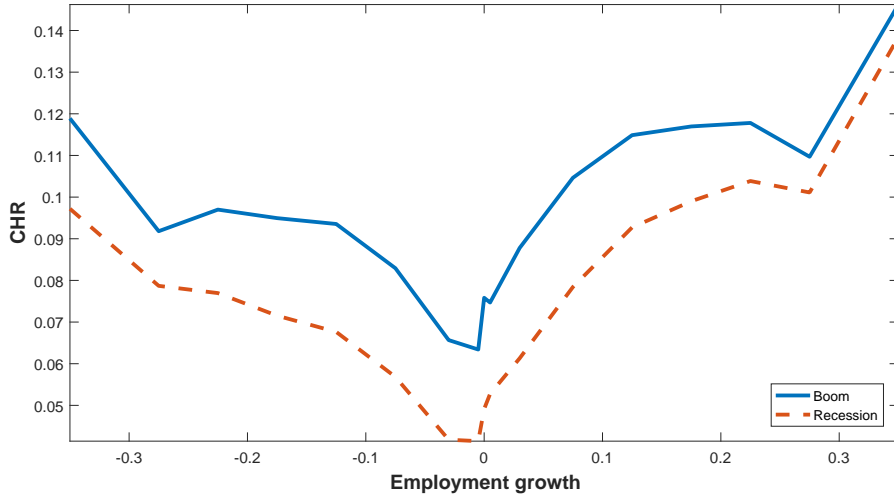
Note: The figure shows churning rates as a function of the plant-specific employment growth rates. We represent the employment growth category by its midpoint as an estimate of the average growth in that category. The blue solid line is the average churning rate in the ten quarters with the lowest HP(100,000)-filtered unemployment rate (boom). The red dashed line is the average churning rate in the highest HP(100,000)-filtered unemployment rate (recession).

5.1 Statistical Models of Procyclical Churn

Before investigating the sources of higher separations during booms in Section 5.2, we analyze, in a statistical sense, what drives cyclical movements in the aggregate churning rate. More specifically, we quantify the relative importance of two channels. First, the parallel shift of the churning rate over the cycle (Figure 6). Second, the employment growth distribution shifts over the cycle, which interacts with the U-shaped pattern of the churning rate (Figure 2). Let $chr(j)_t$ be the churning rate of the j-th employment growth category/bin. Note that

$$CHR_t = \sum_{j=1}^{J} chr(j)_t \underbrace{\frac{n_t(j)}{N_t}}_{es_t(j)}, \tag{7}$$

where $es_t(j)$ is the share of overall employment in an employment growth rate bin. In order to understand the importance of the two channels of cyclical churn, consider the following statistical models:

$$CHR_t^{d-fix} = \sum_{j=1}^{J} chr_t(j)\overline{es(j)} \tag{8}$$

$$CHR_t^{f-fix} = \sum_{j=1}^{J} \overline{chr(j)}es_t(j),$$

15

Figure 7: Contributions to Cyclical Churning
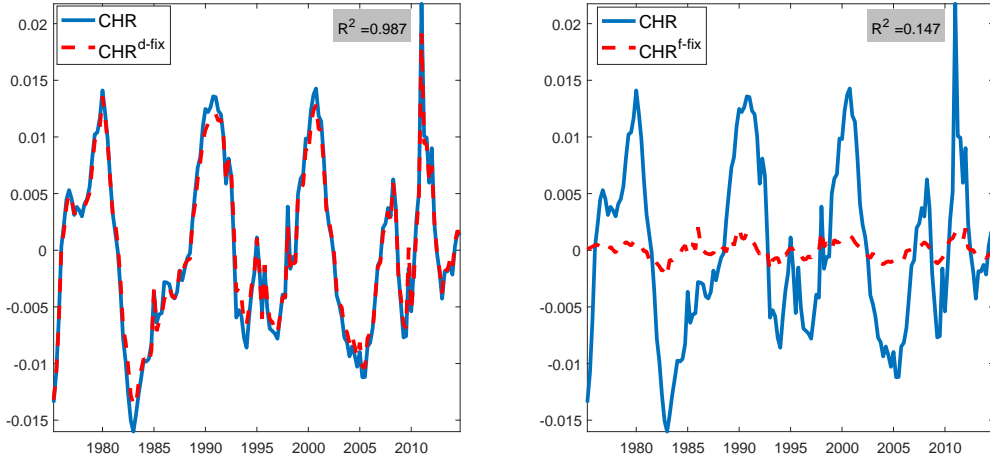
Note: The blue solid lines refer to the empirical churning rates for the West-German sample. The red dashed lines decompose the churning rate into the components described by model (10). R^2: share of the churning rate explained by rate x_t computed as $1 - (\sum(CHR_t - x_t)^2/(\sum CHR_t^2))$, where x_t is either CHR_t^{d-fix} or CHR_t^{f-fix}. All series are plotted in deviations from the HP(100,000)-filter.

where a bar denotes time-average values of employment shares and churning rates, respectively. According to the first model, churn would be procyclical because plants at all employment growth categories increase their churn during a boom (cyclical movements in $chr(j)$), and cyclical changes in the employment growth distribution do not contribute to churn. According to the second model, churn would be procyclical because the employment growth distribution shifts during booms towards employment growth categories with higher average churning rates (cyclical movements in $es_t(j)$). Given the U-shaped behavior of the churning rate, this latter channel would be potentially large, if booms were characterized by a shift away from marginally adjusting plants towards rapidly-adjusting plants.

Figure 7 displays the cyclical components of CHR_t^{d-fix} and CHR_t^{f-fix} along with the actual cyclical churning rate. The churning rate with fixed employment shares is almost identical to the aggregate churning rate. By contrast, the churning rate with fixed growth-specific churning rates explains almost none of the aggregate dynamics in the churning rate. Put differently, to understand aggregate procyclical churn, it is not necessary to jointly study the dynamics in the employment growth distribution and conditional worker flows.

The result may surprise given the findings of Davis et al. (2012). Using US data, they show, that cyclical movements in the employment growth distribution and movements in conditional worker flows are both important to understand movements in aggregate worker flow rates, a finding we replicate in Appendix A.3 for the German data. Intuitively, the difference arises because the variation in the churning for a given employment growth rate bin over the cycle (Figure 6) trumps the variation across employment growth bins (Figure 2) compared to that same relative variation for worker flow rates with their pronounced hockey-stick behavior (Figure A5 in

16

Appendix A.3). This is, however, not to say that the U-shape of the churning rates is unimportant, because, as we have shown, it identifies the underlying shocks and frictions in the first place.

5.2 Sources of Procyclical Churn

As shown in Section 3, procyclical churn is linked to rising separations during booms. To understand the reason for these separations, we differentiate separations (hires) based on their destination (source). In our data, we have information whether a separating worker is employed the next quarter at a different plant. Denote such separations/hires as job-to-job transitions, JTJ. We decompose total worker flows as those resulting from job-to-job transitions, and those resulting from non-employment transitions:

$$HR_t = JTJR_t + HR^{N-emp} \text{ and } SR_t = JTJR_t + SR^{N-emp}, \qquad (9)$$

where HR^{N-emp} denotes the hiring rate from non-employment and SR^{N-emp} denotes the separation rate into non-employment. Figures 8A and 8B shows how separations split up into flows to other employment and non-employment, and how new hires split into hires from employment and non-employment. During booms, the separation (hiring) rate to (from) employment shifts up in an almost parallel fashion over the employment growth distribution. Put differently, along the employment growth distribution, during a boom, more workers leave plants to work for another plant and plants increase their hiring from other plants. At the same time, the separation rate into non-employment and the hiring rate from non-employment show much less cyclical dynamics. On the contrary, rapidly-shrinking plants separate significantly more into non-employment, and rapidly-growing plants hire more from non-employment during recessions. Only plants that keep their employment level constant slightly increase hiring from and separations to non-employment during booms.

We can also decompose cyclical movements in the churning rate into movements in the job-to-job transition rate and the worker turnover rate through non-employment:

$$CHR_t = (HR^{N-emp} + SR^{N-emp} + 2JTJR_t) - (JCR_t + JDR_t). \qquad (10)$$

Figure 8C shows that the aggregate churning rate (divided by two) is almost identical with the job-to-job transition rate. What is more, Figure 8D shows that cyclical movements in the worker turnover rate through non-employment show no relationship with the aggregate churning rate.

Hence, equation (10) implies that worker turnover through non-employment must equal job turnover. Put differently, during booms rising job-to-job transitions do not lead to a rising job destruction rate, but to rising churn. The simultaneous rise in job creation is made possible through a rise in hiring from non-employment. Figure A3 in Appendix A.1 shows that, as a consequence, the job creation rate explains over 60% of the dynamics in the hiring rate from non-employment. Similarly, the job destruction rate explains over 80% of the dynamics in the separation rate to non-employment. Recall from Table 1 that the aggregate separation rate is procyclical, but the job destruction rate is countercyclical. By contrast, the separation rate to non-employment is countercyclical. Also recall from Table 1 that the hiring rate is

Figure 8: Churning Rates, Worker Flows and Job-to-Job Transitions

(A) Job-to-Job Transitions   (B) Non-employment
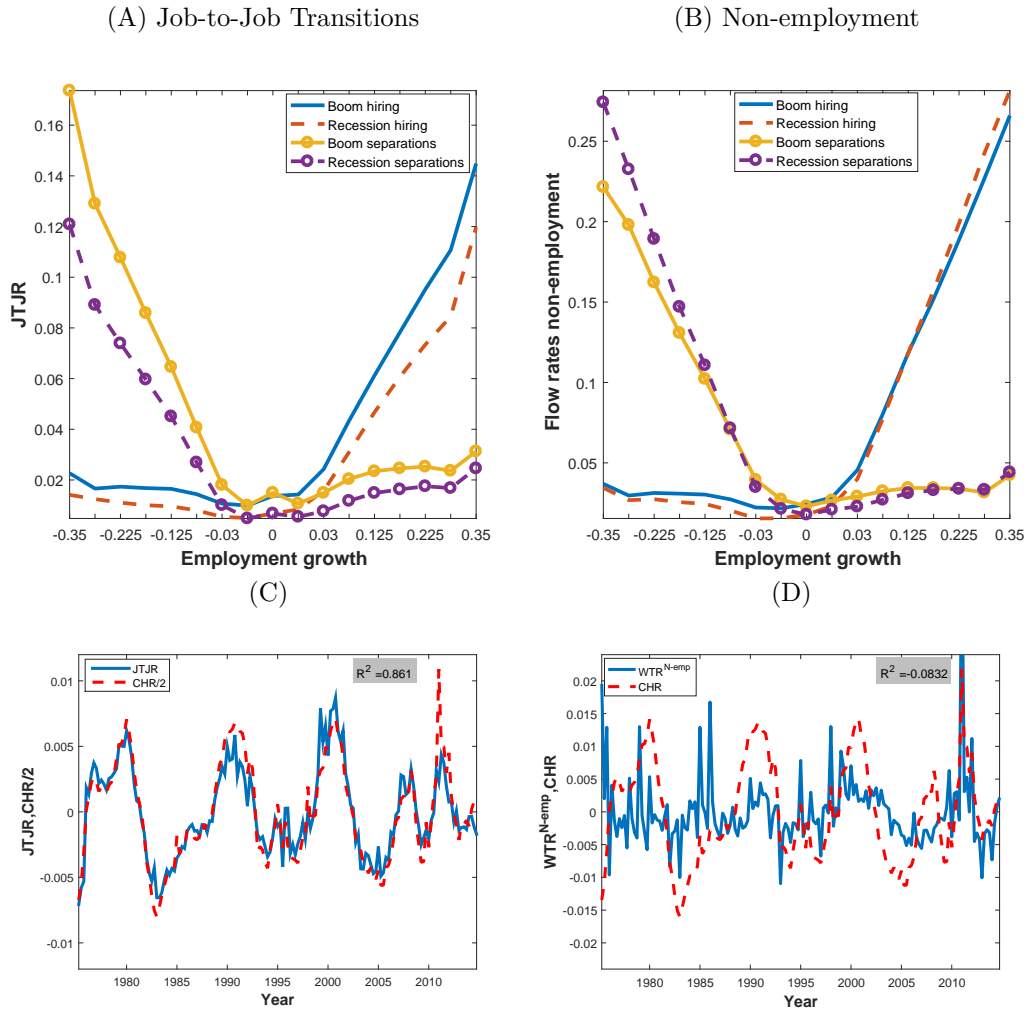
Note: Panel (A) and (B) display worker flow rates by employment growth for the West-German sample. We represent the employment growth category by its midpoint as an estimate of the average growth in that category. Panel (A): The separation (hiring) rate to (from) employment in the ten quarters with the lowest HP(100,000)-filtered unemployment rate (boom) and the highest (recession). Panel (B): The separation (hiring) rate to (from) non-employment in the ten quarters with the lowest HP(100,000)-filtered unemployment rate (boom) and the highest (recession). Panel (C) and (D) plot, respectively, the aggregate job-to-job transition rate and the aggregate worker turnover rate through non-employment (solid) against (0.5 times) the aggregate churning rate (dashed). R^2: share of churning rate explained by rate x_t computed as $1 - (\sum (CHR_t - x_t)^2/(\sum CHR_t^2))$, where x_t is either the job-to-job transition rate or the worker turnover rate through non-employment. All series are HP(100,000)-filtered.

substantially more procyclical than the job creation rate. The reason is that we can write the hiring rate as the sum of the job creation rate and the strongly procyclical job-to-job transition rate.

Taken all this together, booms are times of high job creation and high churn (not high job destruction), which means that churn-induced separations, that is,

job-to-job transitions, get ultimately replaced by some plants through hiring from non-employment; and vice versa for recessions. What is more, in terms of timing, job creation and churn both start early in a boom, but churn is more persistent and continues to increase into the maturing boom (see Table 1).

6   Existing Models of Worker Reallocation

How do our results relate to the existing theoretical and empirical literature on labor market flows? Models with a one-to-one link between worker and job flows (such as Mortensen and Pissarides (1994)) miss the large amount of procyclical churn. We show that these cyclical dynamics in worker churn result from changes in job-to-job transitions, not changes in the rate workers are churned through non-employment.

A recent literature links observable plant characteristics to cyclical job-to-job transitions and resulting plant growth. Moscarini and Postel-Vinay (2012, 2013) develop a framework where large firms grow during booms at the expense of small firms by poaching workers from small firms in a procyclical way. Haltiwanger et al. (2015) question such poaching behavior, and show that plant pay is a better predictor for cyclical employment growth patterns.

Several recent papers interpret these observable plant differences as representing underlying plant productivity. Moscarini and Postel-Vinay (2013), Schaal (2015), and Fujita and Nakajima (2016) all develop theories where during times of high production potential, vacancy posting is high, and workers flow from low- to high-productivity firms. Our findings support the idea of using job-to-job transitions as the key cyclical worker reallocation mechanism.

At the same time, our findings do not support the idea that the bulk of procyclical job-to-job transitions is driven by a common ranking (productivity) of plants for all workers and, thus, a systematic reallocation of workers from low to high productivity plants during booms. In such a set-up, highly ranked plants have low separation rates (and churning rates) on average. Moreover, during a boom, higher ranked plants grow more than during recessions. Therefore, we should observe that during booms, the separation rate (and churning rate) increases by more at shrinking plants than at growing plants. As shown in the preceding section, we find no evidence of this. We therefore view our empirical evidence as one for job-ladders, where booms foster reallocation of workers, moving to jobs they like better. Yet, our results are not in line with workers having a single common ranking across plants.

7   Conclusion

This paper studies the link between worker churn and establishment growth using a newly assembled plant-level dataset from Germany. We show that churn occurs along the entire employment growth distribution; most pronounced at rapidly-adjusting plants. Stochastic separation rate shocks that lead to planning errors by establishments do a good job explaining cross-sectional churn behavior.

These separation rate shocks become larger in booms leading to procyclical churn along the entire employment growth distribution. Rising separation rates represent

workers reallocating to other establishments, not into non-employment. Separations (and hiring) to (from) other establishments rises by a similar amount along the entire employment growth distribution.

This uniform behavior along the employment growth distribution is at odds with the idea that booms are times where workers systematically reallocate to plants which are desired by all workers. One promising way to rationalize churn across the employment growth distribution may be found in theories that stress the presence of match quality, as in Barlevy (2002), instead of productivity differences between plants. This match quality may also be time-varying because of idiosyncratic productivity shocks or changes in the optimal employment composition. As such, our paper relates to Gulyas (2016) who presents some evidence that the desired workforce composition may change when plants grow or shrink in size. This idea of an optimal workforce composition might also explain why workers do not have a common ranking of firms.

References

Abowd, J. M. and Vilhuber, L. (2011). National estimates of gross employment and job flows from the quarterly workforce indicators with demographic and industry detail. Journal of Econometrics, 161(1):82–99.

Bachmann, R. and Bayer, C. (2014). Investment dispersion and the business cycle. American Economic Review, 104(4):1392–1416.

Bachmann, R., Bayer, C., Seth, S., and Wellschmied, F. (2013). Cyclicality of job and worker flows: new data and a new set of stylized facts. IZA Discussion Paper, 7192.

Barlevy, G. (2002). The sullying effect of recessions. The Review of Economic Studies, 69(1):65–96.

Bloom, N. (2014). Fluctuations in uncertainty. Journal of Economic Perspectives, 28(2):153–176.

Burgess, S., Lane, J., and Stevens, D. (2000). Job flows, worker flows, and churning. Journal of Labor Economics, 18(3):473–502.

Davis, S. J., Faberman, J., Haltiwanger, J. C., Jarmin, R. S., and Miranda, J. (2010). Business volatility, job destruction, and unemployment. American Economic Journal: Macroeconomics, 2(2):259–287.

Davis, S. J., Faberman, J. R., and Haltiwanger, J. C. (2006). The flow approach to labor markets: new data sources and micro-macro links. Journal of Economic Perspectives, 20(3):3–26.

Davis, S. J., Faberman, J. R., and Haltiwanger, J. C. (2012). Labor market flows in the cross section and over time. Journal of Monetary Economics, 59(1):1–18.

Davis, S. J., Haltiwanger, J. C., and Schuh, S. (1996). Job Creation and Destruction. MIT Press.

Dundler, A., Stamm, M., and Adler, S. (2006). The establishment history panel BHP 1.0. FDZ Datenreport 3/2006, IAB.

Fujita, S. and Nakajima, M. (2016). Worker flows and job flows: a quantitative investigation. Review of Economic Dynamics, 22(3):1–20.

Gartner, H., Merkl, C., and Rothe, T. (2012). Sclerosis and large volatilities: two sides of the same coin. Economics Letters, 117(1):106–109.

Gulyas, A. (2016). Identifying labor market sorting with firm dynamics. Mimeo, UCLA.

Haltiwanger, J. C., Hyatt, H., and McEntarfe, E. (2015). Cyclical reallocation of workers across eemployers by firm size and firm wage. Working Paper 21235, NBER.

Hyatt, H. R. and Spletzer, J. R. (2015). The recent decline of single quarter jobs. IZA Discussion Paper, 8805.

Jung, P. and Kuhn, M. (2014). Labor market institutions and worker flows: comparing Germany and the U.S. Economic Journal, 124(581):1317–1342.

Lazear, E. P. and McCue, K. (2017). Hires and separations in equilibrium. Working Paper 23059, NBER.

Lazear, E. P. and Spletzer, J. R. (2012). Hiring, churn, and the business cycle. American Economic Review: Papers & Proceedings, 102(3):575–579.

Merkl, C. and Stüber, H. (2016). Wage cyclicalities and labor market dynamics at the establishment level: theory and evidence. FAU Discussion Papers in Economics, 12/2016.

Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. Review of Economic Studies, 61(3):397–415.

Moscarini, G. and Postel-Vinay, F. (2012). The contribution of large and small employers to job creation in times of high and low unemployment. American Economic Review, 102(6).

Moscarini, G. and Postel-Vinay, F. (2013). Stochastic search equilibrium. Review of Economic Studies, 80(1):1–44.

Muehlemann, S. and Pfeifer, H. (2016). The structure of hiring costs in germany: evidence from firm-level data. Industrial Relations, 55(2):193–218.

Schaal, E. (2015). Uncertainty and unemployment. mimeo, New York University.

Seth, S. and Stüber, H. (2017). Administrative wages and labor market flow panel (AWFP) 1975-2014. FAU Discussion Papers in Economics, 1/2017.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. American Economic Review, 95(1):25–49.

A   Appendices

A.1   Further Tables and Figures

Table A1: Dynamics of the Churning Rate

growth rate	*mean*	*SD*	*AC*(1)	*CorrU*
-2 to -0.75	3.00%	0.53%	-0.17	-0.02
-0.75 to -0.4	7.77%	0.59%	0.58	-0.66^*
-0.4 to -0.3	11.11%	1.31%	0.66	-0.70^*
-0.3 to -0.25	8.56%	0.82%	0.76	-0.68^*
-0.25 to -0.2	8.98%	0.91%	0.84	-0.71^*
-0.2 to -0.15	8.65%	1.03%	0.86	-0.69^*
-0.15 to -0.1	8.25%	1.14%	0.87	-0.66^*
-0.1 to -0.05	7.12%	1.05%	0.93	-0.68^*
-0.05 to -0.01	5.49%	0.83%	0.92	-0.72^*
-0.01 to 0	5.22%	0.63%	0.85	-0.75^*
0	6.06%	0.66%	0.90	-0.84^*
0 to 0.01	6.15%	0.59%	0.76	-0.79^*
0.01 to 0.05	7.20%	0.69%	0.88	-0.83^*
0.05 to 0.1	9.03%	0.82%	0.83	-0.75^*
0.1 to 0.15	10.39%	0.89%	0.76	-0.65^*
0.15 to 0.2	10.84%	0.90%	0.77	-0.55^*
0.2 to 0.25	11.33%	0.89%	0.66	-0.52^*
0.25 to 0.3	10.83%	0.84%	0.55	-0.40^*
0.3 to 0.4	14.88%	1.43%	0.34	-0.31^*
0.4 to 0.75	10.01%	0.71%	0.44	-0.39^*
0.75 to 2	3.88%	0.34%	0.32	-0.23

Note: The table displays the HP(100,000)-filtered churning rate over the employment growth distribution. Mean: non-filtered time-average churning rate, SD: standard deviation, $AC(1)$ autocorrelation coefficient, $CorrU$: correlation with unemployment. A * indicates significance at the 5% level obtained by non-parametric block-bootstrapping with a block length of 20.

Figure A1: Empl. Growth Distribution   Figure A2: Aggregate $CHR_t[-0.4, 0.4]$
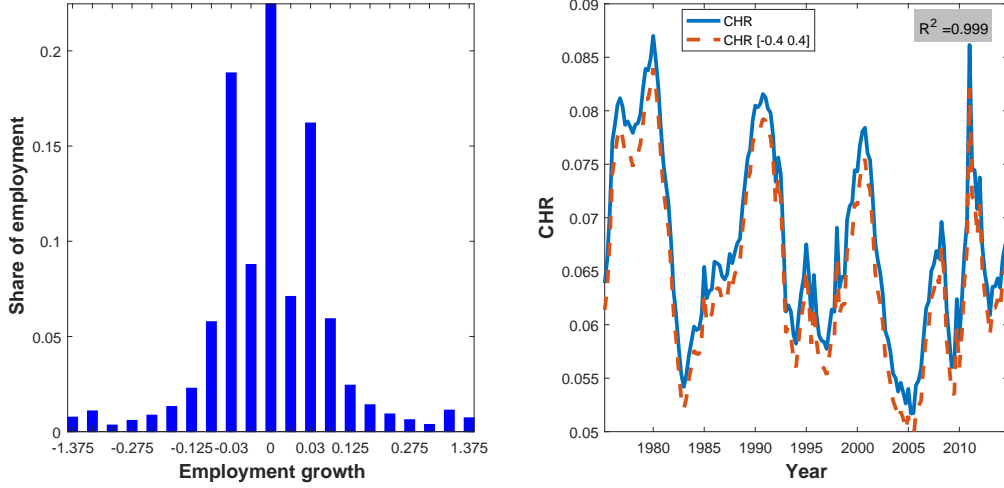

Note: Left figure: The time-average employment share for West-Germany (1975-2014) for each employment growth category. We represent the employment growth category by its midpoint as an estimate of the average growth in that category. Right figure: The blue straight line is the churning rate. The red dashed line is the churning rate resulting from churn occurring in employment growth categories $[-0.4, 0.4]$. R^2: share of the churning rate explained by the churning rate from the plant-growth interval $[-0.4, 0.4]$ computed as $1 - (\sum (CHR_t - CHR_t[-0.4, 0.4])^2 / (\sum CHR_t^2))$.

Figure A3: Aggregate Flows from Non-Employment and Aggregate Job Flows

(A) HR without JTJ   (B) SR without JTJ
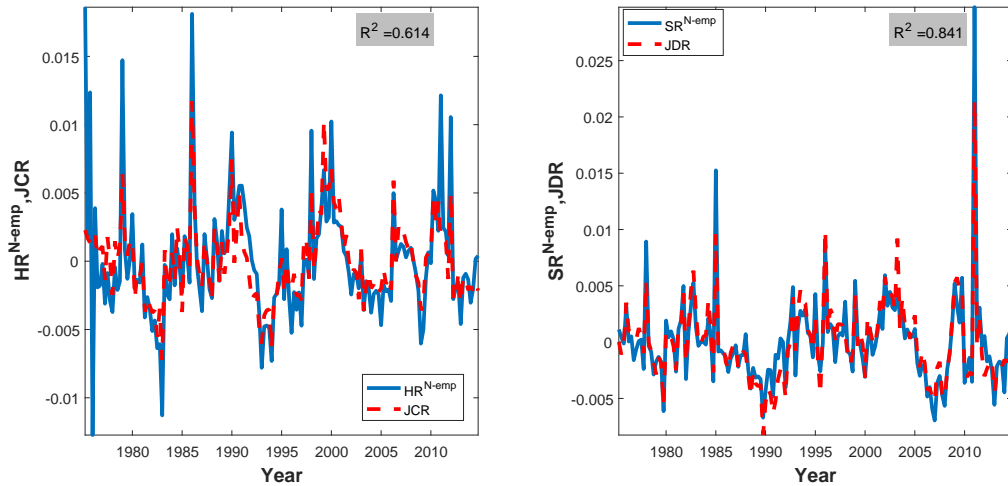
Note: The blue solid lines refer to the empirical hiring rate from non-employment (left) and separation rate to non-employment (right) in West-Germany. The red dashed line is the corresponding job creation rate (left) and job destruction rate (right). R^2: share of the hiring rate from non-employment explained by the job creation rate computed as $1 - (\sum (HR_t^{N-emp} - JCR)^2 / (\sum HR_t^{N-emp^2}))$; analogously for the separation rate to non-employment and the job destruction rate. All series are HP(100,000)-filtered.

24

A.2 Relationship to US Data

For our comparison with the US, we obtain seasonally adjusted US quarterly job flows from the Business Employment Dynamics (*BED*) data for the period of 1992–2014. *BED* contains information on the universe of US establishments, excluding household employment, government employees, the self-employed, and small-farm workers.[18] The *BED* data does not contain information on worker flows. Therefore, we obtain seasonally adjusted worker flows from *JOLTS* for the years 2001–2014. *JOLTS* samples every month 16,000 establishments from the universe of US establishments with the exception of agriculture and private households. We aggregate the monthly flows to the quarterly frequency.

Figure A4 compares German job and worker flows to those in the US. Job and worker flows are substantially larger in the US than in Germany. Average quarterly job flows in Germany are 0.036, compared to 0.071 in the US. Similarly, the average worker flow rate in Germany is 0.070, compared to 0.118 in the US. The second major difference between the countries is that job flows show a negative trend in the US over time, but there is no such trend in Germany. Davis et al. (2010) attribute this trend to declining business dynamism in the US. Hyatt and Spletzer (2015) show that about half of the decrease can be explained by a decrease in the amount of jobs lasting less than a quarter. Such short-lasting jobs have always been rare in Germany; where they exist (e.g., internships, student jobs, etc.), they are not counted as regular workers and hence do not enter our data.

Table A2 displays the cyclical properties of job flow rates in the US. The cyclical volatility of the job-creation rate, *JCR*, and the the job-destruction rate, *JDR*, are similar in the two countries. Remember that both flow rates are substantially lower in Germany. As a result, these flow rates are more than 50 percent more volatile in Germany when using log deviations: the *JCR* and *JDR* are, respectively, 2.5 and 3.7 times more volatile than output in the US. For Germany we find ratios of 4.3 and 5.4. This reflects that the Shimer (2005) puzzle is even more evident in Germany compared to the US (see Gartner et al. (2012) and Jung and Kuhn (2014)).

Table A3 computes the correlations between job and worker flows in US data. As in the German data, the job creation and destruction rate are negatively correlated, and the hiring and separation rate are positively correlated. Moreover, the job creation rate is positively correlated with the hiring rate, and the job destruction rate is positively correlated with the separation rate.

---

[18]The two concepts of establishments are not quite the same. In the US, an establishment is a single physical location where business is conducted, or where services or industrial operations are performed. In our dataset, each firms' production unit located in a county (Kreis) receives an establishment identifier based on an industry classification. When each production unit within a county has a different industry classification, or a firms' production unit are located in different counties, the two definitions coincide. When a firm has more than one production unit within the same county that are classified by the same industry, they may receive the same establishment identifier. The employer may decide, however, to have different identifiers assigned (see Dundler et al. (2006)).

Figure A4: Job and Worker Flows in the US and Germany
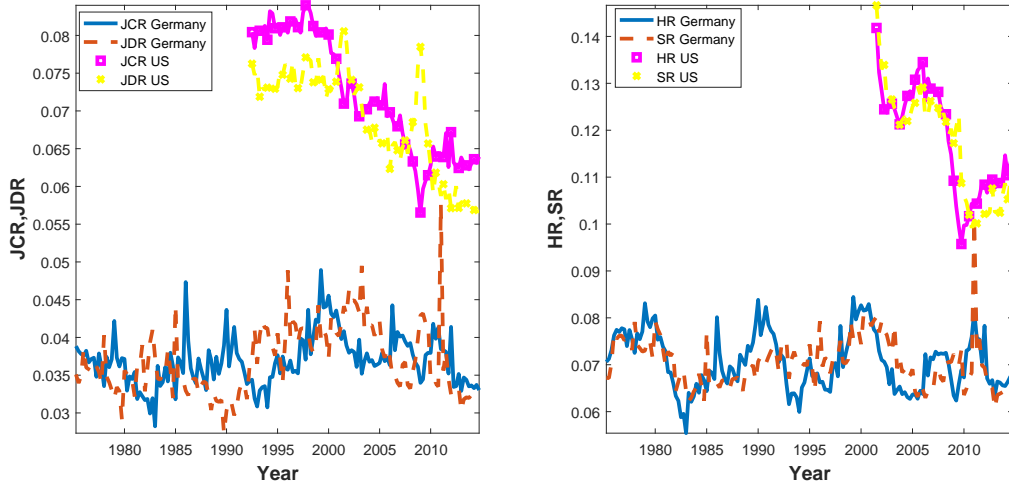
Note: The figure displays job and worker flows in Germany and the US. *JCR*: job creation rate, *JDR*: job destruction rate, *HR*: hiring rate, *SR*: separation rate.

Table A2: Job and Worker Flows in the US and Germany

		Mean	*SD*	*AC*(1)	$j = -2$	$j = 0$	$j = +2$
					Correlation with U_{t+j}		
JCR	GER	3.69%	0.29%	0.52	0.19*	−0.04	−0.28*
JCR	US	7.16%	0.27%	0.81	−0.16	−0.45*	−0.63*
JDR	GER	3.69%	0.36%	0.40	−0.02	0.15	0.29*
JDR	US	6.84%	0.34%	0.81	−0.32*	0.02	0.30*
HR	GER	7.06%	0.57%	0.82	−0.26*	−0.53*	−0.27*
HR	US	11.82%	0.82%	0.93	−0.63*	−0.87*	−0.94*
SR	GER	7.06%	0.47%	0.47	−0.46*	−0.51*	−0.48*
SR	US	11.68%	0.67%	0.87	−0.91*	−0.86*	−0.68*

Note: The table displays the properties of the HP(100,000)-filtered job and worker flow rates. *SD*: standard deviation, *AC*(1): first-order auto correlation. A * indicates significance at the 5% level obtained by non-parametric block-bootstrapping with a block length of 20.

Table A3: Correlations of Job and Worker Flows in the US

	JCR	JDR	HR	SR
JCR	1.00			
JDR	-0.75^*	1.00		
HR	0.64^*	-0.44^*	1.00	
SR	0.08	0.25	0.71^*	1.00

Note: The table displays correlation coefficients of HP(100,000)-filtered job and worker flow rates. A * indicates significance at the 5% level obtained by non-parametric block-bootstrapping with a block length of 20.

A.3 Relationship with Davis et al. (2012)

Figure A5: Worker Flows and Employment Growth
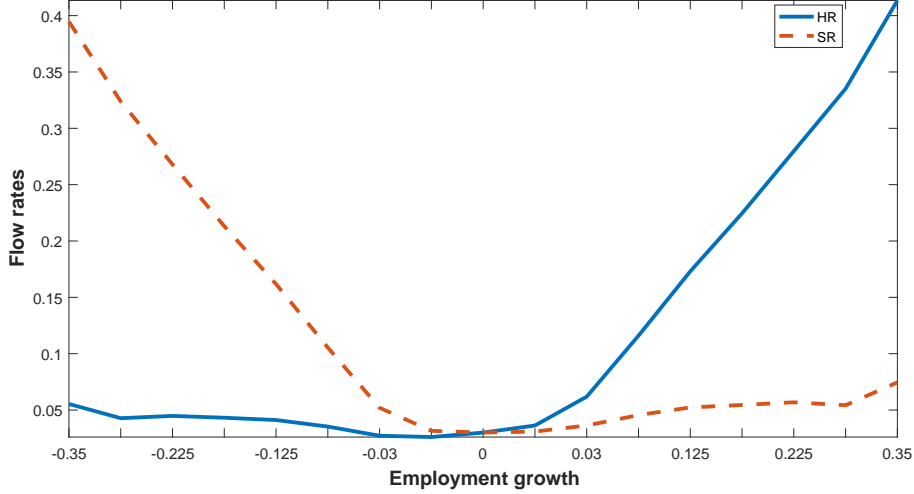
Note: The figure displays time averaged flow rates by employment growth for the West-German sample 1975-2014. We represent the employment growth category by its midpoint as an estimate of the average growth in that category. The blue solid line is the hiring rate, the red dashed line the separation rate, and the yellow dotted line the churning rate.

This appendix shows that our German data lead to very similar findings as in Davis et al. (2012) as regards *worker flows*. The most widely used framework to understand worker flows are variants of the Mortensen and Pissarides (1994) model. In this framework, all worker flows result from job flows, a characteristic which Davis et al. (2012) label the "iron link" between job and worker flows. Figure A5 shows the relationship between job and worker flows. The accession and separation rates are positive along the entire employment growth distribution. The accession rate grows close to linearly with positive employment growth, and the separation rate grows close to linearly with negative employment growth, a relationship Davis et al. (2012) call hockey-stick behavior. Furthermore, similar to Davis et al. (2012), we quantify the importance of shifts in the employment growth distribution for worker flows using the following statistical model:

$$HR_t^{f-fix} = \sum_{j=1}^{J} \overline{hr(j)} es_t(j) \tag{A.1}$$

$$SR_t^{f-fix} = \sum_{j=1}^{J} \overline{sr(j)} es_t(j),$$

where a bar denotes time-averaged values and $es_t(j)$ is the share of overall employment in an employment growth rate bin. According to this model, given plant-level employment growth, worker flows do not vary over time. Therefore, cyclical changes

Figure A6: Fixed Worker Flow Rates Over the Cycle
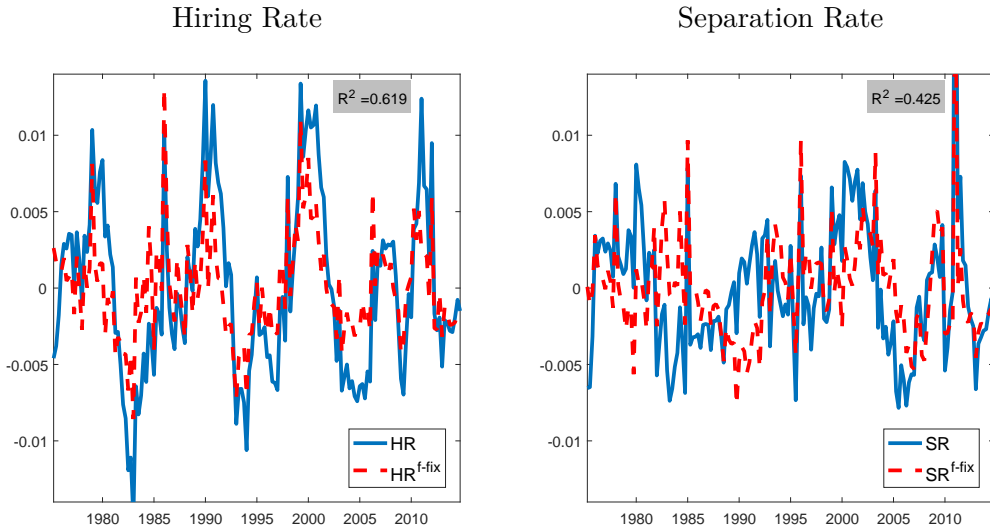
Hiring Rate   Separation Rate

Note: The blue solid lines is the empirical hiring and separation rate for the West-German sample. The red dashed lines display the corresponding synthetic series described by model (A.1). R^2: share of hiring rate explained by rate HR_t^{f-fix} computed as $1-(\sum(HR_t-HR_t^{f-fix})^2/(\sum HR_t^2))$; analogously for separation rates. All series are plotted as deviations from the HP(100,000)-filter.

in worker flow rates result from cyclical shifts in the employment growth distribution only. The specification is more general than the pure "iron link", because it allows shrinking establishments to have positive hires and growing establishments to have positive separations. Moreover, we allow the series to have a time varying trend component.

Figure A6 plots the synthetic flow rates from our statistical model against the true hires and separation rate. Job flows explain a substantial fraction of cyclical worker flows. Movements of the employment growth distribution capture all major movements in the hiring rate. In a statistical sense, the synthetic series explains 64% of the movements in the hiring rate. For the separation rate, the synthetic series with fixed conditional flow rates explains 43%.

We also consider a second model where worker flows fluctuate because for a given amount of employment adjustment, at least some plants increase their worker flows in booms relative to recessions:

$$HR_t^{d-fix} = \sum_{j=1}^{J} hr_t(j)\overline{es(j)} \tag{A.2}$$

$$SR_t^{d-fix} = \sum_{j=1}^{J} sr_t(j)\overline{es(j)}.$$

Figure A7 displays the resulting synthetic series from this exercise. The series are quite a good fit for the realized rates. The synthetic series explains 65% of the hiring rate. The hiring rate is not sufficiently volatile, but the timing of periods with high

29

Figure A7: Fixed Employment-Growth Distribution Over the Cycle
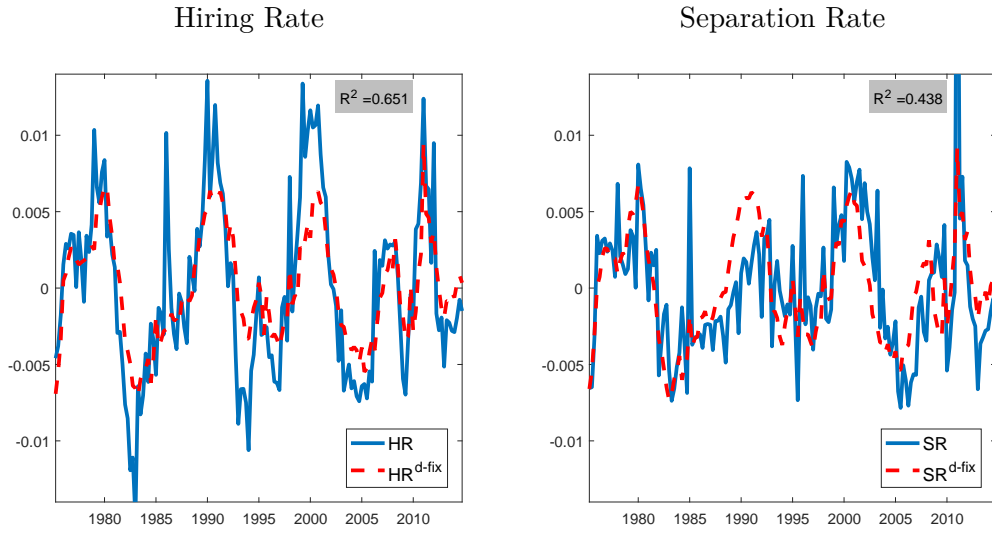
Hiring Rate   Separation Rate

Note: The blue solid lines is the empirical hiring and separation rate for the West-German sample. The red dashed lines display the corresponding synthetic series described by model (A.2). R^2: share of the hiring rate explained by rate HR_t^{d-fix} computed as $1 - (\sum(HR_t - HR_t^{d-fix})^2/(\sum HR_t^2))$; analogously for the separation rate. All series are plotted as deviations from the HP(100,000)-filter.

and low rates is almost identical. The statistical model explains 44% of the separation rate. Taken together, in a statistical sense, the model with the fixed employment growth distribution and the model with the fixed conditional worker flows explain similar amounts of the volatility in aggregate worker flow rates. Particularly for the separation rate, the model with the fixed employment growth distribution explains mainly major changes in the rate, and the model with fixed conditional flows explains quarter to quarter spikes.