Online Appendix to “Worker Churn in the Cross Section and Over Time: New Evidence from Germany”

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

May 18, 2020

A Relationship to U.S. Job and Worker Flows

A.1 Aggregate Dynamics

Figure A1: Job and Worker Flows in the United States and Germany

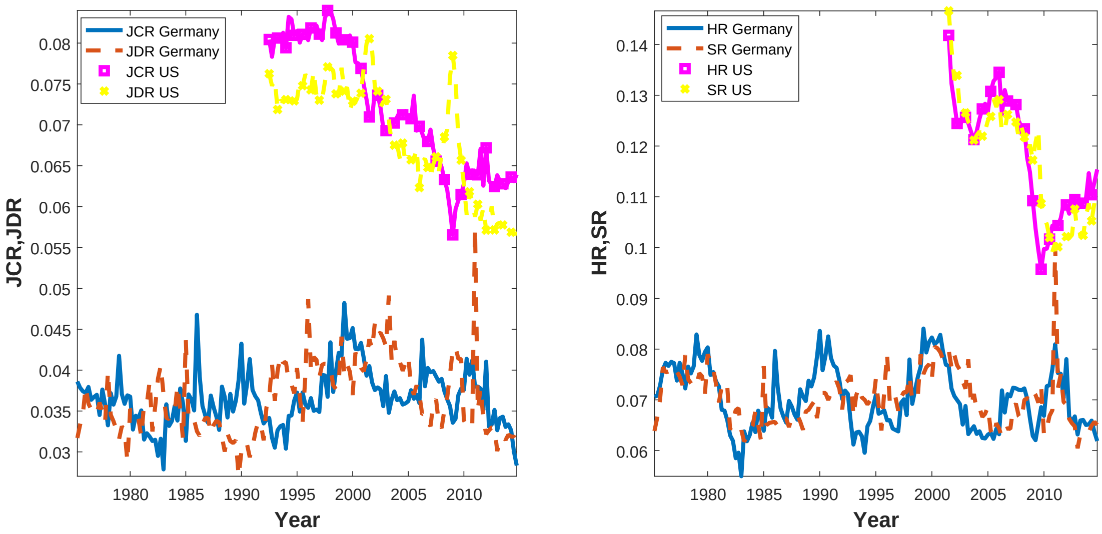

Note: This figure displays job and worker flows in West Germany and the United States. JCR: job creation rate, JDR: job destruction rate, HR: hiring rate, SR: separation rate.

For our comparison with the United States, we obtain non-seasonally-adjusted U.S. quarterly job flow rates from the Business Employment Dynamics (*BED*) data for the period of 1992-2014 and apply the same X-12 ARIMA CENSUS procedure as for German data. *BED* contains information on the universe of U.S. establishments, excluding household employment, government employees, the self-employed, and small-farm workers.[1] The *BED* data does not contain information on worker flows. Therefore, we construct seasonally adjusted worker flows from *JOLTS* for the years 2001-2014. Specifically, we use the monthly non-seasonally-adjusted flow rates and aggregate flows to back out monthly employment levels, E_t. By summing over the monthly aggregate flows to quarterly flows as the numerator and using the average of the first-month and third-month (of each quarter) employment levels to approximate N_t in the denominator, we construct quarterly flow rates that we, finally, seasonally adjust with the X-12 ARIMA CENSUS procedure. *JOLTS* samples every month 16,000 establishments from the universe of U.S. establishments with the exception of agriculture and private households.

Figure A1 compares German job and worker flows to those in the United States. Job and worker flows are substantially larger in the United States than in Germany. Average quarterly job flows in Germany are 3.69%, compared to the 7.16% job creation rate (6.84% job destruction rate) in the United States. Similarly, the average worker flow rate in Germany is 7.06%, compared to an 11.82% hiring rate (11.68% separation rate) in the United States.

[1]The two concepts of establishments are not quite the same. In the United States, an establishment is a single physical location where the business is conducted, or where services or industrial operations are performed. In our data set, 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 units 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)).

Table A1: Job and Worker Flows in the United States and Germany

					Correlation with U_{t+j}		
		Mean	*Std*	*AC*(1)	$j=-2$	0	+2
JCR	GER	3.65%	0.30%	0.54	0.18^{**}	-0.05	-0.29^{***}
JCR	U.S.	7.16%	0.27%	0.81	-0.16	-0.45^{***}	-0.63^{***}
JDR	GER	3.65%	0.36%	0.40	-0.03	0.14	0.28^{**}
JDR	U.S.	6.84%	0.34%	0.81	-0.32^{***}	0.02	0.30^{*}
HR	GER	7.02%	0.58%	0.81	-0.26^{***}	-0.53^{***}	-0.72^{***}
HR	U.S.	11.82%	0.82%	0.93	-0.63^{***}	-0.87^{***}	-0.94^{***}
SR	GER	7.02%	0.47%	0.47	-0.46^{***}	-0.51^{***}	-0.48^{***}
SR	U.S.	11.68%	0.67%	0.87	-0.91^{***}	-0.86^{***}	-0.68^{***}

Note: This table displays the properties of a number of aggregate labor market flow rates. The third column, *Mean*, displays the time-averaged rates. The subsequent columns display moments of the HP(100,000)-filtered rates. *JCR*: job creation rate, *JDR*: job destruction rate, *HR*: hiring rate, *SR*: separation rate. *Std*: standard deviation, *AC*(1): first-order auto correlation. Stars indicate significance at the 1%, 5% and 10% level obtained by non-parametric block-bootstrapping with a block length of 20.

The second major difference between the countries is that job flows show a negative trend in the United States over time, but there is no such trend in Germany.[2] Davis et al. (2010) attribute this trend to declining business dynamism in the United States. Hyatt and Spletzer (2017) show that about half of the decrease can be explained by a decrease in the number 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 A1 displays the cyclical properties of job flow rates in the United States. The cyclical volatility of the job creation rate, *JCR*, and the job destruction rate, *JDR*, are similar in the two countries. Recall 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 United States. In Germany, these ratios are, respectively, 4.3 and 5.4. This means that the Shimer (2005) puzzle is even more evident in Germany compared to the United States (see Gartner et al. (2012) and Jung and Kuhn (2014)).

Table A2 computes the correlations between job and worker flows, respectively, for the two countries. In both countries, the job creation and destruction rates are negatively correlated. The job creation rates and hiring rates, and the job destruction rates and the separation rates are positively correlated. Nonetheless, the hiring and separation rates are also positively correlated.

[2] There is such a negative trend in former East-Germany. The initially high flows after re-unification might arise from low productivity plants exiting the market, yet, in light of our results, may also represent a high level of initial mismatch.

Table A2: Correlations of Job and Worker Flows in the United States and Germany

		JCR	JDR	HR	SR
JCR	U.S.	1.00			
	Germany	1.00			
JDR	U.S.	-0.75^{***}	1.00		
	Germany	-0.32^{***}	1.00		
HR	U.S.	0.64^{***}	-0.44^{***}	1.00	
	Germany	0.81^{***}	-0.29^{***}	1.00	
SR	U.S.	0.08	0.25^{**}	0.71^{***}	1.00
	Germany	0.11	0.61^{***}	0.49^{***}	1.00
CHR	U.S.	-	-	-	-
	Germany	0.45^{***}	-0.18	0.89^{***}	0.67^{***}

Note: This table displays correlation coefficients of HP(100,000)-filtered job and worker flow rates. JCR: job creation rate, JDR: job destruction rate, HR: hiring rate, SR: separation rate, CHR: churn rate. Std: standard deviation, $AC(1)$: first-order auto correlation. Stars indicate significance at the 1%, 5% and 10% level obtained by non-parametric block-bootstrapping with a block length of 20.

A.2 Statistical Models of Worker Flows and Churn

This section goes beyond simple correlations, and analyzes, in a statistical sense, the drivers of the cyclical movements in aggregate worker flows and churn. The German data reveal very similar patterns to those in Davis et al. (2012) as regards worker flows. We also extend their analysis to churn.

Figure A2 shows the relationship between job and worker flows at the plant level. The hiring and separation rates are positive across the entire plant-level employment growth distribution. The hiring 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. Similar to Davis et al. (2012), we quantify the importance of quarter-to-quarter changes 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 $J = 21$ are employment growth categories/bins,[3] $es_t(j) = \frac{N_t(j)}{N_t}$ is the share of overall employment in an employment growth rate bin, and a bar denotes time-averaged values.

[3] These bins are the intervals: -2 to -0.75, -0.75 to -0.4, -0.4 to -0.3, -0.3 to -0.25, -0.25 to -0.2, -0.15 to -0.1, -0.1 to -0.05, -0.05 to -0.01, -0.01 to 0, 0, and symmetrically for positive employment growth. We allow each employment growth category to have its own seasonal component. To derive the aggregate series for West Germany, we finally sum over the seasonally adjusted series for all employment growth categories.

Figure A2: Worker Flows and Employment Growth

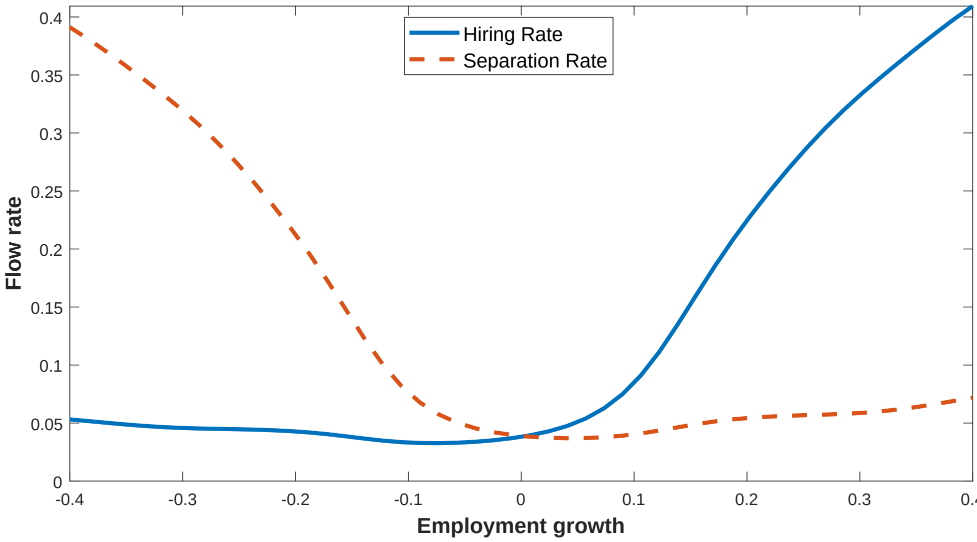

Note: This figure displays average worker flow rates of a plant as a function of its employment growth rate, estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The blue solid line is the hiring rate, the red dashed line the separation rate. West German plants with a quarterly frequency, 1975Q1-2014Q4.

Figure A3: Fixed Worker Flow Rates over the Cycle

Hiring Rate

Separation Rate

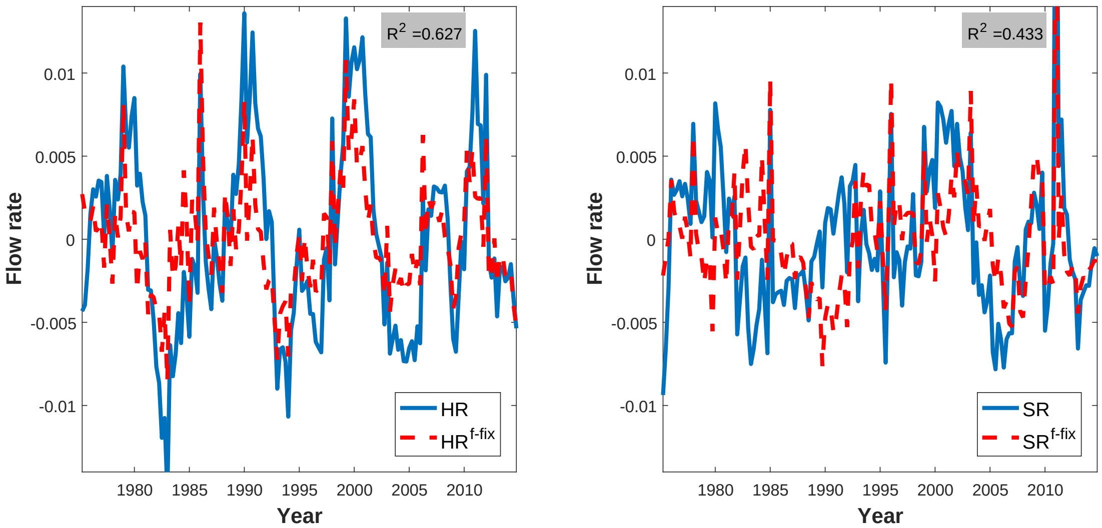

Note: This figure displays the hiring and separation rates in the data together with counterfactual hiring and separation rates where the worker flow rates by employment growth category are fixed over time. The blue solid lines are the empirical hiring and separation rate, respectively. The red dashed 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. West German plants with a quarterly frequency, 1975Q1-2014Q4.

According to this model, given plant-level employment growth, worker flows do not vary over time. Therefore, cyclical changes in worker flow rates result from cyclical shifts in the employment growth distribution only. The specification is more general than a one-to-one link between job and worker flows because it allows shrinking establishments to have positive hires and growing establishments to have positive separations. Moreover, it allows the series to have a time-varying trend component.

Figure A3 plots the synthetic flow rates from our statistical model against the true hiring and separation rates. 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 61% 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 change their worker flows from quarter to quarter:

$$\begin{aligned} HR_t^{d-fix} &= \sum_{j=1}^{J} hr_t(j)\overline{es(j)} \\ SR_t^{d-fix} &= \sum_{j=1}^{J} sr_t(j)\overline{es(j)}. \end{aligned} \tag{A.2}$$

Figure A4: Fixed Employment Growth Distribution over the Cycle

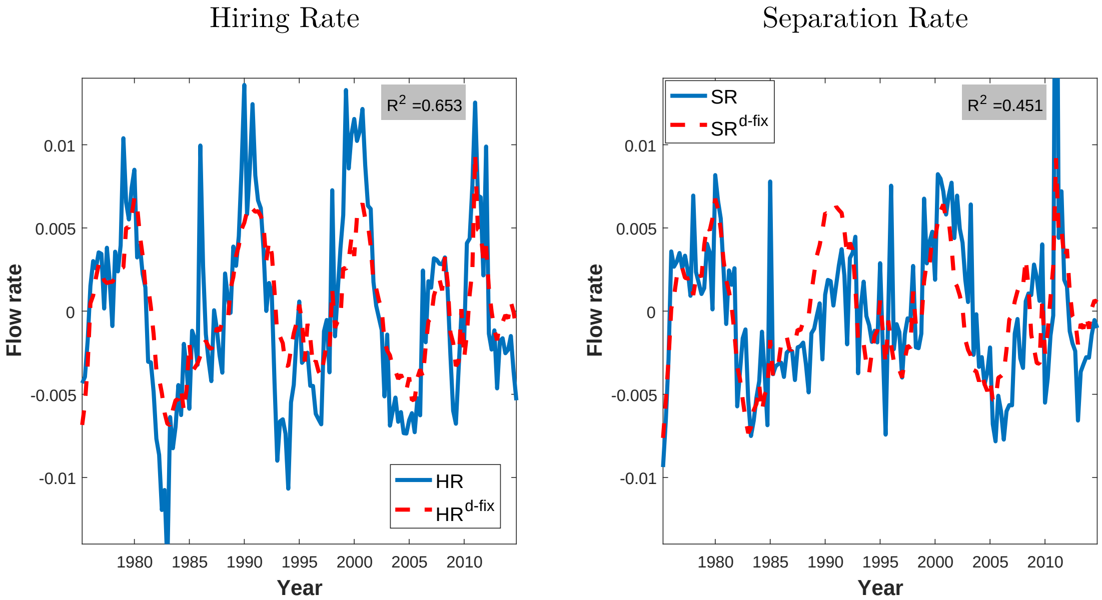

Note: This figure displays the hiring and separation rates in the data together with counterfactual hiring and separation rates with a fixed employment growth distribution over time. The blue solid lines are the empirical hiring and separation rate, respectively. The red dashed lines display the corresponding 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. West German plants with a quarterly frequency, 1975Q1-2014Q4.

Figure A4 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 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 roughly 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.

We now use the same statistical models apply them to worker churn. This allows us to understand the relative importance of the parallel shift of the churn rate over the cycle and shifts of the employment growth distribution over the cycle in explaining procyclical churn. Let $chr(j)_t$ be the churn rate of the j-th employment growth category/bin. Note that

$$CHR_t = \sum_{j=1}^{J} chr(j)_t es_t(j). \tag{A.3}$$

In order to understand the importance of the two channels of cyclical churn, consider the following statistical models:

$$\begin{aligned} CHR_t^{d-fix} &= \sum_{j=1}^{J} chr_t(j)\overline{es(j)} \\ CHR_t^{f-fix} &= \sum_{j=1}^{J} \overline{chr(j)} es_t(j). \end{aligned} \tag{A.4}$$

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 $chr(j)$) while 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 churn rates (cyclical movements in $es_t(j)$). Given the V-shaped nexus of the churn rate with employment growth, this latter channel would be potentially large if booms were characterized by a shift away from marginally adjusting plants towards rapidly adjusting plants.

Figure A5 displays the cyclical components of CHR_t^{d-fix} and CHR_t^{f-fix} along with the actual cyclical churn rate. The churn rate with fixed employment shares is almost identical to the aggregate churn rate. By contrast, the churn rate with fixed growth-specific churn rates explains almost none of the aggregate dynamics in the churn rate.

The result may surprise given the above findings that cyclical movements in the employment growth distribution are important for understanding movements in aggregate worker flow rates. Intuitively, the difference arises because the variation in the churn rate across employment growth is relatively small around zero employment growth, the area most sensitive to the business cycle. In contrast, worker flows follow a hockey-stick behavior (Figure A2) with a large change around zero employment growth.

Figure A5: Contributions to Cyclical Churn

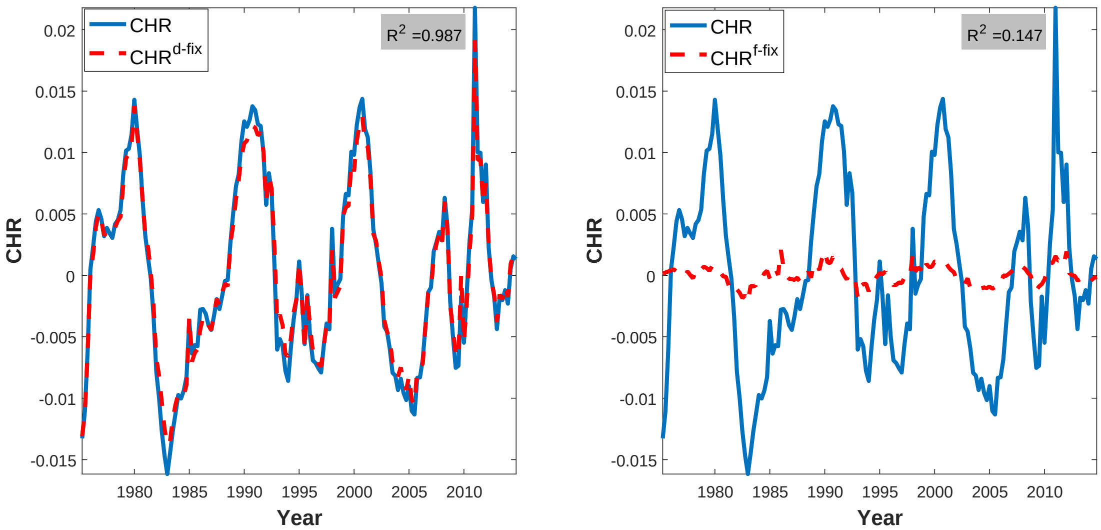

Note: This figure displays the churn rate in the data together with two counterfactual churn rates, where, respectively, the employment growth distribution is fixed over time and where the churn rates by employment growth category are fixed over time. The blue solid lines refer to the empirical churn rates. The red dashed lines decompose the churn rate into the components described by model (A.4). R^2: share of the churn 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. West German plants with a quarterly frequency, 1975Q1-2014Q4.

B Additional Tables and Figures

Figure A6: Aggregate Job and Worker Flows and the Churn Rate

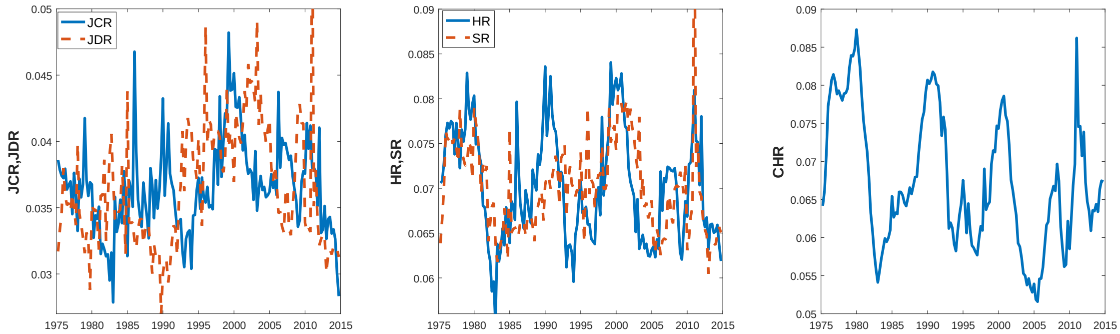

Note: The left panel displays aggregate job flows. JCR: job creation rate, JDR: job destruction rate. The center panel displays aggregate worker flows. HR: hiring rate, SR: separation rate. The right panel displays the aggregate churn rate, CHR. All rates are seasonally adjusted. West German plants with a quarterly frequency, 1975Q1-2014Q4.

Figure A6 displays seasonally-adjusted aggregate job and worker flow rates in Germany from 1975-2014. The left panel shows job flow rates, the center panel plots the worker flow rates, and the right panel displays the resulting churn rate.

Figure A7 displays the relationship between plant-level employment growth and churn across different ten-year samples. It shows that the V-shaped relationship between plant-level employment growth and churn is stable from the 1970s to the 2010s.

The model we present in Section 4 describes plants that, resulting from transitory separation rate shocks, fluctuate around their long-run optimal employment level. Section 3 studies the relationship between plant-level churn and employment growth without differentiating long-run employment growth trends from transitory deviations from these trends. Figure A8 shows that the relationship between plant-level churn and employment growth

Figure A7: Churn and Employment Growth over Four Ten-Year Samples

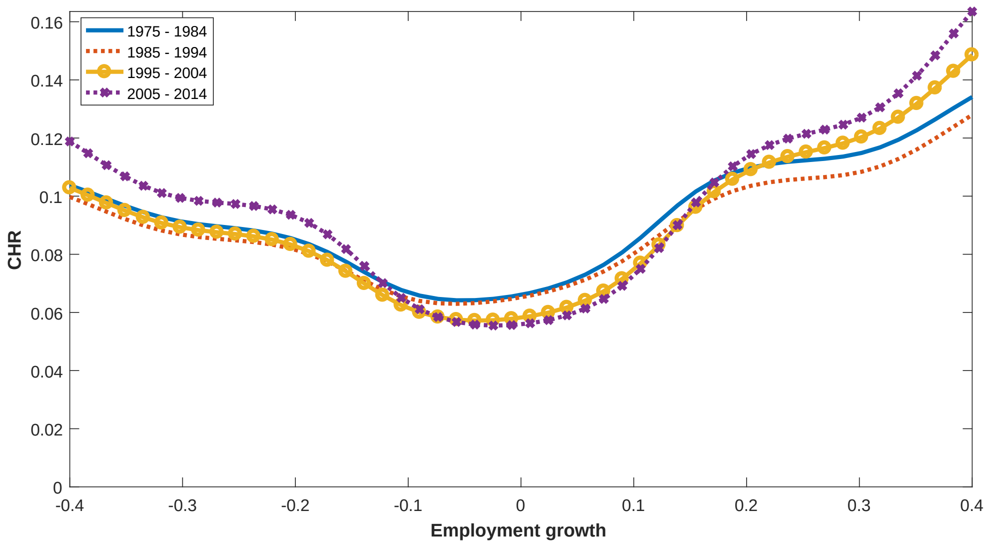

Note: This figure displays the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The estimate is done for four different ten-year samples. West German plants only, 1975Q1-2014Q4.

Figure A8: Churn and Employment Growth within Plants

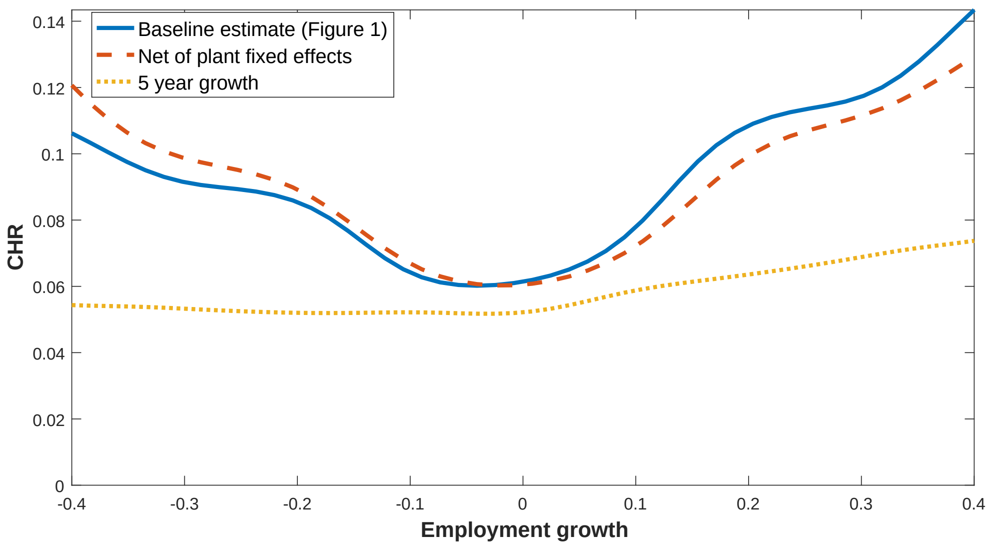

Note: This figure displays the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The blue line uses quarterly employment growth rates. The red dashed line controls for plant fixed effects in quarterly employment growth rates. The yellow dotted line uses the 5-year employment growth rate prior to the date of churn. West German plants only, 1975Q1-2014Q4.

changes little when controlling for long-run employment growth trends. The red dashed line shows the relationship when we control for plant fixed effects in plants' employment growth rates. The relationship becomes a bit more symmetric but the differences relative to the baseline specification are small. The figure also displays the plant-level relationship between employment growth and (quarterly) churn when we consider long-run employment growth, i.e., the 5-year average employment growth rate prior to the date of churn. In this case, the V-shaped relationship is no longer visible.

Section 3 studies the relationship between plant-level churn and worker characteristics by comparing workers with different education levels, tasks, and ages. Wages provide us with another, now continuous, measure that reflects worker characteristics. Figure A9 plots a non-parametric estimate of the relationship between plant-level churn rates and the absolute difference in the plant level average earnings of new hires and separating workers. The slope of this relationship is negative, i.e., plants with high churn rates have smaller differences in the earnings of incoming and outgoing workers. This negative relationship is the opposite of what one would expect if churn was driven by plant reorganization where plants substitute workers of different types that command different wage levels for another.

Next, Section 3 also shows that churn is higher at plants with many low-tenured workers and that these plants display a more pronounced V-shaped relationship between plant-level churn and employment growth. At a deeper level of analysis, this pattern may be partly driven by plant and worker interaction. Young plants have, by definition, many low-tenured workers and these plants may have less experience in managing their workforce leading to planning mistakes that are independent of worker tenure. Figure A10A studies this interaction between worker and plant characteristics. It displays the relationship between plant-level churn and employment growth for workers of low- and high-tenured workers, differentiating between plants younger and older than 5 years. The first takeaway message is that the patterns in Section 3 also hold conditional on plant age: Low-tenured workers have higher churn rates, and their V-shape is more pronounced. Furthermore, there is a small

Figure A9: Churn Rates and Wage Dynamics

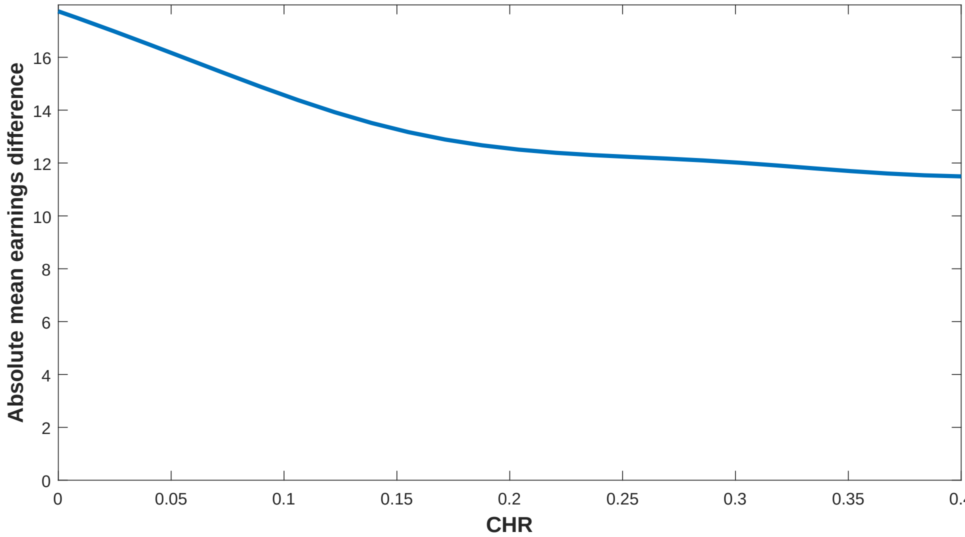

Note: This figure displays the absolute difference in average daily earnings of in- and out-flowing workers of a plant as a function of its churn rate. The estimate results from a semi-parametric regression which estimates a non-parametric relationship between churn and wage-differences (using a Gaussian kernel with 0.05 bandwidth), controlling linearly for the effect of flow sizes by including the square-root of total worker flows. It plots the function up to a churn rate of 0.4 which represents 97 percent of all employment. Each observation is weighted by the total worker flows at the plant. West German plants only with quarterly frequency, 1975Q1-2014Q4.

additional positive effect on the level of churn coming from young plants.

Figure A10B repeats the same analysis conditional on plant size. Again, we find the same patterns as in Section 3. First, even conditional on plant size, plants with many low-tenured workers have higher churn rates and a more pronounced V-shaped relationship between plant-level churn and employment growth. Second, large plants (at least 50 workers) display an even more pronounced V-shape than small plants.

In addition to predicting an overall negative autocorrelation of employment growth, our model in Section 4, importantly, also predicts that the strength of the autocorrelation varies systematically with the level of employment growth. The calibrated average separation rate shock in the model is close to 7%, that is, a plant at its optimal size today expects to lose 7% of its workforce on average. Therefore, plants that shrink by more than 7% are plants that experience a larger than expected separation rate shock this period (period t) and end the period with fewer than the optimal amount of workers. As a result, these plants hire many workers the following period (period $t+1$) and tend to grow, i.e., there is a strong negative correlation between today's and tomorrow's employment growth.

On the other extreme, plants that grow by more than 7% are plants that experience a larger than expected separation rate shock in the previous period (period $t-1$) and, thus, ended the previous period with fewer than the optimal amount of workers. Therefore, in period t they re-hire and are, on average, at their optimal employment level. As a result, their employment growth next period ($t+1$) shows little correlation with their employment growth this period (period t).

Plants with small negative or positive employment growth this period t tend to be plants without large separation rate shocks in the current or the previous period and, thus, end the current period close to their optimal employment level. Therefore, their employment growth today also shows little correlation with their employment growth next period.

Taken together, the model makes the stark prediction that the observed negative autocorrelation of employment growth comes mainly from plants that shrink by a substantial

Figure A10: Churn Rates and Plant-level Employment Growth By Tenure Composition and Plant Age and Plant Size

(A) Tenure Composition and Plant Age

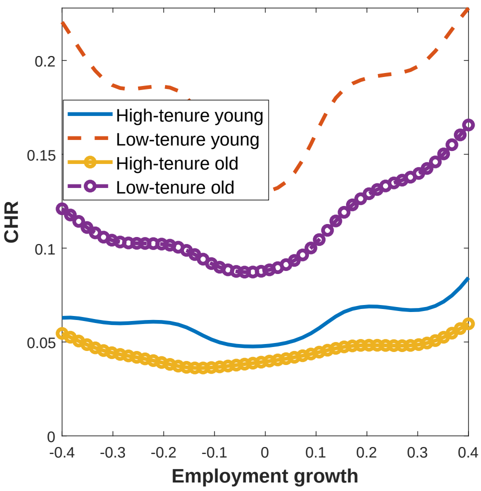

(B) Tenure Composition and Plant Size

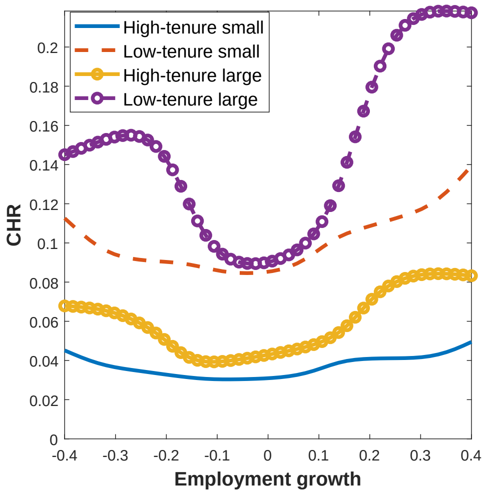

Note: Panel A displays the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). Plants are grouped in four categories based, first, on whether their share of workers with more than one year of tenure is below or above the overall plant-level median and based, second, on whether their age is strictly larger than 20 quarters (old plants) or smaller than this threshold (young plants). Panel B displays the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). Plants are grouped in four categories based, first, on whether their share of workers with more than one year of tenure is below or above the overall plant-level median and based, second, on whether their size is at least 50 workers (large plants) or below (small plants). West German plants with a quarterly frequency, 1975Q1-2014Q4.

amount in this period and grow by a substantial amount the following period. Figure A11 shows that this pattern holds in the data.

Finally, Table A3 displays moments of the churn rate across the employment growth rate distribution. The mean churn rates of the discrete employment growth rate bins correspond to the group means we display in Figure 1 in Section 3. The other columns display moments of the HP-filtered churn rates conditional on the employment growth rate bin. These results confirm the cyclical dynamics of the churn rate conditional on employment growth that we displays in Figure 5A in Section 5 using non-parametric estimates.

Figure A11: Autocorrelations of Employment Growth

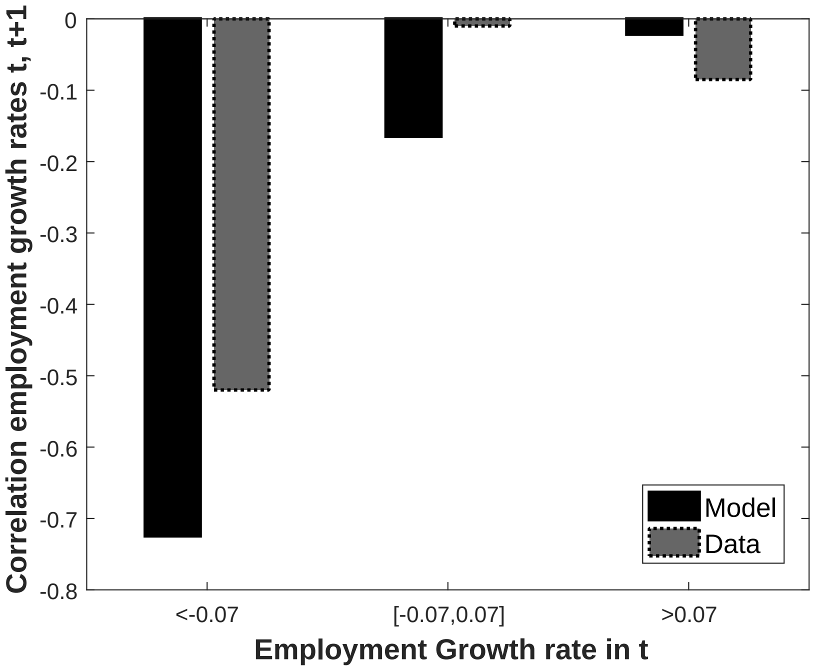

Note: This figure displays the correlation between contemporaneous and next quarter employment growth rates as a function of plants contemporaneous employment growth rate. The black bars are the data for the West-German sample 1975-2014. The gray bars are the calibrated baseline model.

Table A3: Cyclical Dynamics of the Churn Rate and Employment Growth

Employment growth rate	*Mean*	*Std*	$AC(1)$	*CorrU*
-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^{**}

Note: This table displays moments of churn over the employment growth distribution. The second column, *Mean*, displays the time-averaged churn rate. The subsequent columns display moments of the HP(100,000)-filtered churn rate. *Std*: standard deviation, $AC(1)$ autocorrelation coefficient, *CorrU*: correlation with the filtered unemployment rate. ** and *** indicate significance at the 1% and 5% level, respectively, obtained by non-parametric block-bootstrapping with a block length of 20. The aggregate churn rate in each employment growth category has been separately seasonally adjusted. West German plants with a quarterly frequency, 1975Q1-2014Q4.

C Frequency and Seasonality

Figure A12: Churn and Employment Growth with Seasonal Adjustment

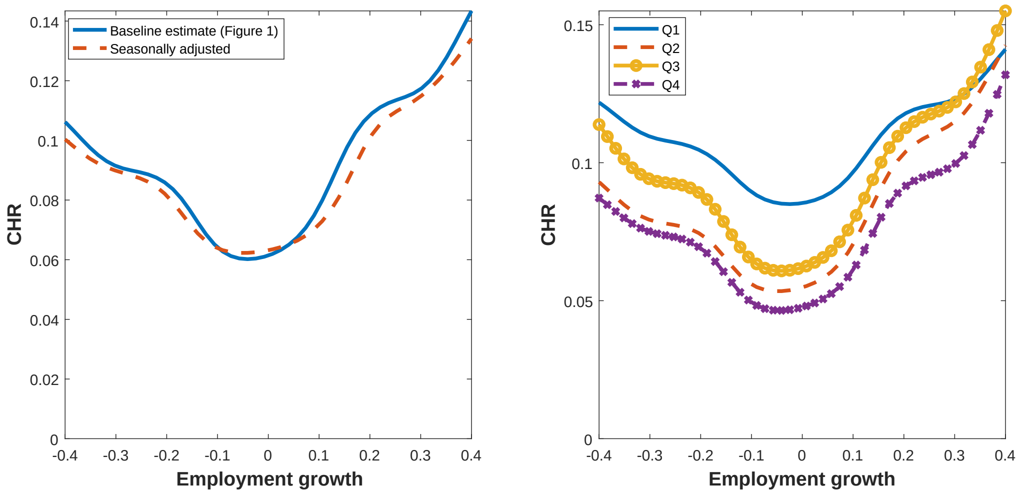

Note: The left panel displays the average churn rate of a plant as a function of its employment growth rate. The blue line is estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The red dashed line is estimated the same way after regressing plant-level employment growth rates on four quarter dummies. The right panel displays the kernel estimate for each quarter. West German plants only with quarterly frequency, 1975Q1-2014Q4.

This appendix studies the importance of seasonal effects for the relationship between plant-level churn and employment growth and extends the analysis in Section 3 to the yearly frequency.

Our baseline cross-sectional results in Figure 1 of the main text are for non-seasonally-adjusted data. To investigate the importance of seasonal effects, we seasonally adjust plant-level employment growth rates by regressing them on four quarter dummies. The left panel of Figure A12 shows that season adjustment leaves our results essentially unchanged. That is not to say that churn has no seasonality. The right panel of Figure A12 shows that the level of churn is different across the different quarters, yet the V-shaped churn-employment growth nexus is present in all four quarters. Finally, the negative autocorrelation of employment growth rates remains at -0.15 after seasonal adjustment.

Turning to the data at the yearly frequency, Figure A13 shows non-parametric estimates of the relationship between plant-level churn and employment growth at the yearly frequency. As expected, churn is higher at the yearly frequency compared to the quarterly frequency throughout the employment growth distribution. Importantly, churn is also V-shaped at the yearly frequency. That is, plants adjusting their size rapidly from one year to the next have higher churn rates than plants that change their level of employment little. The difference between the peak to trough for shrinking plants is five percentage points, similar to the quarterly churn data. For growing plants, the difference is somewhat larger than in the quarterly data. As in the quarterly data, these results are not driven by small plants. Finally, the autocorrelation of employment growth at the yearly frequency is also negative. It is even more negative (-0.23) than at the quarterly frequency (-0.15).

Figure A13: Churn and Employment Growth at the Yearly Frequency

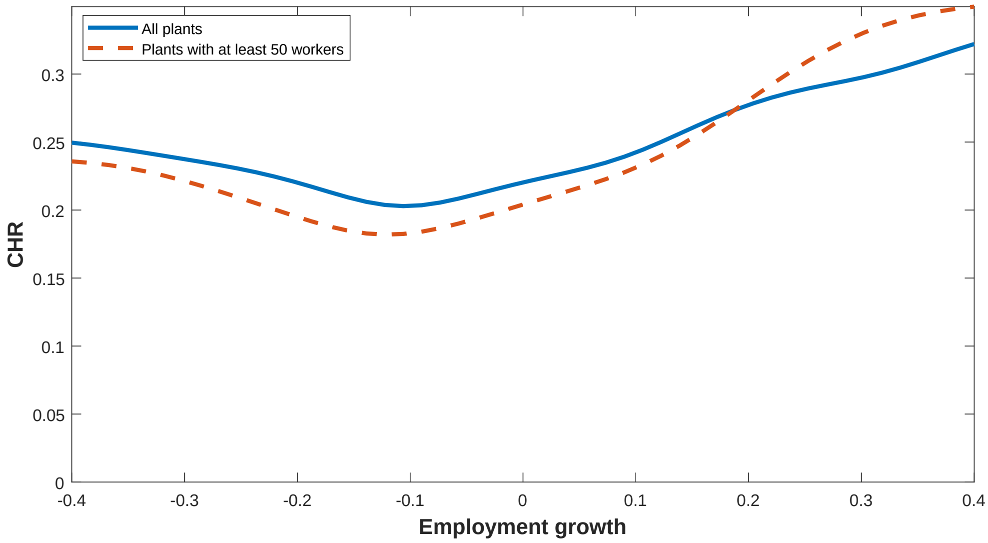

Note: This figure displays the average churn rate of a plant as a function of its employment growth rate. The blue line is estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The red dashed line is estimated the same way for those plants with more than 49 employees. West German plants only with annual frequency, 1975-2014.

D Alternative Models of Employment Dynamics and Churn

In this appendix, in a first step, we discuss basic alternatives to our baseline model and show that stochastic separations and the time-to-hire friction are essential to produce the observed V-shaped plant churn-employment growth nexus. First, we analyze a model with productivity shocks instead of separation rate shocks. Second, we study a model with separation rate shocks and convex employment adjustment costs but no time-to-hire friction.

In a second step, we discuss extensions to our baseline model. Here, we first add separation rates that vary with worker tenure to the baseline model, and, second, we add convex employment adjustment costs to the baseline model.

D.1 Alternatives to the Baseline Model

D.1.1 Productivity Shocks

We consider a model where plants have the same decreasing returns-to-scale production function in employment as in the main text, but 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 production function:

$$Y_{it} = z_{it} E_{it}^{\alpha}, \tag{A.5}$$

where E_{it} is the plant's (end-of-quarter t) 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:

$$\log z_{it} = (1-\rho)\mu_z + \rho_z \log z_{it-1} + \epsilon_{it}, \quad \epsilon_{it} \sim N(0, \sigma_z^2). \tag{A.6}$$

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

$$E_{it} = (1-\bar{s})E_{it-1} + \Delta^a_{E_{it}}. \tag{A.7}$$

The plant decides on $\Delta^a_{E_{it}}$ 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: $c_{it} = \psi\left(\Delta^a_{E_{it}}\right)^2$. Plants choose their active employment adjustment to maximize their value:

$$V(z, E_{-1}) = \max_{\Delta^a_E}\left\{zE^{\alpha} - wE - \psi\left(\Delta^a_E\right)^2 + \frac{1}{1+r}\mathbb{E}V(z', E)\right\}, \text{such that } E = (1-\bar{s})E_{-1} + \Delta^a_E,$$

where r is the discount rate and w is the wage rate. To sum up, the timeline of events at a plant within a period is: revelation of productivity, (deterministic) separations, active employment adjustment, production, and wage payment.

In terms of worker flow accounting, when $\Delta^a_{E_{it}} > 0$, this active adjustment is counted as hires in the model, i.e., $H_{it} = (\Delta^a_{E_{it}})^+$; as for separations, in this case: $S_{it} = \bar{s}E_{it-1}$. When $\Delta^a_{E_{it}} < 0$, $H_{it} = 0$, and we count the active adjustment as additional separations in the model, i.e., $S_{it} = \bar{s}E_{i,t-1} + (\Delta^a_{E_{it}})^-$.

It is straightforward to see that for negative employment growth rates smaller than $-\bar{s}$, there is no hiring and, hence, according to Equation (3) churn is zero, in contrast to the data. For employment growth rates larger than $-\bar{s}$, plants re-hire for the workers lost

Figure A14: Churn and Employment Growth – Model with Productivity Shocks

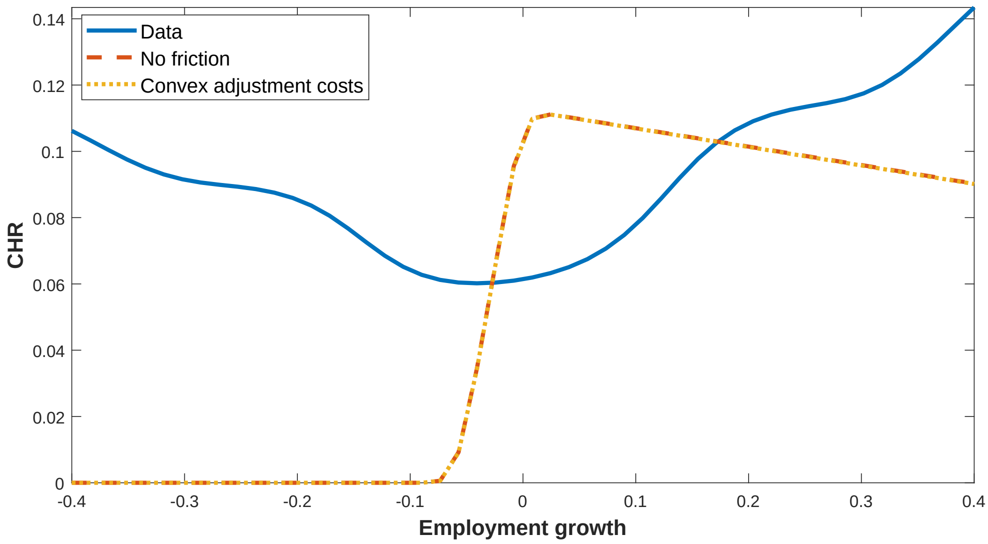

Note: This figure shows the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The blue solid line is the data for the West-German sample 1975-2014. The yellow dotted line, *convex adjustment costs*, displays the churn rates from the calibrated model with productivity shocks and convex adjustment costs. The red dashed line, *No adjustment costs*, displays the churn rates in the same model but adjustment costs are set to zero, with all other parameters recalibrated.

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 $(E_{it-1} + E_{it})/2$ in the denominator of worker flow and churn rates, the churn rate turns out to 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 A14. The parameters of this simple model are the wage, w, the returns to scale parameter, α, the quarterly interest rate r, the mean of the log productivity process, μ_z, the auto-correlation, 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_z to $0.992 = 0.9675^{\frac{1}{4}}$ (the 0.9675 being the estimate of Bachmann and Bayer (2014) for annual data from Germany), and use μ_z to match the average plant size in our data of 12.6. We obtain the remaining three parameters, σ_z, ψ, and $\bar{s}$, by a simulated minimum distance calibration. Our moments are the same as for the model in the main text; specifically, the plant average churn rate (obtained by the kernel estimate) at 50 equally spaced employment growth categories on the interval $[-0.4, 0.4]$, and the aggregate separation rate of 7.0% (with a 50 times larger weight). The left panel in Table A4 displays the calibrated parameters.[4]

Figure A14 compares churn rates as a function of plant-level employment growth rates 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 our framework, as Figure A14

[4] The parameters for the frictionless recalibration are not displayed for the sake of brevity but are available upon request.

Table A4: Parameters of the Alternative Models

Productivity Shocks				Separation Rate Shocks				
μ_z	σ_z %	ψ	$\bar{s}$	z	ψ	μ_s	σ_s	ρ_s
1.34	0.35	0.02	0.06	4.76	1.12	-3.44	0.74	0.24

Note: The left panel shows the parameters for a model with productivity shocks and convex employment adjustment costs. μ_z: mean of log-productivity. σ_z: standard deviation of log-productivity shocks. ψ: scaling parameter of the convex employment adjustment cost function. $\bar{s}$: exogenous separation rate. The right panel shows the parameters for a model with separation rate shocks and convex employment adjustment costs. z: productivity. ψ: scaling parameter of the convex employment adjustment cost function. μ_s: mean of the $AR(1)$ process governing separation rate shocks. σ_s: standard deviation of shocks to the $AR(1)$ process governing separation rate shocks. ρ_s: autocorrelation of the $AR(1)$ process governing separation rate shocks.

also demonstrates.

D.1.2 Separation Rate Shocks Plus Convex Employment Adjustment Costs

To highlight the importance of the time-to-hire friction, consider now a hybrid model with separation rate shocks where deviations from optimal employment are due to quadratic employment adjustment costs instead of the time-to-hire friction.

Using the notation from the main text and the previous section, the dynamic program of a plant in such an environment is given by:

$$
\begin{aligned}
&V(\tilde{s}, E_{-1}) = \max_{\Delta_E^a}\Big\{zE^\alpha - wE - \psi\Big(\Delta_E^a\Big)^2 + \frac{1}{1+r}\mathbb{E}V(\tilde{s}', E)\Big\} \\
&E = (1-s)E_{-1} + \Delta_E^a \\
&s = \min\{1, \exp(\tilde{s})\} \\
&\tilde{s}' = (1-\rho_s)\mu_s + \rho_s\tilde{s} + \epsilon, \quad \epsilon \sim N(0, \sigma_s^2).
\end{aligned}
$$

The timeline of events at a plant within a period is thus basically the same as in the previous section: (stochastic) separations, active employment adjustment, production, and wage payment. Also, the worker flow accounting works the same way as in the previous section, with the deterministic separation rate, $\bar{s}$, replaced by the stochastic s.

To get a quantitative impression of the model's implied churn behavior, we follow again our calibration strategy. We set 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 productivity, z, the adjustment costs, ψ, and the parameters guiding the uncertainty of the separation rate, ρ_s, μ_s, and σ_s. As before, we choose z to match the average plant size in the data and obtain the other parameters by our minimum distance calibration. The right panel in Table A4 displays the calibrated parameters of the separation shock distribution.

Figure A15 compares the model-implied churn-employment growth nexus to that of the data. We find too much churn for rapidly shrinking plants, and too little churn, in particular for moderately growing plants. We finish this section by explaining why this pattern is to be expected from a model where employment adjustment costs instead of the time-to-hire friction generate deviations from the optimal employment target.

The nexus between employment growth and separations rate shocks is similar as in the model with the time-to-hire friction because employment adjustment costs lead to a delayed adjustment to separation rate shocks. That is, plants that shrink rapidly tend to be plants that have suffered a large separation rate shock in the current period. By contrast, plants that grow rapidly tend to be plants that have had a large separation rate shock in the previous

Figure A15: Churn and Employment Growth – Model with Separation Shocks and Convex Employment Adjustment Costs

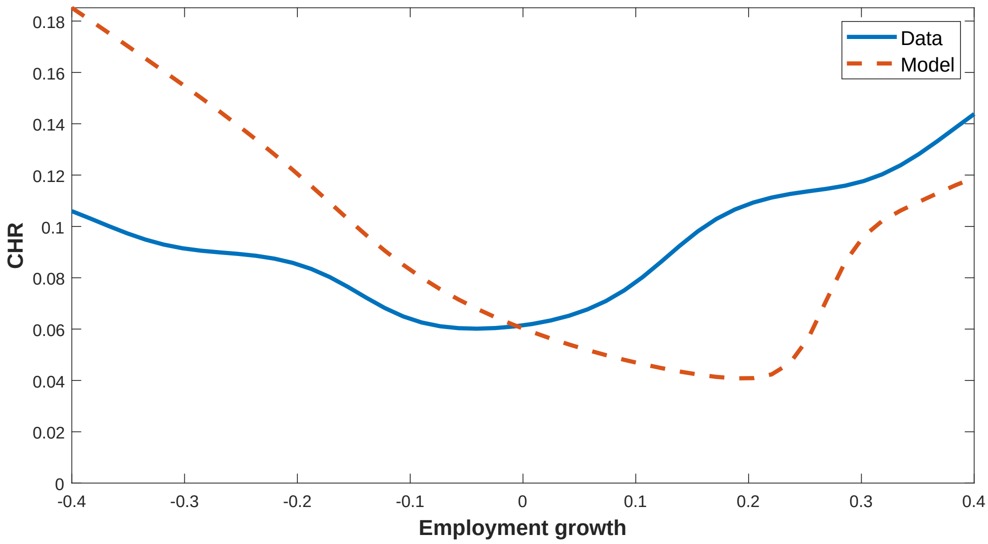

Note: This figure shows the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The blue solid line is the data for the West-German sample 1975-2014. The red dashed line displays the churn rates from the calibrated model with quadratic employment adjustment costs and stochastic separations.

period. Finally, plants that grow or shrink little have had small separation rate shocks both in the previous and the current period. The crucial difference, however, is that, in a model with only employment adjustment costs, the largest part of the adjustment to separation rate shocks is realized within-period. This within-period rehiring leads to a mechanically positive relationship between separation rate shocks and churn. Since plants that shrink rapidly tend to have large separation rates, this tilts the V-shaped relationship of the time-to-hire model towards additional churn in rapidly shrinking plants. Moderately growing plants have now the smallest churn because they tend to have the smallest separations rates. In the time-to-hire model, these plants re-hired in reaction to the expected separation rate shock, whereas in the adjustment cost model they re-hired in reaction to their small realized separation rate shock. For a given moderate employment growth rate, this makes the churn of the time-to-hire plants higher than that of the adjustment cost plants.

D.2 Extensions to the Baseline Model

D.2.1 Adding Heterogeneous Separation Rates by Worker Tenure

We begin by considering the role of worker tenure for separations. In the baseline model, we assume that all workers of a plant have the same separation rate (note that, in the baseline model, separation rates are already heterogeneous across plants). However, in the data, workers with at most one year of tenure have an average separation rate of 15.3%. Workers with longer tenure have an average separation rate of only 4.7%. To model this difference, we now consider two types of workers. Low-tenured workers, L, are those with at most one year of tenure. High-tenured workers, H, are those with more than one year of tenure. As in the baseline model, plants maximize expected profits:

$$\max_{\Delta_E^a} \Big\{ \mathbb{E}_{t-1}\{Y_{it} - wE_{it}\} \Big\}, \tag{A.8}$$

Table A5: Parameters of the Extended Baseline Models

Quadratic Adjustment Costs				Heterogeneous Separation Rates			
μ_s	σ_s	ρ_s	ψ (%)	μ_s	κ	σ_s	ρ_s
-3.16	0.95	0.35	0.64	-2.88	0.51	0.76	0.49

Note: The left panel shows the parameters for a model with separation rate shocks and quadratic employment adjustment costs. μ_s: mean of the $AR(1)$ process governing separation rate shocks. σ_s: standard deviation of shocks to the $AR(1)$ process governing separation rate shocks. ρ_s: autocorrelation of the $AR(1)$ process governing separation rate shocks. ψ: scaling parameter of the quadratic employment adjustment cost function. The right panel shows the parameters for a model with separation rate shocks and heterogeneity in worker tenure. κ: parameter guiding the difference in separation rates between workers of different tenures.

with $E_{it} = E_{it}^L + E_{it}^H$. Stochastic separations are driven by a latent variable:

$$\tilde{s}_{it} = (1-\rho_s)\mu_s + \rho_s \tilde{s}_{it-1} + \epsilon_{it}, \quad \epsilon_{it} \sim N(0, \sigma_s^2). \tag{A.9}$$

Differently from the baseline model, the mean separation rate is now allowed to be different for low- and high-tenured workers, and we capture this difference by the parameter κ:

$$s_{it}^L = \min\{\exp(\tilde{s}_{it} + \kappa), 1\} \tag{A.10}$$

$$s_{it}^H = \min\{\exp(\tilde{s}_{it} - \kappa), 1\}. \tag{A.11}$$

We calibrate the model the same way as the baseline model: We set the returns-to-scale parameter, α, to 0.6, normalize the wage to $w = 1$, and we choose plant productivity, z, to match the average plant size in the data of 12.0. We obtain the parameters guiding the uncertainty of the separation rate, ρ_s, μ_s, σ_s, and κ, by a simulated minimum distance calibration. Our moments are the plant average churn rate (obtained by the kernel estimate) at 50 equally spaced employment growth categories on the interval $[-0.4, 0.4]$. The only difference to the baseline calibration is that we replace the aggregate separation rate with the separation rates of workers with at most one year and more than one year of tenure. Table A5 displays the resulting parameters.

Figure A16 shows that adding heterogeneous worker tenure does not change the model's ability to match the relationship between plant-level churn and employment growth. Moreover, the calibrated model matches almost perfectly the separation rate of high-tenured workers (15.1%) and low-tenured workers (4.7%) and marginally improves the fit of the V-shaped churn-employment growth relationship for rapidly growing plants.

D.2.2 Adding Convex Adjustment Costs

Next, we add convex employment adjustment costs to the baseline model (without heterogeneous separation rates by worker tenure). The number of employed workers has the same law of motion as in the baseline model:

$$E_{it} = (1 - s_{it})(E_{it-1} + \Delta_{E_{it}}^a). \tag{A.12}$$

Differently from the baseline model, plants now pay also convex adjustment costs when actively adjusting their workforce, on top of the time-to-hire friction: $c_{it} = \psi\left(\Delta_{E_{it}}^a\right)^2$. The

Figure A16: Churn and Employment Growth – Baseline Model with Heterogeneous Separation Rates by Worker Tenure

Note: The figure displays the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The blue solid line is the data for the West-German sample 1975-2014. The red dashed line displays the churn rates from the calibrated baseline model with heterogeneous separation rates. The yellow dotted line displays the baseline model without heterogeneous separation rates.

resulting dynamic programing problem becomes:

$$V(E_{-1},\tilde{s}_{-1}) = \max_{\Delta_E^a}\Big\{\mathbb{E}_{-1}z\{E^\alpha - wE\} - \psi\Big(\Delta_E^a\Big)^2 + \frac{1}{1+r}\mathbb{E}_{-1}V(E,\tilde{s})\Big\} \tag{A.13}$$

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

$$s = \min\{\exp(\tilde{s}), 1\}$$

$$\tilde{s} = (1-\rho_s)\mu_s + \rho_s\tilde{s}_{-1} + \epsilon, \quad \epsilon \sim N(0,\sigma_s^2),$$

where $\mathbb{E}_{t-1}$ is an expectation operator with the associated information set $\{E_{-1},\tilde{s}_{-1}\}$. We calibrate the model the same way as the baseline model: We set the returns-to-scale parameter, α, to 0.6, normalize the wage to $w = 1$, and we choose plant productivity, z, to match the average plant size in the data of 12.0. We obtain the parameters guiding the uncertainty of the separation rate, ρ_s, μ_s, and σ_s, by a simulated minimum distance calibration. Our moments are the plant average churn rate (obtained by the kernel estimate) at 50 equally spaced employment growth categories on the interval $[-0.4, 0.4]$, and the aggregate separation rate of 7.02% (with a 50 times larger weight). This leaves us with the parameter guiding adjustment costs, ψ. Convex adjustment costs reduce active per period employment adjustments overall and, thus, in particular move the mass of active downsizers towards zero employment growth. We choose ψ such that the plant-level churn employment growth relationship has its minimum as in the data. Table A5 displays the resulting parameters.

Figure A17 shows that this extended model matches the V-shaped relationship between plant-level churn and employment growth. What is more, convex adjustment costs imply that plants adjust to their optimal employment level slowly and, thus, induce a positive correlation in employment growth that reduces the overall negative autocorrelation of employment growth in the model from -0.48 to -0.37, thus, bringing the model also in this dimension closer to the data. Overall, introducing convex employment adjustment costs as an additional layer of realism improves the model fit but does not alter the overall insights of the paper.

Figure A17: Churn and Employment Growth – Baseline Model Plus Convex Employment Adjustment Costs

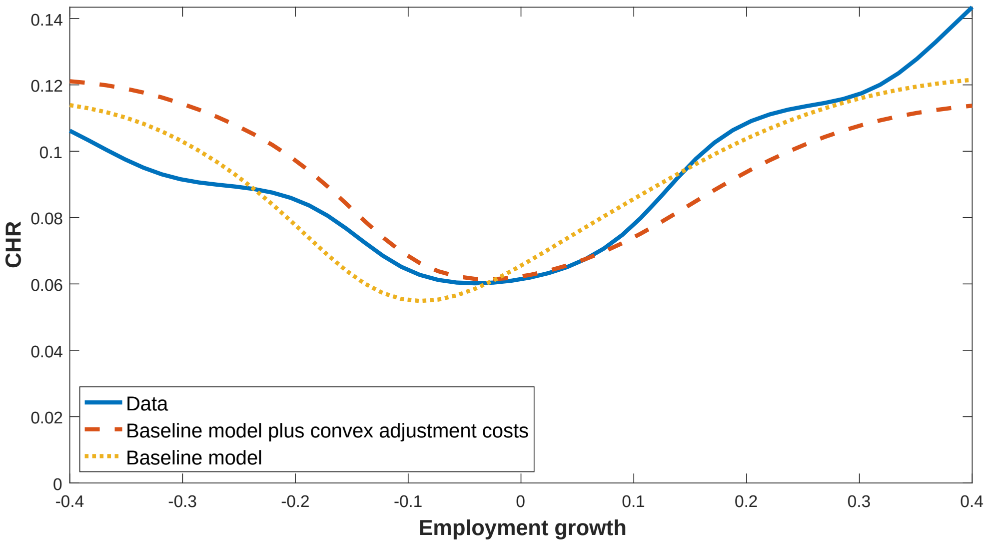

Note: The figure displays the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it}-weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The blue solid line is the data for the West-German sample 1975-2014. The red dashed line displays the churn rates from the calibrated baseline model plus convex adjustment costs. The yellow dotted line displays the baseline model without convex adjustment costs.

References

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

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. (2012). Labor market flows in the cross section and over time. *Journal of Monetary Economics*, 59(1):1–18.

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

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

Hyatt, H. R. and Spletzer, J. R. (2017). The recent decline of single quarter jobs. *Labour Economics*, 46(2):166–176.

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

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