

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*^{†‡}

April 2, 2020

Abstract

Worker churn, that is, worker flows in excess of job flows, is procyclical in the German labor market. To understand this procyclicality, we study the plant-level connection of churn and employment growth, using the new Administrative Wage and Labor Market Flow Panel from 1975 to 2014, and find that churn rises in the absolute value of employment growth. Analyzing this V-shaped churn-employment growth nexus by worker skill, age, and tenure, we establish that churn is unlikely to result from plant reorganization but rather from uncertainty about match quality. Using a simple dynamic labor demand framework, we argue that the cross-sectional evidence on churn can be interpreted as manifestations of idiosyncratically stochastic separation shocks in conjunction with a time-to-hire friction. These shocks become larger and more predictable during booms leading to procyclical churn, which, as we show, (1) increases almost uniformly across the employment growth distribution, and (2) stems almost exclusively from job-to-job transitions. Procyclical churn, thus, reflects a more active reshuffling of workers towards individually better matches in booms.

Key Words: worker churn, employment growth, job flows, worker flows, labor demand, separation shocks, job-to-job transitions, aggregate fluctuations

JEL Classification: E20, E24, E32, J23, J63

*Corresponding author: Rüdiger Bachmann, University of Notre Dame, CEPR, ifo, and CESifo, rbachman@nd.edu.

[†]Christian Bayer, University of Bonn and CEPR. Christian Merkl, Friedrich-Alexander Universität Erlangen-Nürnberg (FAU), Institute for Employment Research (IAB), CESifo, and IZA. Stefan Seth, IAB. Heiko Stüber, FAU, IAB, and IZA. Felix Wellschmied, Universidad Carlos III de Madrid and IZA.

[‡]We thank seminar participants at the Royal Economic Society Meeting, the Verein für Socialpolitik, the EES Conference, the 2nd IZA@DC Young Scholar Program, the 29th Annual Conference of the European Association of Labour Economists, the Annual Conference of the Scottish Economic Society, the ZEW Research Seminar, the 6th DFG SPP 1764 Workshop, the 24th Annual Meeting of the Society of Labor Economists, the ECB/CEPR 2019 Labour Market Workshop, the Brown Bag seminar of the BMF, and the IMK seminar. We thank James Costain, Alexander Herzog-Stein, Andreas Kostøl, Jim Spletzer, and two anonymous referees for helpful comments. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FTP/2007-2013) / ERC Grant agreement no. 282740. Felix Wellschmied gratefully acknowledges support from the Spanish Ministry of Economics through research grants *ECO2014-56384-P*, *MDM 2014-0431*, and *Comunidad de Madrid MadEco-CM (S2015/HUM-3444)* and thanks the Department of Economics at ITAM for its hospitality. Heiko Stüber gratefully acknowledge support from the German Research Foundation (DFG) under priority program "The German Labor Market in a Globalized World" (SPP 1764). Christian Merkl gratefully acknowledge support from SPP 1764 and the Hans Frisch Stiftung. This paper supersedes the earlier IZA Discussion Paper 7192 by a subset of the authors of this paper, entitled: "Cyclicality of job and worker flows: new data and a new set of stylized facts," as well as CEPR DP 12343, CESIFO WP 6702, and IZA DP 11063 "Worker Churn and Employment Growth at the Establishment Level" by the same set of co-authors.

1 Introduction

Many establishments hire new workers while separating from some of their existing workforce within relatively narrow windows of time. This leads to worker turnover in the economy that is larger than the observed job creation and destruction; in short, there exists worker churn (see Burgess et al., 2000; Davis et al., 2006, 2012). We document that behind this churn are worker flows that are twice as large as job flows not only in the United States but also in Germany—with the aggregate job and worker flow rates being about half the size in Germany compared to the United States. In our sample, the aggregate German worker churn rate is around 6.7% of employment each quarter. Statistically, almost 50% of worker turnover is due to worker churn while the rest is due to job flows. Moreover, during boom times, when the unemployment rate is low, the worker churn rate is about 3 percentage points (or 40 percent) higher than during recession times. The churn rate is also highly persistent over the business cycle. To sum up: worker churn is strongly procyclical and persistently so.

To better understand the procyclical behavior of worker churn, we start by studying the cross-sectional relationship between worker churn and employment growth at the establishment/plant level.¹ For our analysis, we use the new plant-level 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-level employment growth to hiring decisions (from other establishments and non-employment) and separation decisions (to other establishments and non-employment).² The AWFP also provides plant-level worker stocks and flows by worker characteristics, such as education, job task, tenure, and age.

We find that the churn-employment growth nexus is V-shaped. Plants that grow fast on average separate from more workers than plants that grow little; and plants that shrink fast hire more workers on average than plants that shrink little. On average, the smallest churn occurs in plants with a relatively stable workforce.

Conceptually, worker churn can arise for two distinct reasons. Plants may wish to restructure their production process and, hence, replace workers with some ex-ante known characteristics with workers with other ex-ante known characteristics, and this re-organization could be linked to plant growth. For instance, plant reorganization may be connected to large changes in (optimal) plant size when a plant replaces production workers with information technology (IT) specialists and robots, and, as a result, shrinks. In a different version of the story, the plant may grow through reorganization because it adds sales representatives as the robots produce larger quantities.

¹In this paper, we use “establishment” and “plant” interchangeably.

²See Davis et al. (2012) for a comprehensive overview of similar U.S. data. The AWFP has some advantages over U.S. data. One major obstacle for studying links between job and worker flows in the United States is the availability of data sets that provide information on establishment characteristics, worker flows, and job flows. The most commonly used U.S. data source is the Job Openings and Labor Turnover Survey (JOLTS), used, for example, by Davis et al. (2006), sampling 16,000 establishments in the United States every month. 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 not only the cross section but also the cyclical behavior of job and worker flows and their interaction.

Alternatively, churn could result from uncertainty about the quality of the matches workers and plants form, particularly at the beginning of a match, and from the (partial) resolution of said uncertainty. If this match quality uncertainty is resolved in a somewhat random fashion and if it takes time to replace separations, then plants might grow or shrink because, at least in the short run, they cannot control the size of their workforce exactly. In other words, uncertain separations lead to both churn and planning mistakes that manifest themselves in short-term, plant-level employment fluctuations.

To distinguish between these two explanations for plant-level churn and its link to employment growth, we analyze the churn-employment growth nexus by observable worker or job characteristics, specifically worker education, job tasks, and worker age. We find that most of the level of churn and all of the V-shaped nexus between churn and employment growth occur within a job or worker category. Moreover, we find that new hires and separating workers have more similar wages at plants that churn more. Under the reorganization hypothesis one would expect the opposite. Therefore, we interpret these facts altogether as pointing away from reorganization as the driver of churn. Furthermore, plants with a low-tenure workforce, which we use as a proxy for exposure to high match uncertainty, have higher levels of churn, a more pronounced V-shape, and a more dispersed employment growth rate distribution than plants with a high-tenure workforce. We interpret these facts as pointing towards a story of uncertain match quality that is stochastically resolved.

In light of these findings, we set up a simple heterogeneous-plant dynamic labor demand framework. This model has two key elements: (1) plants are hit by persistent idiosyncratic separation rate shocks to which (2) they cannot react immediately (time-to-hire friction). In the model, large job destruction is mainly driven by large contemporaneous separation rate surprises, whereas job creation is mainly driven by corrections of past separation shocks. Important for worker flows, plants not only re-hire to compensate for past separations, but they also hire because of expected separations. Both effects together with separation shocks in the current period create churn and a V-shaped relationship between churn and employment growth. We view these two elements as a parsimonious, semi-structural way to capture the main insights from our cross-sectional analysis of the data.

A calibrated version of this model replicates the churn-employment growth nexus not only qualitatively but also quantitatively. At the same time, the stochastic separations cum time-to-hire model is also consistent with the dispersion of the plant-level employment growth distribution in the data, an untargeted object. Through the lens of our model, this, in turn, means that plant-level short-term employment growth can be driven by transitory planning mistakes as regards the size of a plant's workforce.³

³By way of an appendix, we also show that in a perhaps more standard model with idiosyncratic productivity shocks at the plant-level, churn only occurs at plants that have not suffered too large negative productivity shocks and wish to replace exogenously leaving workers. This is fundamentally at odds with the V-shaped relationship between plant-level employment growth and churn. To be clear, we do not mean to say that productivity shocks are not important ingredients to understand plant-level employment dynamics. Nevertheless, our analysis does suggest that another shock, stochastic separations, together with a time-to-hire friction, is important to understand the joint behavior of plant growth and churn.

Next, we study the cyclical properties of worker churn. As mentioned above, churn increases substantially during booms. This increase happens almost uniformly across the entire plant-level employment growth distribution. As a consequence, the procyclical behavior of worker churn, differently from hiring and separation rates, stems entirely from these uniform changes in churn within employment growth categories and not from cyclical changes of the employment growth rate distribution.

When viewed through the lens of our simple model, the decline in churn at all levels of employment growth in recessions is the result of two factors: separation rates are smaller on average in recessions but also more uncertain. While lowering average stochastic separation rates makes the model match the lower recession separation rates in the data, lower average stochastic separation rates are not sufficient to fully explain the lower churn rates in recessions. To rationalize these lower churn rates, separations also need to be more uncertain. Higher separation uncertainty renders what is worker turnover job turnover: An unexpected separation, in the model, produces a job destruction. Since churn is the difference between worker flows and job flows, higher uncertainty, for a given average separation rate, reduces churn.

Digging deeper, we decompose separations (and hires) into those going to other establishments and those going to (coming from) non-employment. We show that a parallel upward shift of separations (and hires) to other establishments drives the increased worker churn in a boom relative to a recession. Worker transition rates through the non-employment pool show no such cyclical behavior. As a result, cyclical aggregate worker churn is almost identical to procyclical job-to-job transitions. What is more, we show that hiring from non-employment (separation into non-employment) is very close to job creation (job destruction) in terms of its cyclical behavior. In other words, the observed increase in worker reallocation in (at least mature) booms create little additional job reallocation, which lines up well with our finding above that plants are more certain about separations in booms.

Our paper contributes to the literature that studies employment dynamics at the establishment level. The most common interpretation of these dynamics is one through plant-level idiosyncratic productivity shocks (see Hopenhayn, 1992). We show that separation rate shocks also contribute to short-run employment dynamics and are important to understand their nexus with churn. Regarding the business-cycle properties of employment dynamics, Davis and Haltiwanger (1992) show that job reallocation in excess of net employment growth is close to acyclical. Our contribution is to link this acyclicity to the underlying worker flows and the uncertainty plants face about those flows. A recent literature focuses on job-to-job transitions as one particular procyclical worker flow (see Barlevy (2002), Moscarini and Postel-Vinay (2012), Schaal (2017), and Fujita and Nakajima (2016)). Moscarini and Postel-Vinay (2013) develop a framework that links establishment *size* to cyclical job-to-job transitions. Haltiwanger et al. (2018) argue that establishment *pay* is a better explanation for cyclical job-to-job transitions. Tanaka et al. (2019) link *pay growth* to job-to-job transitions and churn. Our contribution to this literature is twofold. First, we show that cyclical job-to-job transitions are not systematically linked to establishment *growth*, that is, job-to-

job transitions rise by a similar amount across the employment growth distribution during booms. Second, we show that procyclical job-to-job transitions lead ultimately to worker churn and not to procyclical job reallocation. Consistent with the absence of a systematic reallocation of workers towards particular plants, Lindenlaub and Postel-Vinay (2017) develop a job-ladder model where workers have idiosyncratic rankings of plants.

Finally, our paper is related to the very recent literature on mismatch, reorganization, and plant-level employment dynamics: Borovicková (2016), building on an earlier model of mismatch by Pries and Rogerson (2005),⁴ provides a structural framework where separations arise from mismatch and learning about individual match quality in multi-worker firms. In her model, positive productivity shocks at the firm level lead to less precise knowledge of match quality (and thus more separations) because a larger fraction of workers are new hires, and all productivity shocks change the information set about existing workers as they get assigned to new tasks. Gulyas (2018), by contrast, models assortative matching (but with certain match quality) as the central reason for the response of worker turnover to productivity shocks. With assortative matching, the ideal worker type for a plant changes after a shock to productivity and, hence, plants that change productivity have more reason to separate from existing workers, as their match quality, the distance of plant and worker type, typically declines. Scaling his single-worker plant model up to a multi-worker environment, the model suggests a link of plant growth to churn as we see it in the data. Therefore, Gulyas (2018) provides a worker-plant interaction based micro-foundation for the surprise separations that play a key role in our model.

The remainder of this paper is organized as follows. Section 2 introduces the new AWFPP data set and explains the main concepts that we use to analyze the data. It also provides aggregate statistics about job turnover, worker turnover, and churn in Germany. Section 3 analyzes the relationship of employment growth and churn in the cross section of establishments, overall and for a variety of worker and job characteristics. Section 4 provides our simple model framework of worker churn that serves to rationalize our findings from Section 3 with separation rate shocks and a time-to-hire friction. Section 5 discusses the cyclical behavior of churn in light of this model and our previous cross-sectional findings. Section 6 concludes. An extensive appendix provides further details on data and measurement as well as additional details on job and worker flows in our data both along the cross-sectional as well as along the time-series dimension. It also provides a detailed comparison to U.S. data and serves as a bridge to the existing U.S. literature on job and worker flows. Finally, it discusses a number of alternatives and extensions to our baseline model.

⁴Pries and Rogerson (2019) use a similar model to explain the secular decline of worker turnover in the U.S. through better screening mechanisms available to employers which improve initial match qualities. It is important to note that Germany experienced no secular decline in worker turnover, as shown in Appendix A.

2 Data Set and Variable Definitions

2.1 The Administrative Wage and Labor Market Flow Panel

The new Administrative Wage and Labor Market Flow Panel (AWFP) provides employment, labor flow, and earnings data⁵ for the universe of German establishments (*Betriebe*) for the years 1975-2014 (see Stüber and Seth (2017)). The AWFPs main data source is the Employment History (*Beschäftigten Historik*, BeH) of the German Institute for Employment Research (IAB). The BeH is an individual-level data set covering all workers in Germany subject to social security.⁶ 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 and end date 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.⁷ To ensure consistency over time, most variables in the AWFP—and all variables used in this 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 reunified Germany thereafter). It is available at an annual and a quarterly frequency. For most of our analyses, we use the AWFP at the quarterly frequency and drop all establishments that are on the territory of former East-Germany and Berlin or for which we cannot determine the German state (*Bundesland*) in which an establishment is located, to avoid a break in the series. A big advantage of the AWFP is that it allows us to aggregate important characteristics of worker stocks and flows onto the establishment level, such as worker education, job tasks that workers are assigned to, worker tenure, worker age, and, in the case of flows, whether they came from (left to) non-employment or employment.

2.2 Variable Definitions

In the AWFP, a worker is considered to be working for a given plant in a given quarter when she is employed at this plant at the end of the quarter.⁸ From this definition follows the number of jobs at a plant i at the end of a quarter, E_{it} , the number of hires, H_{it} (a worker that was not working for that plant at the end of the previous quarter), as well as

⁵Merkel and Stüber (2016) use the AWFP to analyze the effects of wage dynamics on labor flows.

⁶Marginal part-time workers (*geringfügig Beschäftigte*) have been covered since 1999. The main types of employees not covered by the BeH are civil servants (*Beamte*), military personnel, and the self-employed.

⁷Before this aggregation, the data on individuals undergo numerous validation procedures. Further details on the data set are described in Stüber and Seth (2017). Conceptual differences between the AWFP and U.S. data are discussed in Appendix A.

⁸It turns out that, in Germany, most workers leave or join a plant at the end/the beginning of a quarter.

the number of separations, S_{it} (a worker that no longer works for a plant at the end of the quarter). These are the basic data from which all data series are derived.

We compute the net job flow at a plant as $JF_{it} = E_{it} - E_{it-1}$. When a plant decreases employment within a quarter ($JF_{it} < 0$), 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. To capture the extent of such worker reallocation in excess of job flows, Burgess et al. (2000) introduce the concept of *worker churn*.⁹ Worker churn is defined as the sum of plants' hirings in excess of job creation and their separations in excess of job destruction:

$$CH_{it} = (H_{it} - JC_{it}) + (S_{it} - JD_{it}). \quad (1)$$

We next define flow rates, where we use the average of contemporaneous and lagged end-of-quarter employment as the denominator:¹⁰

$$N_{it} = [E_{it} + E_{it-1}]/2.$$

For example, the hiring rate is given by:

$$HR_{it} = \frac{H_{it}}{N_{it}}. \quad (2)$$

The separation rate, (SR_{it}), the job creation rate (JCR_{it}), the job destruction rate (JDR_{it}), and the churn rate (CHR_{it}) are defined analogously.

We obtain aggregate time-series rates (JCR_t , JDR_t , HR_t , SR_t , CHR_t) as N_{it} -weighted averages of individual plant rates and use the X-12 ARIMA CENSUS procedure to seasonally adjust these aggregate rates. We then compute their cyclical component employing an HP-filter with a smoothing parameter of 100,000 (following Shimer, 2005). We use the HP filtered seasonally adjusted West-German unemployment rate as a cyclical reference series.¹¹

2.3 Aggregate Summary Statistics

Table 1 shows summary statistics for several aggregate labor market flow rates (Figure A6 in Appendix B displays the corresponding time series). The time average quarterly job creation and destruction rates are both around 3.7%.¹² Worker flows are substantially larger. The time average quarterly hiring and separation rate are both around 7.0%. Thus, worker turnover in Germany is about twice as high as job turnover. The resulting worker churn rate is around 6.7% of employment each quarter. Thus, in a statistical sense, 47% of worker turnover is due to worker churn while the rest is due to job flows.

⁹See also Lazear and Spletzer (2012) and Lazear and McCue (2017) who also study worker churn in the United States.

¹⁰See Davis et al. (1996) for a thorough discussion of these rates.

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

¹²It is by chance that the time average flow rates are almost equal. Employment has grown in Germany over the last decades, yet, this net growth was essentially all in part-time employment.

Table 1: Job and Worker Flows and the Churn Rate

	<i>Mean</i>	<i>Std</i>	<i>AC</i> (1)	Correlation with U_{t+j}				
				$j = -2$	-1	0	$+1$	$+2$
JCR	3.65%	0.30%	0.54	0.18**	0.07	-0.05	-0.18**	-0.29***
JDR	3.65%	0.36%	0.40	-0.03	0.05	0.14	0.22***	0.28**
EJTR	6.87%	0.40%	0.51	0.20**	0.17*	0.13	0.10	0.08
HR	7.02%	0.58%	0.81	-0.26***	-0.40***	-0.53***	-0.64***	-0.72***
SR	7.02%	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: This table displays the properties of a number of aggregate labor market flow rates. The second 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, *EJTR*: excess job turnover rate = $JCR + JDR - |JCR - JDR|$, *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. West German plants only with quarterly frequency, 1975Q1-2014Q4.

Appendix A provides a detailed comparison between the German and U.S. data. Although job and worker flows are substantially larger in the United States, the relative magnitudes between job and worker flows and the cyclical patterns of those rates are remarkably similar in both countries. Davis et al. (2012) provide statistical models that link U.S. job and worker flows, and we show in Appendix A that the same statistical relationships also hold in the German data (see also Bellmann et al. (2018), a paper which focuses on job and worker flows at the annual frequency, using mostly survey data). A stark difference to the U.S. is that the German flow rates do not show a significant downward trend but are stable despite demographic and institutional changes.

Returning to Table 1, the churn rate has a negative contemporaneous correlation with the unemployment rate (-0.77) and fluctuates substantially over the cycle: as Figure A6 in Appendix B shows, during boom times, it is about 3 percentage points higher than during times of recessions. With an autocorrelation of 0.92, it is also highly persistent. To better understand the sources of this persistent procyclical churn, the table also displays time series properties of the underlying job and worker flows. 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, leading the latter. In contrast, the job destruction rate moves with the unemployment rate, again leading it. The excess job turnover rate defined as $EJTR = JCR + JDR - |JCR - JDR|$, i.e., job turnover in excess of net employment growth, is close to acyclical. Worker flows are more persistent than job flows and more volatile. Moreover, both worker flow rates are procyclical. Taken together, in a boom

(recession), job creation is high (low) and job destruction is low (high). However, worker flows, both hirings and separations, stay high (low) longer throughout the boom (recession) leading to procyclical churn but no procyclical excess job reallocation. The hiring rate is more procyclical than the job creation rate because the separation rate is procyclical, too. Thus, procyclical churn results from hirings rising more than job creation, and from separations rising, yet job destruction declining during booms.

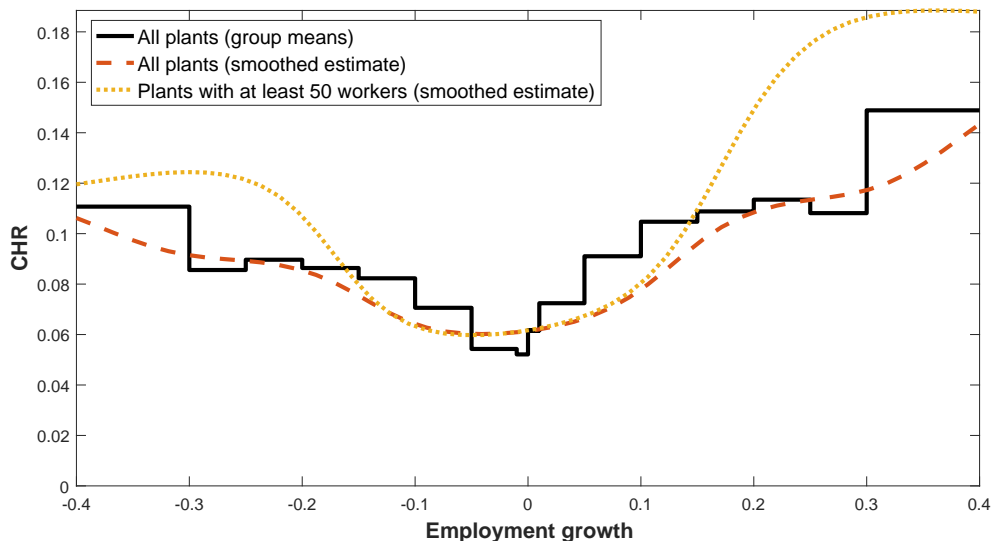
In the remainder of the paper, we will provide the following account of the procyclicality of the churn rate: starting from the plant level, we establish that churn is rising in the absolute value of employment growth and that churn is mainly the result of (the somewhat stochastic correction of) mismatches in the labor market. We further argue that these facts can be interpreted as manifestations of stochastic idiosyncratic separation shocks in conjunction with a time-to-hire friction, which we show to be important features of the labor market, at least in a semi-structural sense. From this plant-level view, procyclical churn is then the result of more separations in booms, which we show (1) to occur almost uniformly across the employment growth distribution and (2) to come almost exclusively from job-to-job transitions. Procyclical churn thus reflects a more active reshuffling of workers towards individually better matches in booms.

3 Worker Churn and Employment Growth in the Cross Section

Conceptually, churn can occur for two reasons. First, it can occur because plants reorganize (reorganization hypothesis). That is, they hire new workers while separating from others, because they desire to change the composition of their workforce. For example, they may wish to change their skill composition; or, they may wish to make their workforce younger. Alternatively, churn can also occur because plants and workers learn about their match quality which is uncertain at the time of hiring, and with the uncertainty being resolved somewhat stochastically over time (uncertainty hypothesis).

Both of these hypotheses could be related to plant growth. For example, under the reorganization hypothesis, churn may be linked to large changes in (optimal) plant size when a plant replaces production workers with IT specialists and robots, and, as a result, shrinks. In a different version of the story, the plant may grow through reorganization because it adds sales representatives as the robots produce larger quantities. In this view, churn results from replacing workers with some no longer desired ex-ante known characteristics by workers with other ex-ante known and now desired characteristics. By contrast, under the uncertainty hypothesis, plants might simply shrink because more than expected workers or plants realize that a match they deemed valuable no longer is so. Vice versa, plants might grow when less than the expected number of workers separate. Hence, under the uncertainty hypothesis, separations are stochastic, and plants, at least in the short run, cannot control the size of their workforce exactly when it takes time to replace separations.

Figure 1: Churn Rates and Plant-level Employment Growth



Note: This figure displays the average churn rate of a plant as a function of its employment growth rate. The red dashed line is estimated by an N_{it} -weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The yellow dotted line is estimated the same way for those plants with more than 49 employees. The black straight line displays the N_{it} -weighted group average churn rate for 17 discrete plant employment growth categories (see Table A3 in Appendix B for the exact definition of these plant-level employment growth bins). West German plants only with quarterly frequency, 1975Q1-2014Q4.

Figure 1 shows non-parametric estimates of the relationship between plant-level churn rates and employment growth rates. This relationship is V-shaped. The larger the absolute rate of employment growth, the larger is the churn rate, which is at least six percent of employment along the entire employment growth distribution. The figure also shows that this pattern is not driven by small plants, where small absolute worker flows may imply large flow rates. The pattern is even somewhat more pronounced for plants with at least 50 workers.¹³

We have subjected our core empirical finding of the V-shaped churn-employment growth relationship to a battery of robustness checks. First, we find the V-shape also separately for all four decades in our sample, see Figure A7 in Appendix B.¹⁴ Second, it is also present for all four quarters separately or when we adjust the plant-level employment growth rates seasonally, see Figure A12 in Appendix C. Third, Figure A13 in the same appendix shows that the V-shape is also present in yearly data. Finally, Figure A8 in Appendix B shows that controlling for plant fixed effects in quarterly employment growth does not alter the V-shaped relationship.

¹³Our results focus on growth rates between -0.4 and 0.4 which represent more than 96% of total employment. In other words, we abstract in this paper from exiting, near-exiting, or entering plants.

¹⁴We define the decades as 1975-1984, 1985-1994, 1995-2004, 2005-2014. This means in particular that the Hartz labor market reforms in the first half of the 2000s appear not to have had any impact on the plant-level churn-employment growth nexus.

To better understand this V-shaped relationship, note that we can write Equation (1) as:

$$CH_{it} = \begin{cases} 2H_{it} & \text{if } E_{it-1} \geq E_{it} \\ 2S_{it} & \text{if } E_{it} \geq E_{it-1}. \end{cases} \quad (3)$$

Hence, through the lens of this formula, a V-shape implies that, while growing plants hire a lot, they also separate from a large share of workers. Indeed, they separate from a larger fraction of their workers than plants with a constant workforce. Vice versa, rapidly shrinking plants hire more workers than plants with a constant workforce.

Next, to test the reorganization hypothesis, we use a number of worker-related features of our plant-level data set and investigate how much of the churn rate by plant-level employment growth category occurs within workers' task¹⁵, education¹⁶, and age¹⁷ groups. Let a worker's observable characteristic be indexed by $m = 1 \dots M$. We have that churn at plant i in period t is:

$$\begin{aligned} CH_{it} &= H_{it} + S_{it} - (JC_{it} + JD_{it}) = \sum_{m=1}^M [H_{itm} + S_{itm}] - (JC_{it} + JD_{it}) \\ &= \sum_{m=1}^M [H_{itm} + S_{itm} - (JC_{itm} + JD_{itm}) + (JC_{itm} + JD_{itm})] - (JC_{it} + JD_{it}) \\ &= \sum_{m=1}^M CH_{itm} + \left[\sum_{m=1}^M (JC_{itm} + JD_{itm}) - (JC_{it} + JD_{it}) \right], \end{aligned} \quad (4)$$

where CH_{itm} is churn arising from plants replacing a worker within a particular observable category with another worker with the same characteristic, and the second term in brackets represents between-group churn. Dividing by N_{it} gives:

$$CHR_{it} = \frac{\sum_{m=1}^M CH_{itm}}{N_{it}} + \frac{\sum_{m=1}^M (JC_{itm} + JD_{itm}) - (JC_{it} + JD_{it})}{N_{it}}, \quad (5)$$

i.e., a plant's churn rate is the sum of the churn occurring within groups (the first term) and the churn occurring from shifting around workers between groups (the second term).

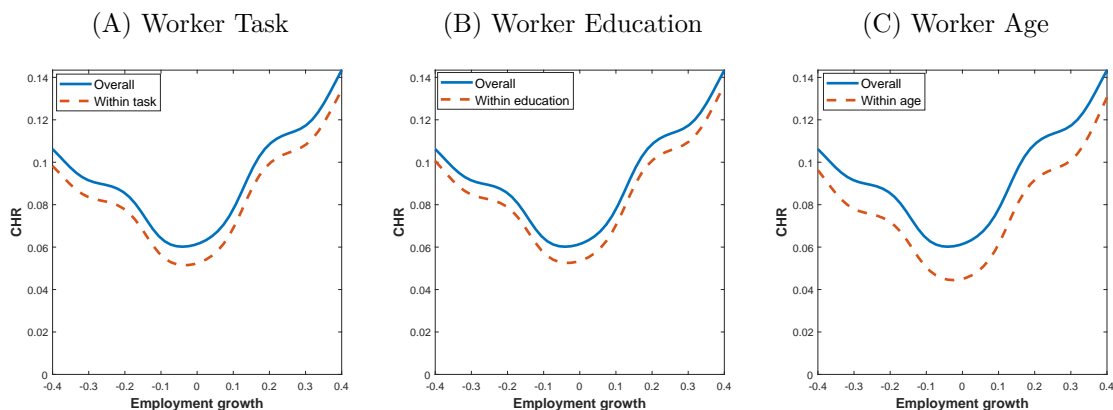
Figure 2 plots the churn rates by employment growth category arising from within-task/education/age-group churn. It shows that almost all churn and the entire V-shaped relationship with employment growth arises from within-group churn. Thus, churn does not appear to come mainly from plants replacing production workers with IT specialists, or workers with little formal training with workers with a university degree, or older workers, perhaps separating predictably through retirement, with younger workers. Moreover, comparing wages of the newly hired and separating workers, we find that the higher the churn

¹⁵Category 1: Agricultural occupations, elementary manual occupations, elementary personal services occupations, elementary administrative occupations. Category 2: Skilled manual occupations, skilled services occupations, skilled administrative occupations. Category 3: Technicians, associate professionals. Category 4: Professional occupations, managers.

¹⁶Category 1: Workers without formal vocational training. Category 2: Workers with formal vocational training and/or higher education entrance qualification. Category 3: Workers with a university degree.

¹⁷Category 1: 15-29 years old. Category 2: 30-44. Category 3: 45-59. Category 4: 60+.

Figure 2: Churn Rates and Plant-level Employment Growth By Worker Characteristics



Note: This figure displays the average churn rate of a plant as a function of its employment growth rate by worker characteristics, estimated by an N_{it} -weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). Overall: total churn rate. Within task/education/age: aggregate churn due to churn within worker groups. Panel A measures skills by tasks, panel B measures skills by education; panel C groups workers by age. West German plants with a quarterly frequency, 1975Q1-2014Q4.

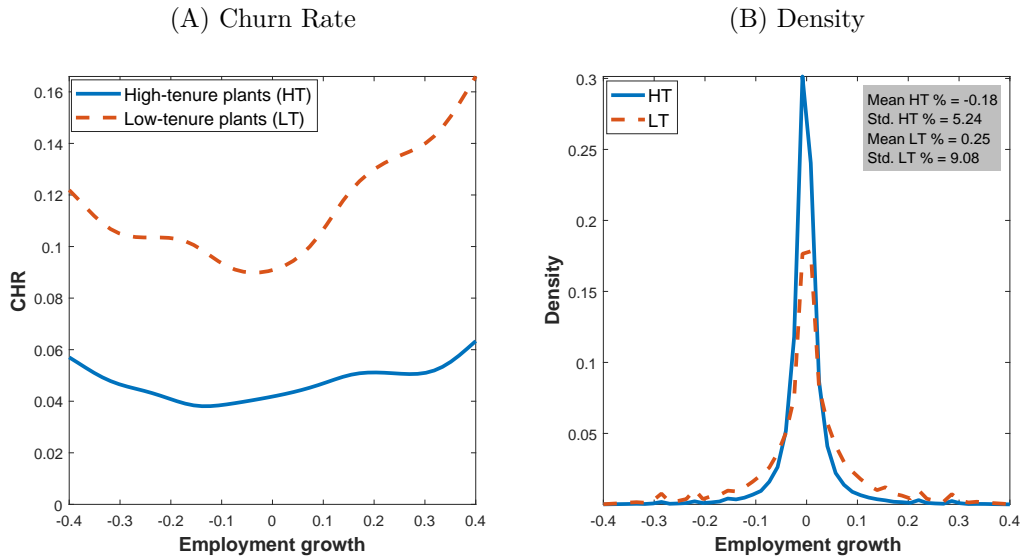
rate of a plant, the more similar are the wages of new hires and separating workers. This can be seen in Figure A9 in Appendix B. Taken together, this body of evidence makes it unlikely that reorganization, the replacement of workers with observable and ex-ante known characteristics, or rejuvenation are the key drivers of churn.

If, by contrast, churn reflects learning over ex-ante unknown match quality—be it because the worker is unproductive in the match or because the worker finds a better match elsewhere—then churn should be larger at plants with putatively higher uncertainty about their existing match qualities and, thus, the number of workers separating. Moreover, when plants are not able to predict exactly how many workers will separate in the short-run, plants may actually be shrinking because more workers separated from the plant than expected ex-ante and the plant is not able to re-hire the desired amount of workers. Similarly, plants may be growing in the short run to make up for past separations and, with that, have a large inflow of new workers with uncertain match qualities. That is, under the uncertainty hypothesis, the relationship between churn and employment growth should depend on the amount of the underlying uncertainty. Incidentally, this second view of churn is consistent with the bulk of churn occurring within worker groups with the same observable characteristics.

To gauge the importance of this channel, we split plants into two groups: a “high-tenure” group for which the share of workers with more than one year of tenure is above the plant-level median and a “low-tenure” group for which the share is below the plant-level median. Figure 3A shows that for each employment growth category, churn is twice as high at “low-tenure” plants with putatively the highest uncertainty about separations. What is more, the V-shape of churn is more pronounced at plants with many low-tenure workers. That is, when uncertainty about match quality is high, plants achieve rapid employment growth despite many workers leaving, and plants downsize considerably despite many hires.

Figure A10 in Appendix B shows that “low-tenure” plants have higher average churn and a more pronounced V-shaped churn pattern even conditional on plant age and plant size.

Figure 3: Churn Rates and Plant-level Employment Growth By Tenure Composition



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 two categories based on whether their share of workers with more than one year of tenure is below or above the plant-level median. Panel B displays for these two groups the density of employment growth estimated by a kernel smoother (Gaussian kernel with a 0.005 bandwidth). West German plants with a quarterly frequency, 1975Q1-2014Q4.

This analysis by plant age demonstrates that part of the match uncertainty causing churn is related to young (not older than 5 years) plants which are presumably more inexperienced in finding good matches. However, plant age explains only a minor fraction of the difference between low- and high-tenured workers. Even young plants have less and relatively flat churn rates across the employment growth rate distribution when they are high-tenure plants. By contrast, even old and presumably experienced plants have high and highly V-shaped churn rates when they have many workers that have been with them under a year. In sum, both sides of the match, plants and workers, appear to be contributing to the match uncertainty related to workers with a short tenure.

The second moment of the employment growth distribution also supports the view that uncertain separations due to a stochastic resolution of initially uncertain match qualities might be an important driver of short-run employment fluctuations. Figure 3B shows that the employment growth distribution of “low-tenure” plants is substantially more dispersed (with a similar mean) than that of “high-tenure” plants. This suggests that, in the short run, workforce levels are highly uncertain for “low-tenure” plants.

4 Cross-sectional Worker Churn as a Result of Stochastic Separations

In this section, we develop a simple formal framework to illustrate how uncertainty about the number of workers separating from a plant leads to short-run plant-level employment dynamics and churn that are consistent with the data. The model has two key elements:

separations are idiosyncratic and in part stochastic, and a plant cannot react intra-period to these separations (“time to hire”).¹⁸ We note that in this paper we stop short of attributing a deep structural meaning to these two elements—mismatch and thus separations, in reality, will occur for a variety of reasons on both the worker and the employer side—but we do argue that an environment that looks like stochastic separations leading to workforce planning mistakes is required to match the data.

To be specific, consider a model where plant i has the following decreasing returns-to-scale production function in employment:

$$Y_{it} = zE_{it}^\alpha, \quad (6)$$

where E_{it} is the plant’s (end-of-quarter t) employment level, z is (for simplicity fixed) plant-level productivity, and α (with $0 < \alpha < 1$) is the curvature of the production function. The plant chooses to actively adjust employment $\Delta_{E_{it}}^a$ to maximize each period:

$$\begin{aligned} \max_{\Delta_{E_{it}}^a} & \left\{ \mathbb{E}_{t-1} \{ zE_{it}^\alpha - wE_{it} \} \right\} & (7) \\ E_{it} & = (1 - s_{it})(E_{it-1} + \Delta_{E_{it}}^a) \\ s_{it} & = \min\{\exp(\tilde{s}_{it}), 1\} \\ \tilde{s}_{it} & = (1 - \rho_s)\mu_s + \rho_s\tilde{s}_{it-1} + \epsilon_{it}, \quad \epsilon_{it} \sim N(0, \sigma_s^2), \end{aligned}$$

where w is the (exogenous) wage and \mathbb{E}_{t-1} is an expectation operator with the associated information set $\{E_{it-1}, \tilde{s}_{it-1}\}$. \tilde{s}_{it} is an autocorrelated stochastic latent variable with autocorrelation ρ_s that determines the separation rate s_{it} . These separations occur after the plant decides on its current-period employment adjustment. Thus, adjustment decisions take only into account the expected amount of separations.¹⁹ In this way, plants can make planning mistakes in their employment adjustment decisions. To sum up, the timeline of events at a plant within a period is: active employment adjustment (that is, active hiring of or separation from workers), stochastic separations, production, and wage payment.

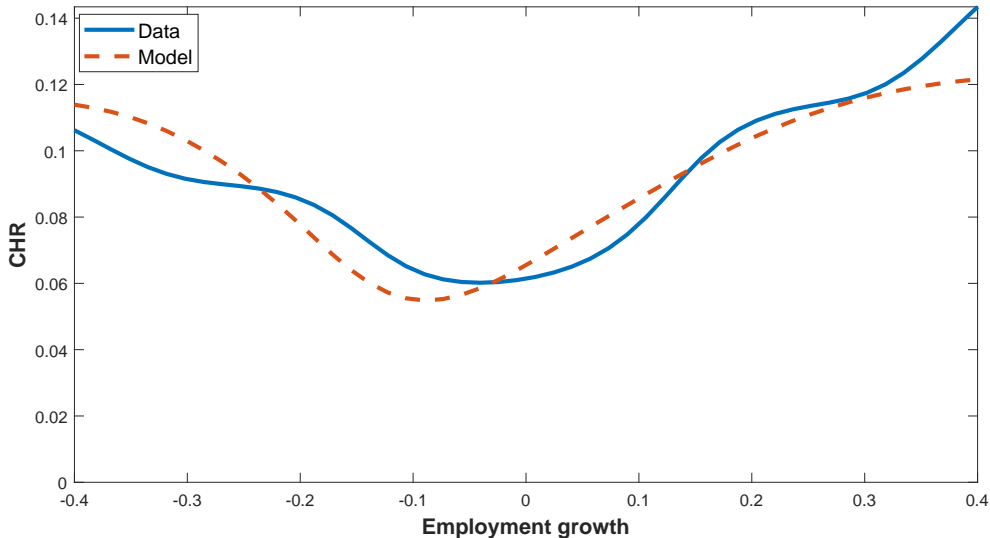
In terms of worker flow accounting, we follow the same procedure as in the data. That is, when $\Delta_{E_{it}}^a > 0$, we count as new hires only those that do not separate within the same period, i.e., $H_{it} = (1 - s_{it})(\Delta_{E_{it}}^a)^+$; as for separations, in this case: $S_{it} = s_{it}E_{it-1}$. When $\Delta_{E_{it}}^a < 0$, $H_{it} = 0$, and we count the active adjustment as additional separations in the model, i.e., $S_{it} = s_{it}(E_{it-1} + \Delta_{E_{it}}^a) + (\Delta_{E_{it}}^a)^-$.

To quantify the amount of uncertainty about separations, we calibrate the model as follows: We set the returns-to-scale parameter, α , to 0.6, normalize the wage to $w = 1$,

¹⁸Appendices D.1.1 and D.1.2 show that stochastic separations and time to hire are essential. Standard idiosyncratic plant-level productivity shocks with or without employment adjustment costs cannot generate the observed cross-sectional facts about worker churn. Of course, this is not to say that a deeper modeling of productivity shocks in a richer model environment might not give rise to what we identify as stochastic separation shocks.

¹⁹The baseline model assumes separations are homogeneous for workers with different tenure. Appendix D.2.1 shows that incorporating a downward-sloping hazard for the separation rate in worker tenure leaves the baseline results almost unaffected.

Figure 4: Churn Rates and Employment Growth in a Model with Time-to-Hire and Stochastic Separations



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 model.

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 in Figure 1) 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 yields $\mu_s = -3.07$, $\sigma_s = 0.85$, and $\rho_s = 0.36$.

Figure 4 compares the V-shaped churn-employment growth pattern of our calibrated model with that in the data. The figure shows that the model is able to replicate this V-shaped pattern of the churn rate rather well. Particularly, churn is largest at rapidly shrinking and rapidly growing plants. Through the lens of Equation (3) above, this means that hiring rates increase the faster plants shrink, and separations increase the faster plants grow.

Rapidly growing plants tend to be those which have experienced high unexpected separations last period and, thus, start the period well below their optimal size. These plants, because of autocorrelated separation rates, also tend to have high separation rates in the current period. Such plants, because of below-target employment levels (and because of high expected separations), will now hire many workers in order to return to their optimal employment size. They have, thus, both large positive employment growth and high churn.

Next, looking at plants with moderate positive employment growth, recall from Equation (3) that these must be plants with low separation rates to rationalize their low churn. Some plants aim to grow moderately because they ended the previous period close to but slightly below their optimal size. For them to actually grow moderately, these plants must have experienced a close to expected separation shock today. Since the expected separation rate

is approximately 7%, this implies few separations and their churn is small. Other plants grow moderately because they did not aim to grow but their separation rate shock today was smaller than expected, and, again, since the expected separation rate is approximately 7% on average, this surprise growth can only be moderate. There is a third group of plants that grow moderately because they wanted to grow a lot and are surprised by high separation rates. Their churn is large, but this group among the moderately growing plants is small because belonging to this group requires two large separation rate shocks in a row, a rare event.

Conversely, by the same argument as above, there is among the moderately shrinking plants a relatively large fraction of plants that have experienced a smaller than expected separation rate in the previous period. These plants want to shrink in the current period so they hire little, if at all, and according to the logic of Equation (3), their churn is small. This effect is amplified by the autocorrelation in separation rates which leads to small expected separation rates for these plants with small past separation rates and, hence, little need for replacement hiring.

Finally and in contrast to moderately shrinking plants, rapidly shrinking plants are those affected by a large separation rate surprise in the current period; in short, they shrink rapidly but do not aim to do so. The typical such plant is a plant that ended the previous period close to employment target because separation surprises affect plants independently of their employment level (and most plants have close to target employment levels). These plants are also plants that hire in the current period because of expected separations. Therefore, they have more churn than the typical moderately shrinking plant which hires little, if at all. Taken together, the model thus implies a V-shaped relationship between employment growth and churn; and churn is lowest at moderately shrinking plants, consistent with Figure 4.

What is more, the calibrated model is also broadly consistent with the dispersion of the employment growth rate distribution. As column “Whole sample” in Table 2 shows, the cross-sectional standard deviation of plants’ employment growth rates is 9.23% in the model (compared to 8.09% in the data).

Finally, the calibrated model is also consistent with other non-targeted plant-level moments. First, the churn rate is positively autocorrelated both in model and data (0.38 and 0.11, respectively). Key for this is the autocorrelation in separation rates (0.13 and 0.05, respectively). Moreover, the model features a negative autocorrelation of employment growth (-0.48 compared to -0.15 in the data). This means that plants indeed partly reverse past employment growth, which means that, through the lens of the model, they reverse past planning mistakes.

The autocorrelation of employment growth is somewhat too negative in the baseline model. Appendix D.2.2 adds to the baseline model convex employment adjustment costs which makes the autocorrelation less negative (-0.37). What is more, the autocorrelation in the data is a mixture of plants fluctuating around their optimal size and plants that have recently entered the economy. Kaas and Kircher (2015) show that the latter feature positively correlated employment growth as they adjust to their optimal size. Our model

Table 2: Parameters and Moments of the Calibrated Models

	Whole sample	“Low tenure”	“High tenure”
Productivity level, z	4.49	4.48	4.78
Location parameter, μ_s	-3.07	-2.72	-3.87
Dispersion parameter, σ_s	0.85	0.74	0.61
Autocorrelation, ρ_s	0.36	0.36	0.36
Implied moments of separation shocks			
Average stoch. separation rate, $\mathbb{E}(s)$ %	6.97	8.97	2.79
Std. stoch. separation rate $\sigma(s)$ %	7.74	8.28	1.98
Implied economic outcomes			
Output loss %	0.10	0.12	0.005
Std. of employment growth %	9.23	10.29	2.78
	(Data: 8.09)	(Data: 9.08)	(Data: 5.24)

Note: This table shows parameters and moments of the calibrated models. z : productivity. 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. $\mathbb{E}(s)$: implied mean of the stochastic separation rate. $\sigma(s)$: implied standard deviation of the stochastic separation rate. *Std. of employment growth*: the cross-sectional standard deviation of plant-level employment growth rates. *Output loss*: per-period percent of lost output due to separation uncertainty. “Low tenure”: plants whose share of workers with more than one year of tenure is below the plant-level median. “High tenure”: plants whose share of workers with more than one year of tenure is above the plant-level median.

rather highlights the behavior of continuing plants. Consistent with this idea, in the data, continuing plants (those that do not enter or exit within a 2-year time window) display an autocorrelation of employment growth that is indeed more negative than that of all plants (-0.29).

In addition to predicting an overall negative autocorrelation of employment growth, our model also predicts that the strength of this correlation varies systematically with the level of employment growth. Plants that shrink a lot in period t do so because of a large separation rate shock and, hence, ex-post, have made a planning mistake in their recruitment. Therefore, they typically undo this in the following period $t + 1$ and, thus, employment growth between periods t and $t + 1$ is strongly negatively autocorrelated for those plants.²⁰ Plants that have positive employment growth end up typically close to their optimal employment level and, hence, their employment growth is, by-and-large, uncorrelated between periods t and $t + 1$. In Figure A11 in Appendix B, we show that this pattern indeed holds both in the model and data.

Finally, column “Whole sample” in Table 2 summarizes the implied moments of the stochastic separation rate. Separation uncertainty is substantial. With a 90% confidence

²⁰In line with this fact, Figure A8 in Appendix B shows that the V-shape disappears when we relate quarterly churn to long-term employment growth.

probability, stochastic separation rates lie between 1.4 and 14.8 percent on a quarterly basis. This uncertainty implies a misallocation of labor, hence an output loss, in the economy, with some plants having too many and other plants having too few workers.²¹ When we compare the baseline equilibrium to an economy without separation uncertainty, which implies in this model that all plants have the same amount of workers, we find per-period output losses of 0.1 percent of total output as a result of this worker misallocation.²²

The columns “Low tenure” and “High tenure” in Table 2 illustrate how more match uncertainty translates into more uncertain separation rates. Here we calibrate the model to the same moments, but of the corresponding plant groups defined in Figure 3 in the previous section.²³ For simplicity, we fix ρ_s at the value for the whole sample (0.36). Due to the higher churn rate of “low-tenure” plants, we calibrate their average separation rate shock, μ_s , to be larger. Moreover, their more pronounced V-shape leads to a higher calibrated dispersion of their separation rate shock, σ_s . Larger and more dispersed shocks imply in the model, in line with the data, a significant fanning out of the employment growth distribution for “low-tenure” plants: 10.29% versus 2.78% for “high-tenure” plants. More uncertain separations also imply that planning mistakes are larger for those plants. As a result, the output loss arising from this separation-shock-induced worker misallocation is more than twenty times larger at “low-tenure” plants than at “high-tenure” plants.

Appendix D.1.1 shows that in a model with only productivity shocks (with or without convex labor adjustment costs), churn merely occurs at plants that have not suffered too large negative productivity shocks and replace exogenously leaving workers. Even if we match this model to the same moments as our model here, we find that the churn rate as a function of plant-level employment growth is counter-factually largest at non-adjusting plants.

To be clear, we do not mean to say that productivity shocks, as for instance in Hopenhayn (1992), are not important to understanding plant-level employment growth data. Nevertheless, our analysis does suggest that shocks leading to stochastic separations, together with a time-to-hire friction, are important to understanding the joint behavior of plant growth and churn. What is more, stochastic separations with a time-to-hire friction can also explain important moments of observed plant-level employment growth dynamics, making them of independent interest beyond the topic of plant-level churn. Finally, these stochastic separation shocks in conjunction with the time-to-hire friction lead to economically meaningful output losses.

²¹To be precise, this output loss is not meant to be the cost of worker mobility. It is purely a cost of the *unpredictability* of separations leading to an uncertain employment level at the plants. Conversely, this cost does not measure a potential output gain from removing worker mobility; it rather measures the gain from letting a plant know about separations before their active employment adjustment decisions. It is a statistic that is, thus, only meant to capture the extent of the information friction with a single number. There is no easy and straightforward policy intervention that would allow a policymaker to eliminate this output loss.

²²While our model, with essentially Gaussian stochastic log separation rates, broadly matches the second moment of the employment growth distribution in the data, it features too little kurtosis compared to the data. This means that large separation shocks and thus large plant-level output losses from uncertain labor mobility are relatively rare. We have also calibrated a Gaussian-mixture model that allows us to better match the fat tails of the employment growth distribution. In that calibration, the output loss increases to 0.4%.

²³Note that we use our model here as a measurement device, not to explain endogenously the existence of low-tenure and high-tenure plants. The total aggregate separation rate for low-tenure plants is 9.00% in the data, and that for high-tenure plants is 4.10%.

Appendix D.1.2 shows that a hybrid model with stochastic separations and convex labor adjustment costs but no time-to-hire friction cannot match the churn pattern in the data, either. Despite adjustment costs, the most (partial) employment adjustment still occurs in the same period when separations take place, that is, shrinking plants have systematically (but counterfactually) more churn than growing plants.

5 Worker Churn and the Business Cycle

In this section, we return to the procyclicality of the aggregate churn rate documented in Table 1, and show that churn increases across the entire employment growth distribution during booms relative to recessions. What is more, this parallel shift of the churn-employment growth nexus explains essentially all of the cyclical dynamics of the aggregate churn rate. Cyclical shifts of the employment growth distribution are nearly irrelevant for aggregate churn dynamics. Through the lens of our model, this almost uniform level shift of the churn-employment growth nexus is explained by an on average higher separation shock in booms that is less uncertain than in recessions. This translates into an essentially acyclical behavior of the dispersion of the plant-level employment growth distribution, as we observe in the data, and, thus, consistent with Table 1, to no additional (excess) job turnover in booms. Digging deeper, we will also show that this procyclical increase in churn has an almost one-to-one link with procyclical job-to-job transition rates. That is, the entire rise in job-to-job transitions during booms leads ultimately to worker churn, but, again, not to (excess) job turnover.

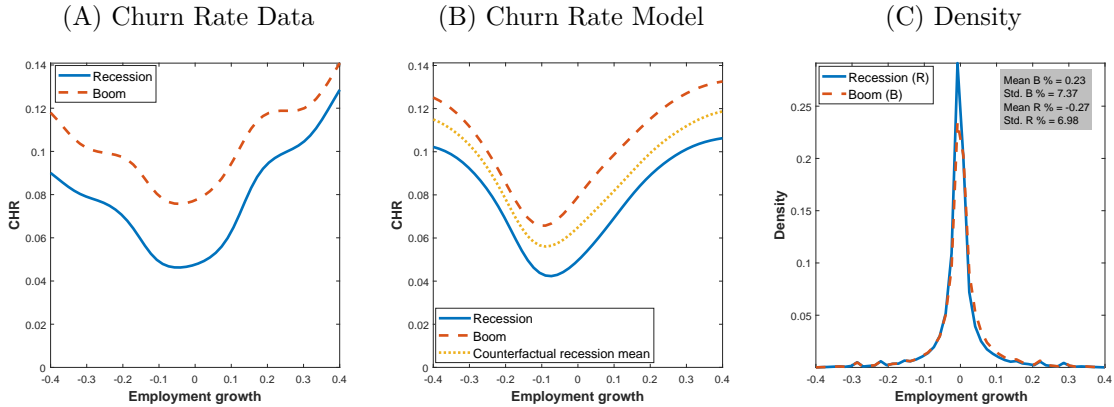
Figure 5A shows the cyclical dynamics of the churn rate as a function of plant-level employment growth rates in the data. To this end, we pool the ten quarters with the lowest cyclical unemployment rate (boom) and the highest cyclical unemployment rate (recession) in our sample. Table A3 in Appendix B displays additional summary statistics of the cyclical dynamics of the churn rate for different employment growth categories. Both the figure and the table show that, across the employment growth distribution, churn moves counter to the unemployment rate. Moreover, in absolute values, the rise during booms is similar across the distribution. Using the statistical techniques from Davis et al. (2012), Appendix A.2 shows that this parallel shift in the distribution explains virtually all of the procyclical churn. Put differently, unlike for worker flows, shifts of the employment growth distribution towards more rapidly adjusting plants are not an important contribution to procyclical churn.

What does this parallel shift imply for uncertainty about separations from the plant's perspective? Table 3 shows the results when we re-calibrate the model from Section 4 separately for boom and recession times.²⁴ Again, we fix ρ_s at the value for the whole sample (0.36). In the data, and consistent with Equation (3), the typical plant faces more separations during a boom, but it replaces these additional separations by new hires.²⁵ Our model interprets this as separations becoming larger on average in booms but also more predictable.

²⁴As in the low/high-tenure case, we use our model again as a measurement device. There are no actual business cycle dynamics in the model.

²⁵The total aggregate separation rate for boom times is 7.59% in the data, and that for recession times is 6.41%, compared to the 7.02% over the whole sample.

Figure 5: Churn Rates and Employment Growth—the Cyclical Dimension



Note: Panel A displays the average churn rate of a plant in the data 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 average churn rate in the ten quarters with the highest HP(100,000)-filtered unemployment rates in the sample (recession): 1983Q1-1983Q4, 1984Q3, 2005Q1-2006Q1. The red dashed line is the average churn rate in the lowest HP(100,000)-filtered unemployment rates in the sample (boom): 1991Q1-1992Q3, 1979Q4-1980Q2. Panel B displays the model-implied 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), calibrated to recession quarters (blue solid line) and boom quarters (red dashed line). The yellow dotted line is the churn-employment growth nexus for a counterfactual calibration of the model, calibrated to match the (lower) average separation rate shock of the recession-calibration, but the (also lower) standard deviation of the boom-calibration. Panel C displays for recession and boom times, respectively, the density of employment growth in the data estimated by a kernel smoother (Gaussian kernel with a 0.005 bandwidth). West German plants with a quarterly frequency, 1975Q1-2014Q4.

Why is the difference between booms and recessions characterized by higher average separation shocks but lower uncertainty of said shocks in booms? Figure 5B helps building the intuition. It displays the model-implied average churn rate of a plant as a function of its employment growth rate for both boom and recession times. It also displays the churn-employment growth nexus for a counterfactual calibration of the model, calibrated to match the (lower) average stochastic separation rate, $\mathbb{E}(s)$, of the recession calibration, but the (also lower) standard deviation, $\sigma(s)$, of the boom calibration (see Table 3 for the numbers). It turns out that this counterfactual calibration essentially already matches the total aggregate separation rates in recession times (6.10%), but, clearly, still produces too much churn compared to the data (as seen in Figure 5B). To further reduce churn, while keeping the total aggregate separation rate basically the same, separations have to become more uncertain from the viewpoint of the plant. Higher separation uncertainty renders what is worker turnover job turnover: An unexpected separation, in the model, produces a job destruction. Since churn is the difference between worker flows and job flows, higher uncertainty, for a given average separation rate, reduces churn. This is an important general insight of our analysis: matching jointly aggregate separations and churn identifies both the mean and the uncertainty of the separation rate shock. And the data, viewed through the lens of our model, suggest a countercyclical uncertainty for these shocks.

Indeed, separations have a larger dispersion during recessions. Losing up to 17 percent of the workforce through a separation rate shock is an event that occurs in the two model

Table 3: Parameters and Moments of the Calibrated Models—the Cyclical Dimension

	Whole sample	Recession	Boom
Productivity level, z	4.49	4.46	4.62
Location parameter, μ_s	-3.07	-3.31	-2.89
Dispersion parameter, σ_s	0.85	0.98	0.71
Autocorrelation, ρ_s	0.36	0.36	0.36
Implied moments of separation shocks			
Average stoch. separation rate, $\mathbb{E}(s)$ %	6.97	6.28	7.52
Std. stoch. separation rate $\sigma(s)$ %	7.74	8.62	7.11
Implied economic outcomes			
Output loss %	0.10	0.15	0.08
Std. of employment growth %	9.23	9.20	9.11
	(Data: 8.09)	(Data: 6.98)	(Data: 7.37)

Note: This table shows parameters and moments of the calibrated models. z : productivity. 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. $\mathbb{E}(s)$: implied mean of the stochastic separation rate. $\sigma(s)$: implied standard deviation of the stochastic separation rate. *Std. of employment growth*: the cross-sectional standard deviation of plant-level employment growth rates. *Output loss*: per-period percent of output lost due to unequal plant sizes. “Recession”: the ten quarters with the highest HP(100,000)-filtered unemployment rates in the sample; 1983Q1-1983Q4, 1984Q3, 2005Q1-2006Q1. “Boom”: the ten quarters with the lowest HP(100,000)-filtered unemployment rates in the sample; 1991Q1-1992Q3, 1979Q4-1980Q2.

calibrations with a 7 percent probability in both the boom- and the recession calibrations. For any cut-off above 17 percent, the probability of the event of losing a larger fraction of the workforce is higher in the recession calibration. The additional dispersion in separation rate shocks during recessions translates into an almost 50 percent difference in output losses between booms and recessions. These countercyclical output losses are particularly interesting in light of the recent debate about the role of time-varying uncertainty in business cycles (see Bloom, 2014, for an overview). This literature typically assumes that productivity shocks are more dispersed in recessions. Here, we find that dispersed separation rate shocks imply that large and unexpected separation events are more likely in recessions and lead to output losses.²⁶

Perhaps interestingly, the higher dispersion of the separation rate shocks during recessions does not lead to a strong countercyclicality of the plant-level employment growth dispersion, as Table 3 shows. This is because both the mean and the dispersion of the separation rate shocks affect the dispersion of the employment growth rate: on average smaller separation rate shocks imply less employment growth dispersion because, on average, planning mistakes in our model are smaller; yet the increase in dispersion implies that those plants that make

²⁶To be clear, our model does not contain continual first- or second-moment business cycle shocks but our result is broadly consistent with the idea of countercyclical uncertainty at the micro-level.

a planning mistake make a particularly large one. Similarly, in the data, we observe very little change in the employment growth distribution between boom and recession times. The distribution shifts right, but the dispersion changes little (see Figure 5C). Taken together, this implies that quarter-to-quarter job reallocation changes little over the business cycle, consistent with the evidence in Table 1. Put differently, the increase in worker turnover during booms does not lead to an increase in job turnover but pure churn.

Next, we investigate from which sources this procyclical churn arises. Conceptually, there are two possibilities. Plants may hire (separate) more workers from (to) other plants during booms. Alternatively, plants may increase hiring and separations through the non-employment pool. To distinguish the two cases, we use information on 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 into those resulting from job-to-job transitions and those resulting from non-employment transitions:

$$HR_t = JTJR_t + HR_t^{N-emp} \text{ and } SR_t = JTJR_t + SR_t^{N-emp}, \quad (8)$$

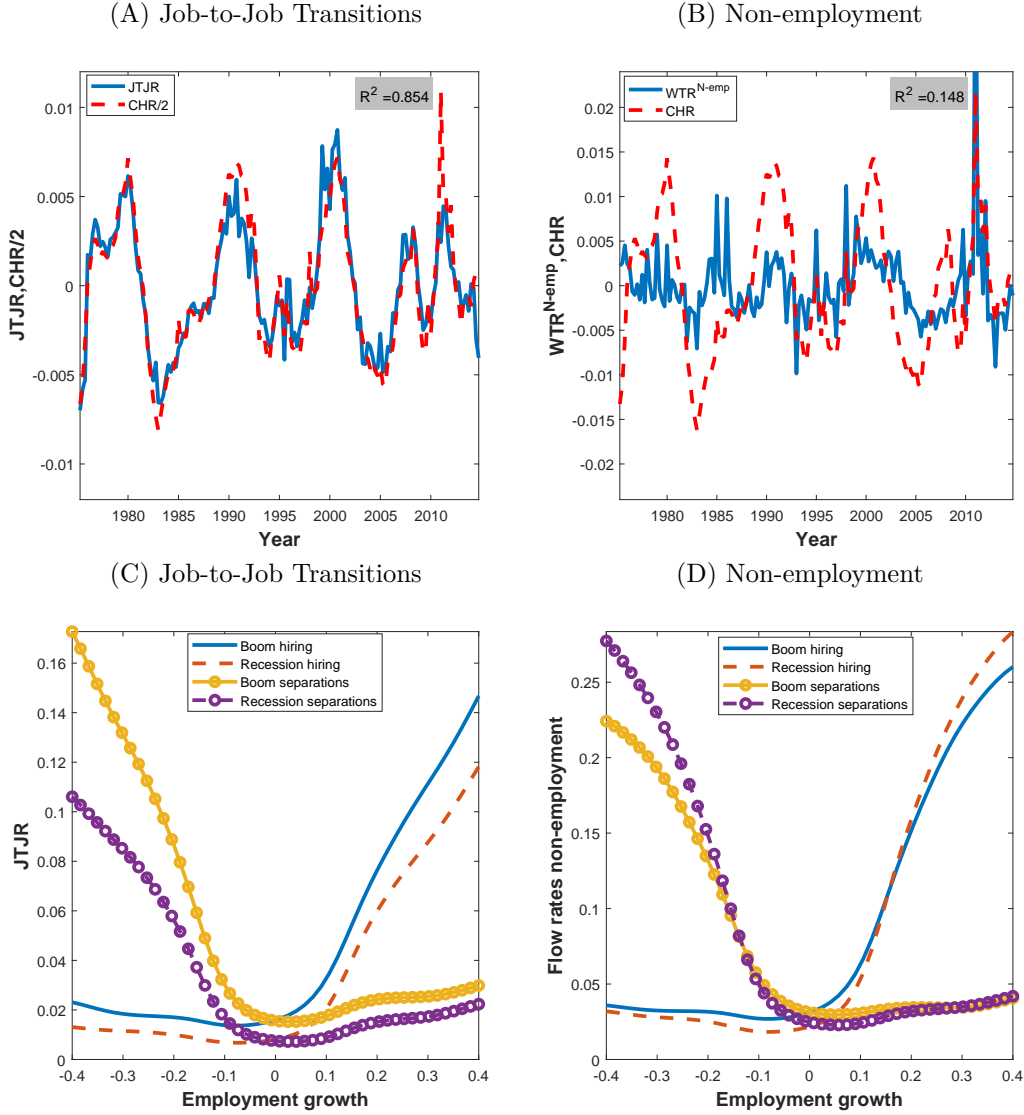
where HR^{N-emp} denotes the hiring rate from non-employment, and SR^{N-emp} denotes the separation rate into non-employment. To quantify the contribution of the cyclical movements in the job-to-job transition rate and the worker turnover rate through non-employment for procyclical churn, we write the churn rate as:

$$CHR_t = (HR_t + SR_t) - (JCR_t + JDR_t) = (HR_t^{N-emp} + SR_t^{N-emp} + 2JTJR_t) - (JCR_t + JDR_t). \quad (9)$$

Figure 6A shows that the aggregate churn rate (divided by two) is almost identical to the job-to-job transition rate. Put differently, the entire rise of the job-to-job transition rate during booms is ultimately worker churn. In fact, Figure 6C shows that during booms, the separation (hiring) rate to (from) employment shifts up in an almost parallel fashion over the entire employment growth distribution. By contrast, Figure 6D shows that the two flows into and out of non-employment, the separation rate into non-employment and the hiring rate from non-employment, do not feature such a parallel shift. Consistent with the cross section, Figure 6B shows that cyclical movements in the aggregate worker turnover rate through non-employment, $WTR^{N-emp} = (HR^{N-emp} + SR^{N-emp})$, have essentially no relationship with the aggregate churn rate. Hence, equation (9) implies that worker turnover through non-employment must equal job turnover. Indeed, as Figure 7 shows, when we split job turnover into its components, job creation and job destruction, we find that the job creation rate explains over 80% of the dynamics in the hiring rate from non-employment. Similarly, the job destruction rate explains over 85% of the dynamics in the separation rate to non-employment. It follows that the characteristic spike in the job destruction rate that is typical for the early phase of a recession results almost exclusively from workers separating into non-employment, not from workers reallocating to other plants through job-to-job transitions.

Recall from Section 2.3 that the beginning of a boom is characterized by a joint increase of the churn rate and job creation rate, the latter, as just shown, being almost equal to the

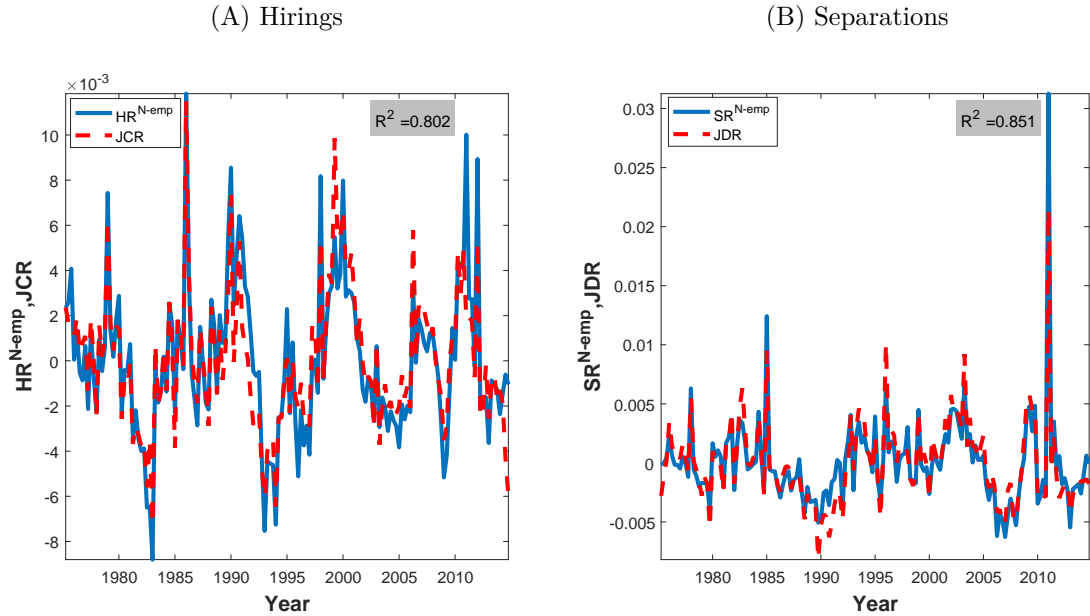
Figure 6: Churn, Worker Flows and Job-to-Job Transitions



Note: Panel (A) and (B) 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 churn rate (dashed). R^2 : share of 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 the job-to-job transition rate or the worker turnover rate through non-employment. All series are HP(100,000)-filtered. Panel (C) and (D) display plants' average flow rates given plants' employment growth rates estimated by an N_{it} -weighted kernel-smoother (Gaussian-kernel with a 0.05 bandwidth). Panel (C): The separation (hiring) rate to (from) employment in the ten quarters with the lowest HP(100,000)-filtered unemployment rates (boom) and the highest (recession). Panel (D): The separation (hiring) rate to (from) non-employment in the ten quarters with the lowest HP(100,000)-filtered unemployment rates (boom) and the highest (recession). "Boom": 1991Q1-1992Q3, 1979Q4-1980Q2. "Recession": 1983Q1-1983Q4, 1984Q3, 2005Q1-2006Q1. West German plants with a quarterly frequency, 1975Q1-2014Q4.

hiring rate from non-employment. That is, early in a boom, plants face rising separation rates due to an increase in job-to-job transitions, yet, those additional separations do not lead to a rising job destruction rate because even though some plants poach from other plants to hire, ultimately there are plants that replace these separations by hiring from non-employment. As the boom matures, hiring from non-employment reverts back to its normal

Figure 7: Aggregate Worker Flows from/to Non-Employment and Aggregate Job Flows



Note: This figure displays worker flows from/to non-employment and job flows. The blue solid lines refer to the empirical hiring rate from non-employment (left) and separation rate to non-employment (right). 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. West German plants with a quarterly frequency, 1975Q1-2014Q4.

level, yet, plants' separation rates, through job-to-job transitions, remain high. At the same time, job destruction remains low implying that plants do not shrink despite losing many workers through job-to-job transitions (and not hiring particularly many workers from non-employment). This is made possible by the same plants also hiring relatively many workers from other plants. This means that, in the aggregate, matured booms are characterized by an elevated reshuffling level of workers across plants but no increased job reallocation. This finding is also consistent with our result that in booms plants are more certain about separations.

This observation ties into the literature that studies cyclical movements in job reallocation (e.g., Davis and Haltiwanger (1992)), and in job-to-job transitions (e.g., Barlevy (2002) and Moscarini and Postel-Vinay (2013)). Our analysis shows that the entire rise in job-to-job transitions during booms leads ultimately to churn and not necessarily to job reallocation. That is, the degree to which more productive plants grow at the expense of less efficient plants does not appear to increase during booms. Rather booms appear to be times where workers and plants find individually better matches, implemented through more yet more predictable separations leading to higher worker churn. Consistent with the absence of systematic reallocation of workers towards particular plants, Lindenlaub and Postel-Vinay (2017) develop a job-ladder model where workers have idiosyncratic rankings of plants.

6 Conclusion

This paper, using a newly assembled plant-level data set from Germany, studies the procyclicality of churn through a micro-to-macro approach. It starts by documenting a cross-sectional link between plant-level worker churn and employment growth. We show that churn occurs across the entire employment growth distribution and is most pronounced at plants which rapidly change their employment level, that is, the plant-level churn-employment growth nexus is V-shaped. We also show that churn, and its relationship with plant-level employment growth, arises from churning workers with similar worker characteristics occupying jobs with similar job characteristics. What is more, churn is largest at plants with a putatively high uncertainty about their workers' match quality. Hence, churn is unlikely to reflect reorganization at the plant level, but rather uncertainty about match qualities that imply uncertainty about quarter-to-quarter separations. In fact, stochastic separation rate shocks that, through a time-to-hire friction, lead to planning mistakes by establishments do a good job explaining the cross-sectional churn-employment growth nexus. We thus show that studying churn can lead to new insights about plant-level employment dynamics.

Larger average separation rate shocks during booms, yet more dispersed separation rate shocks during recessions, do a good job in explaining procyclical churn and largely acyclical job turnover. Also, procyclical churn represents mostly workers reallocating to other establishments, not into non-employment, and these separations (and hiring) to (from) other establishments rise by a similar amount along the entire plant-level employment growth distribution. Early in a boom, additional separations are eventually replaced by some plants hiring workers from non-employment leading to high churn and high job creation but not high job destruction during this phase of the business cycle. Later in a boom, the high separation rates stemming from job-to-job transitions translate into neither job creation nor job destruction, and thus merely into a faster reshuffling of workers across plants.

References

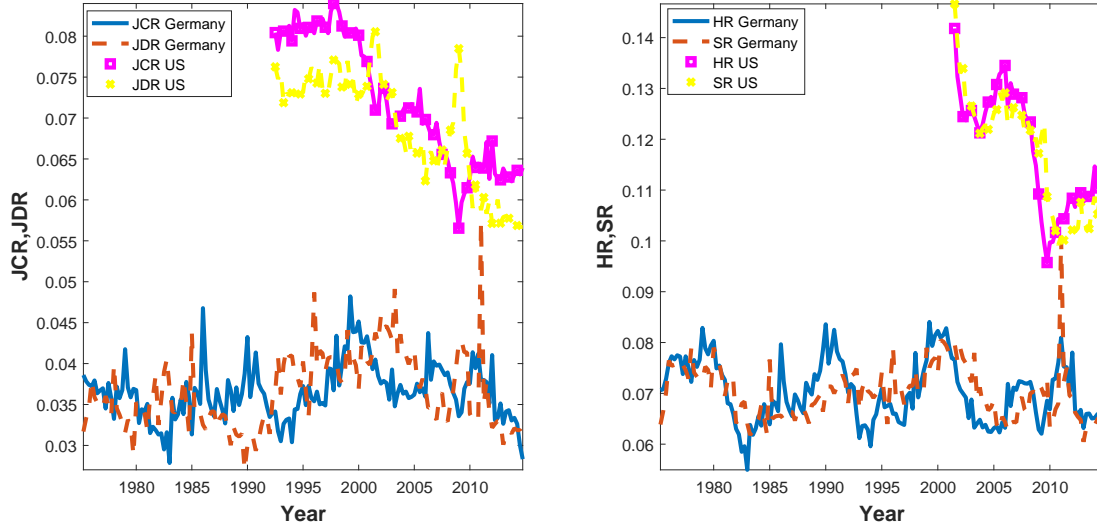
- Abowd, J. M. and Vilhuber, L. (2011). National estimates of gross employment and job flows from the quarterly workforce indicators with demographic and industry detail. *Journal of Econometrics*, 161(1):82–99.
- Bachmann, R. and Bayer, C. (2014). Investment dispersion and the business cycle. *American Economic Review*, 104(4):1392–1416.
- Barlevy, G. (2002). The sullyng effect of recessions. *The Review of Economic Studies*, 69(1):65–96.
- Bellmann, L., Gerner, H.-D., and Upward, R. (2018). Job and worker turnover in German establishments. *The Manchester School*, 86(4):417–445.
- Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28(2):153–176.
- Borovicková, K. (2016). Job flows, worker flows and labor market policies. mimeo.
- 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. and Haltiwanger, J. C. (1992). Gross job creation, gross job destruction, and employment reallocation. *Quarterly Journal of Economics*, 107(3):819–863.
- 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. (2018). Identifying labor market sorting with firm dynamics. mimeo, University of Mannheim.

- Haltiwanger, J. C., Hyatt, H., Kahn, L. B., and McEntarfe, E. (2018). Cyclical job ladders by firm size and firm wage. *American Economic Journal: Macroeconomics*, 10(2):52–85.
- Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60(5):1127–1150.
- 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.
- Kaas, L. and Kircher, P. (2015). Efficient firm dynamics in a frictional labor market. *American Economic Review*, 105(10):3030–3060.
- 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.
- Lindenlaub, I. and Postel-Vinay, F. (2017). Multidimensional sorting under random search. Mimeo, Yale.
- Merkl, C. and Stüber, H. (2016). Wage cyclicality and labor market dynamics at the establishment level: theory and evidence. *FAU Discussion Papers in Economics*, 12/2016.
- 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.
- Pries, M. and Rogerson, R. (2005). Hiring policies, labor market institutions, and labor market flows. *Journal of Political Economy*, 113(4):811–893.
- Pries, M. and Rogerson, R. (2019). Declining worker turnover: the role of short duration employment spells. Working Paper 26019, NBER.
- Schaal, E. (2017). Uncertainty and unemployment. *Econometrica*, 85(6):1675–1721.
- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49.
- Stüber, H. and Seth, S. (2017). The administrative wages and labor market flow panel. *FAU Discussion Papers in Economics*, 1/2017. updated Dec. 2018.
- Tanaka, S., Warren, L., and Wiczer, D. (2019). Earnings growth, job flows and churn. mimeo.

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.²⁷ 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

²⁷The 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}		
		<i>Mean</i>	<i>Std</i>	<i>AC(1)</i>	$j = -2$	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.

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. 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.²⁸ 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)).

²⁸There 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.

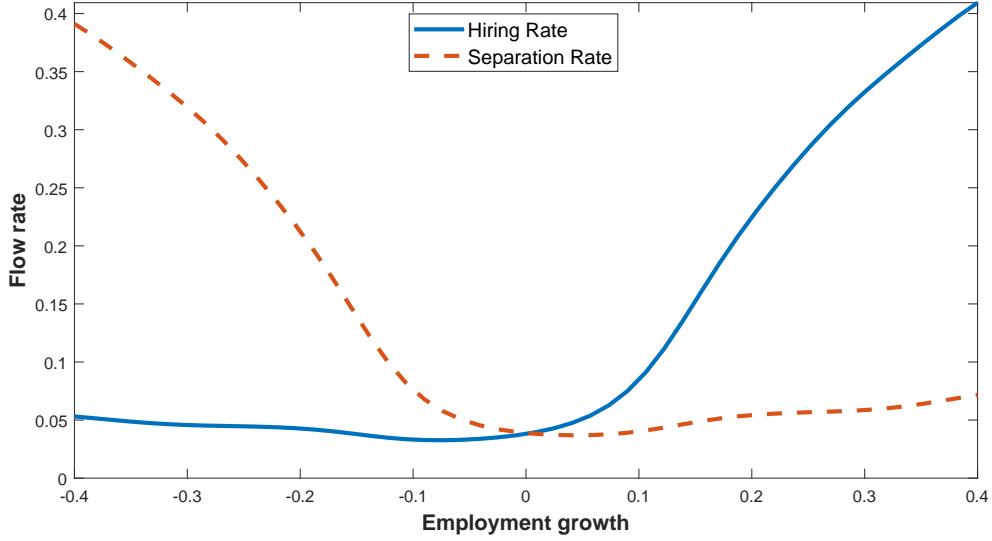
Table A2 computes the correlations between job and worker flows, respectively, for the two countries. In both countries, the job creation and destruction rates are negatively correlated. The job creation rates and hiring rates, and the job destruction rates and the separation rates are positively correlated.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

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:

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.

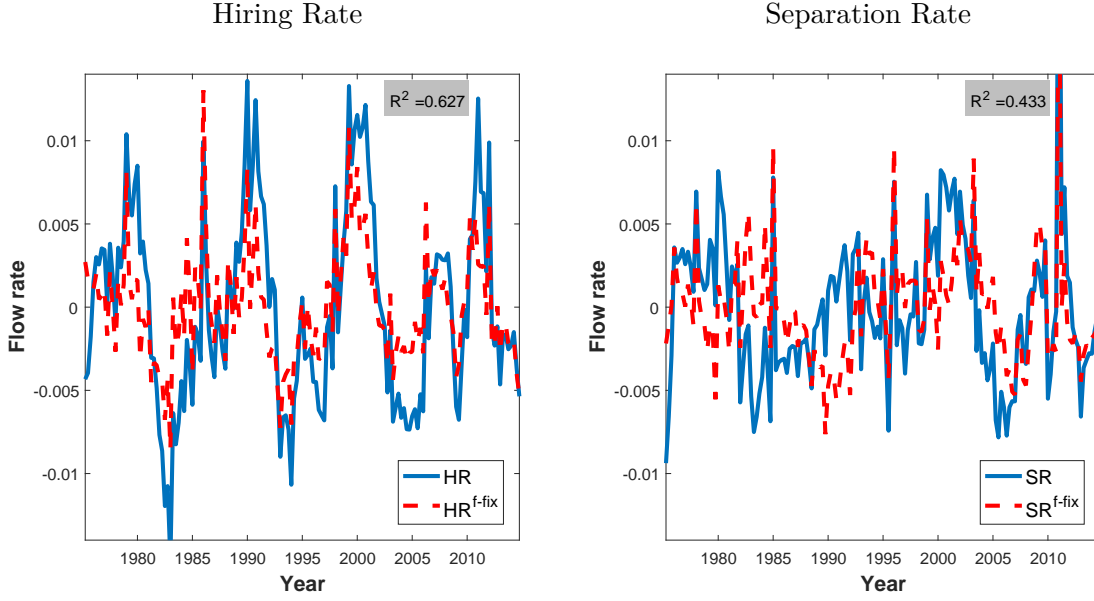
$$\begin{aligned}
 HR_t^{f-fix} &= \sum_{j=1}^J \overline{hr(j)} es_t(j) \\
 SR_t^{f-fix} &= \sum_{j=1}^J \overline{sr(j)} es_t(j),
 \end{aligned}
 \tag{A.1}$$

where $J = 21$ are employment growth categories/bins,²⁹ $es_t(j) = \frac{N_t(j)}{N_t}$ is the share of overall employment in an employment growth rate bin, and a bar denotes time-averaged values. 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%.

²⁹These 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 A3: Fixed Worker Flow Rates over the Cycle



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.

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:

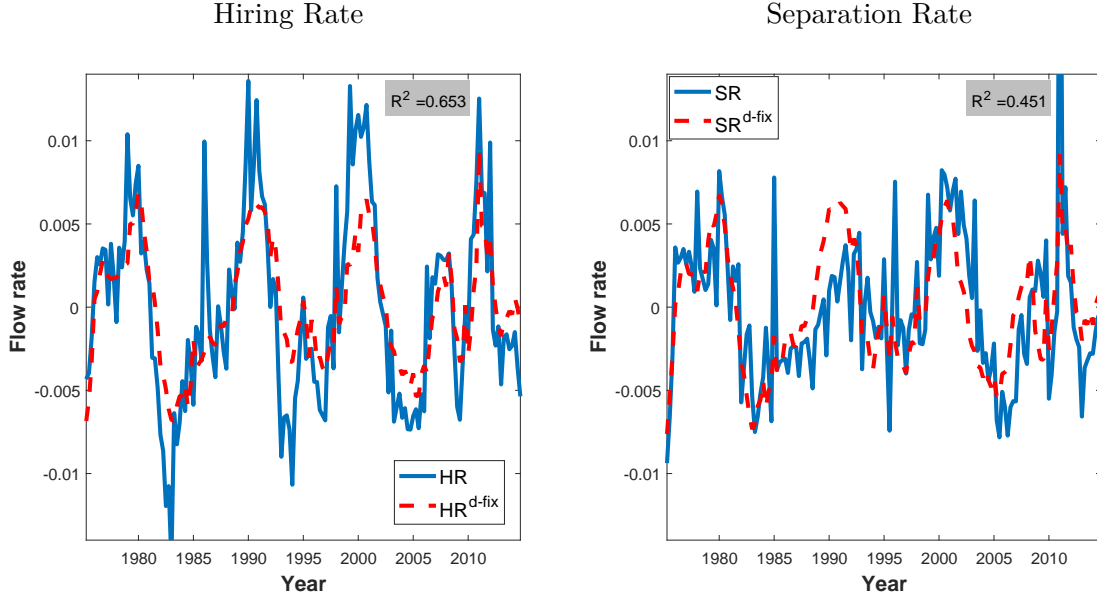
$$HR_t^{d-fix} = \sum_{j=1}^J hr_t(j) \overline{es(j)} \quad (A.2)$$

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

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

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.

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

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

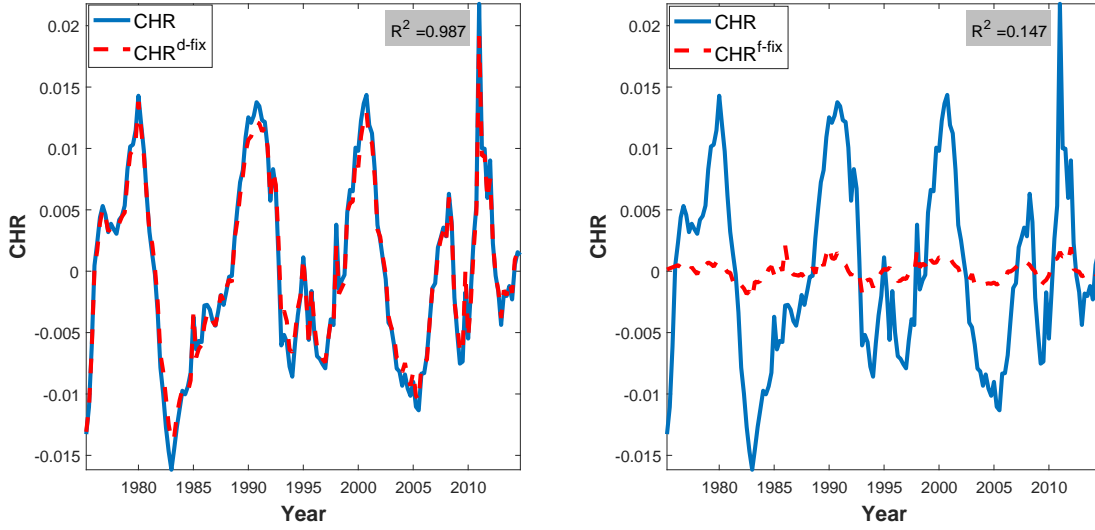
In order to understand the importance of the two channels of cyclical churn, consider the following statistical models:

$$CHR_t^{d-fix} = \sum_{j=1}^J chr_t(j) \overline{es(j)} \quad (\text{A.4})$$

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

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

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.

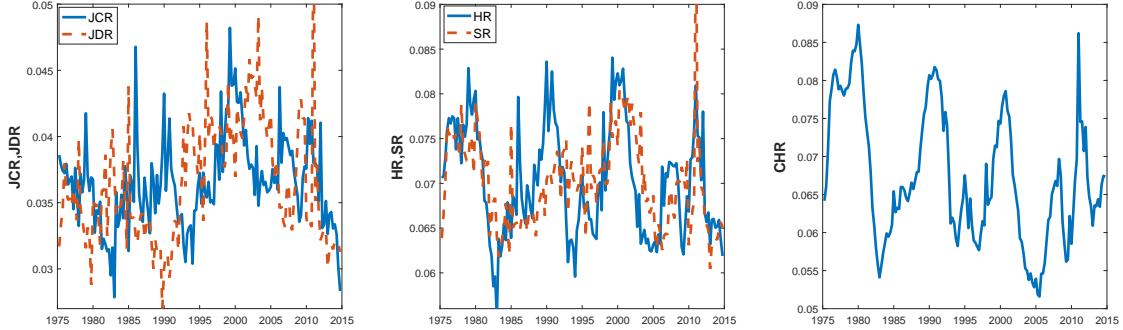
characterized by a shift away from marginally adjusting plants towards rapidly adjusting plants.

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.

B Additional Tables and Figures

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



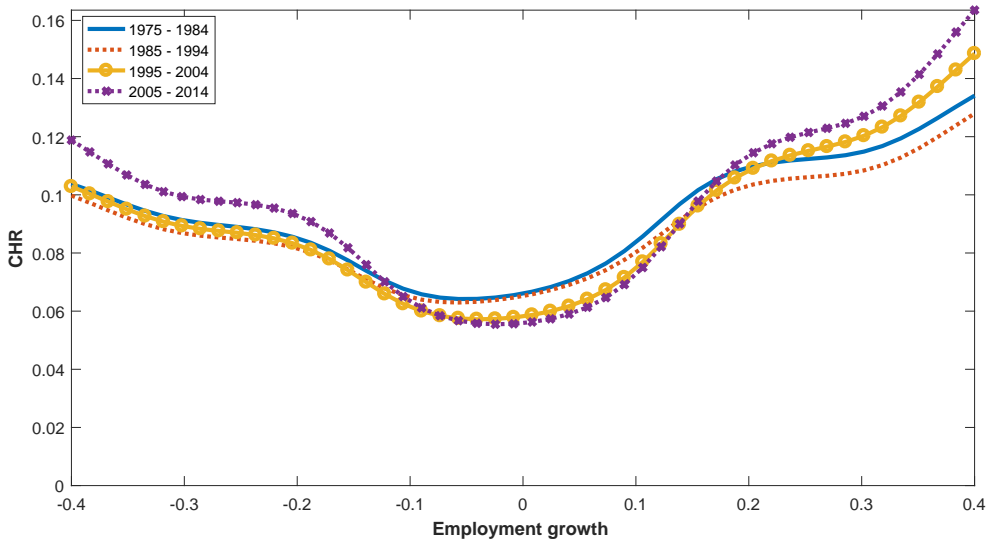
Note: The left panel displays aggregate job flows. JCR : job creation rate, JDR : job destruction rate. The center panel displays aggregate worker flows. HR : hiring rate, SR : separation rate. The right panel displays the aggregate churn rate, CHR . All rates are seasonally adjusted. West German plants 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.

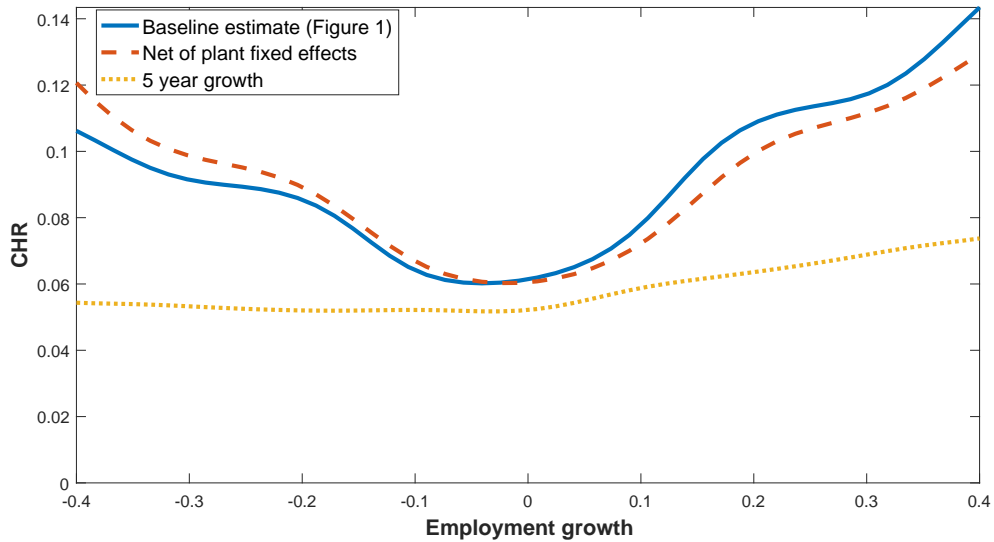
The model we present in Section 4 describes plants that, resulting from transitory separation rate shocks, fluctuate around their long-run optimal employment level. Section 3 studies the relationship between plant-level churn and employment growth without differ-

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



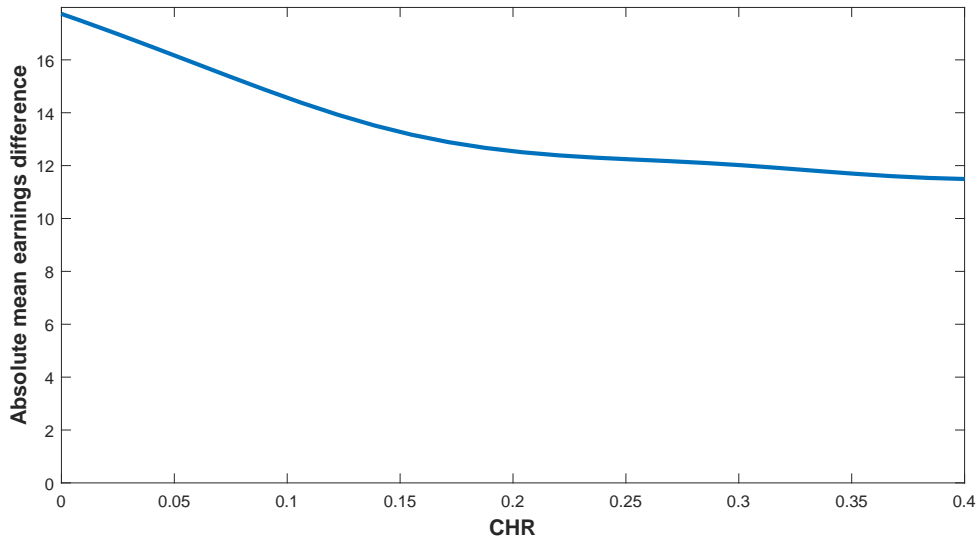
Note: This figure displays the average churn rate of a plant as a function of its employment growth rate estimated by an N_{it} -weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The blue line uses quarterly employment growth rates. The red dashed line controls for plant fixed effects 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.

entiating long-run employment growth trends from transitory deviations from these trends. Figure A8 shows that the relationship between plant-level churn and employment growth 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

Figure A9: Churn Rates and Wage Dynamics



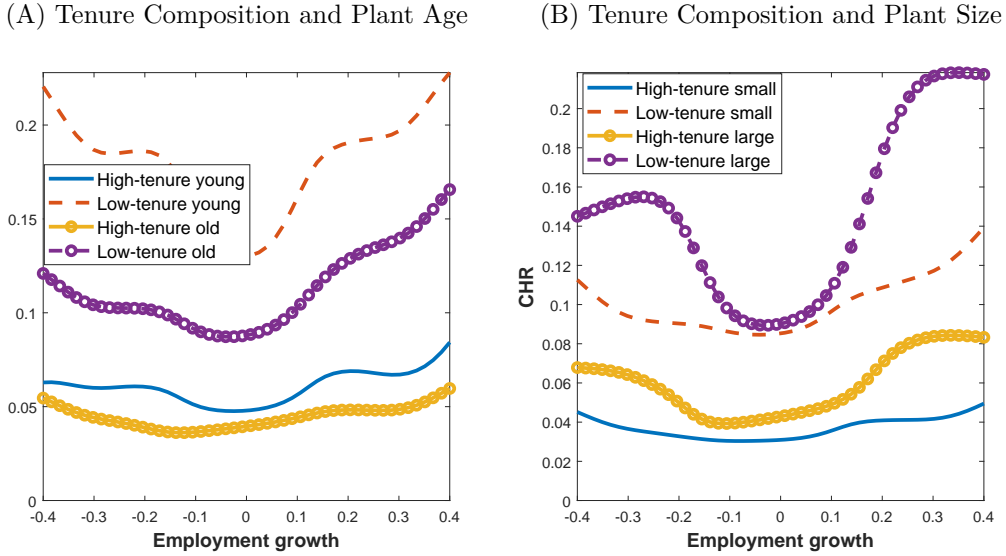
Note: This figure displays the absolute difference in average daily earnings of in- and out-flowing workers of a plant as a function of its churn rate. The estimate results from a semi-parametric regression which estimates a non-parametric relationship between churn and wage-differences (using a Gaussian 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.

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 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.

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

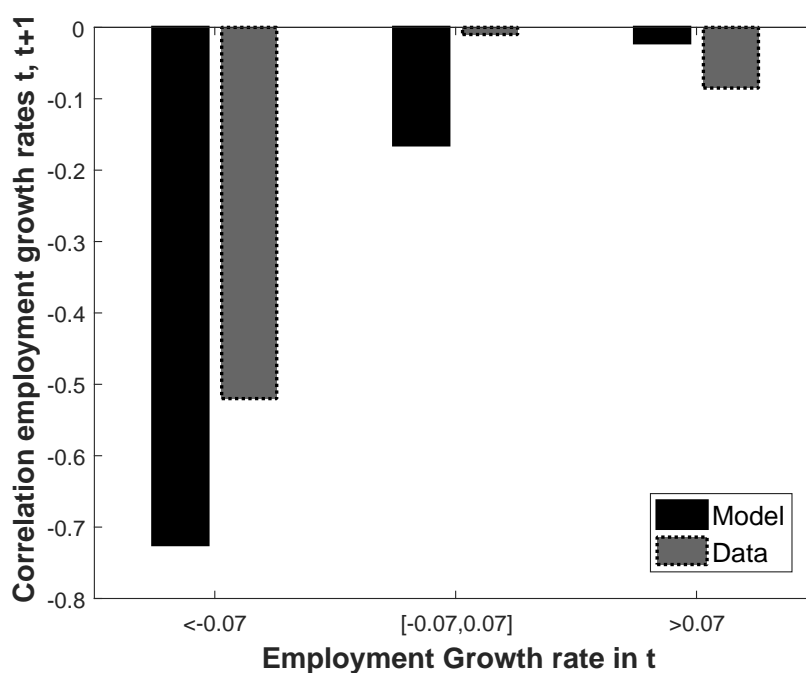
On the other extreme, plants that grow by more than 7% are plants that experience a larger than expected separation rate shock in the previous period (period $t - 1$) and, thus, ended the previous period with fewer than the optimal amount of workers. Therefore, in period t they re-hire and are, on average, 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 auto-correlation of employment growth comes mainly from plants that shrink by a substantial 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

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.

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.

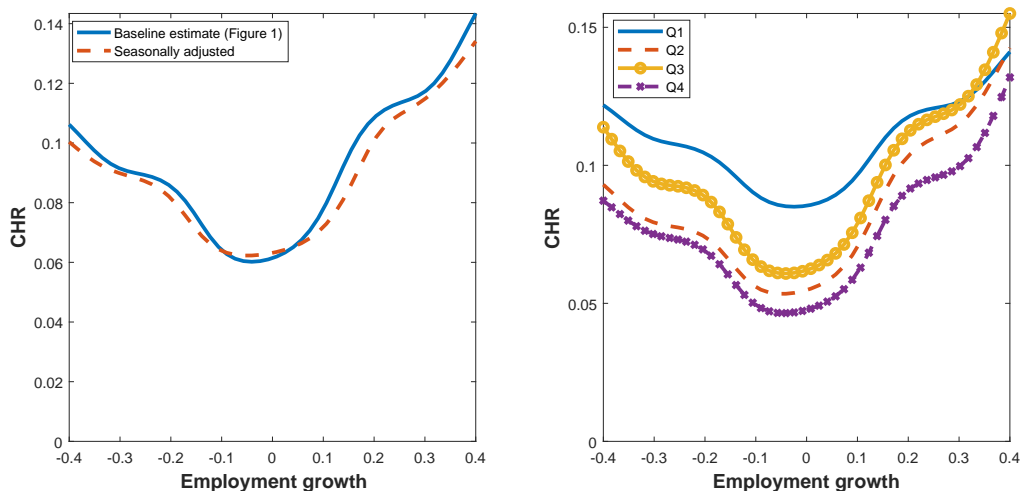
Table A3: Cyclical Dynamics of the Churn Rate and Employment Growth

Employment growth rate	<i>Mean</i>	<i>Std</i>	<i>AC(1)</i>	<i>CorrU</i>
-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



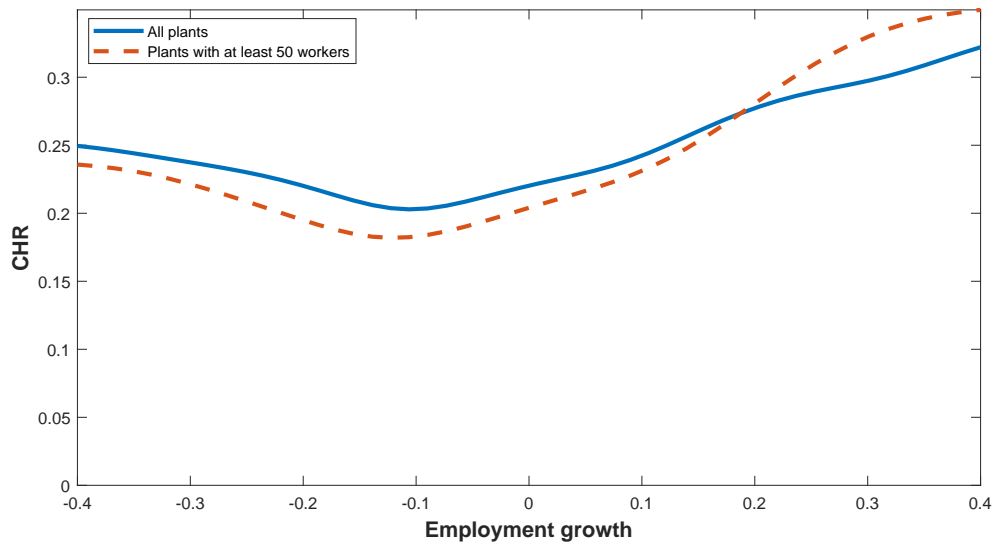
Note: The left panel displays the average churn rate of a plant as a function of its employment growth rate. The blue line is estimated by an N_{it} -weighted kernel smoother (Gaussian kernel with a 0.05 bandwidth). The red dashed line is estimated the same way after regressing plant-level employment 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:

$$Y_{it} = z_{it} E_{it}^\alpha, \quad (\text{A.5})$$

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:

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

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

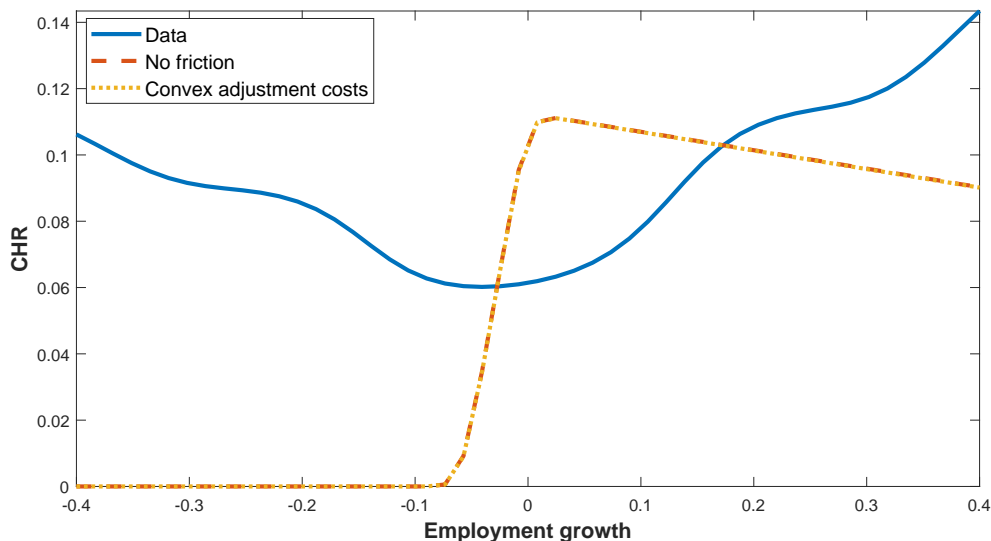
$$E_{it} = (1 - \bar{s})E_{it-1} + \Delta_{E_{it}}^a. \quad (\text{A.7})$$

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 their value:

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

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.

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.

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

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 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

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.

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.³⁰

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 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.

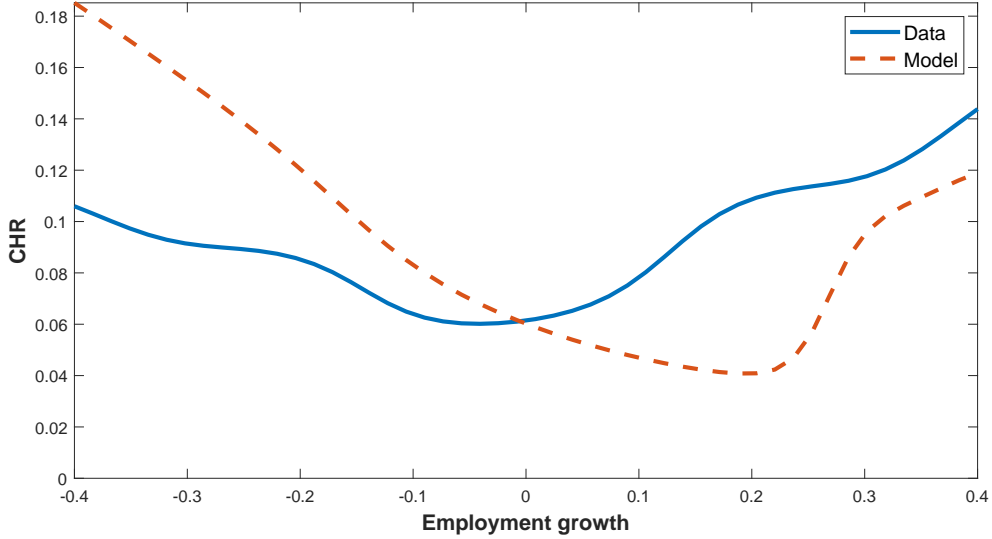
Using the notation from the main text and the previous section, the dynamic program of a plant in such an environment is given by:

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

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 .

³⁰The parameters for the frictionless recalibration are not displayed for the sake of brevity but are available upon request.

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.

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 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

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.

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:

$$\max_{\Delta_E^a} \left\{ \mathbb{E}_{t-1} \{ Y_{it} - w E_{it} \} \right\}, \quad (\text{A.8})$$

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

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

Differently from the baseline model, the mean separation rate is now allowed to be different

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.

for low- and high-tenured workers, and we capture this difference by the parameter κ :

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

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

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:

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

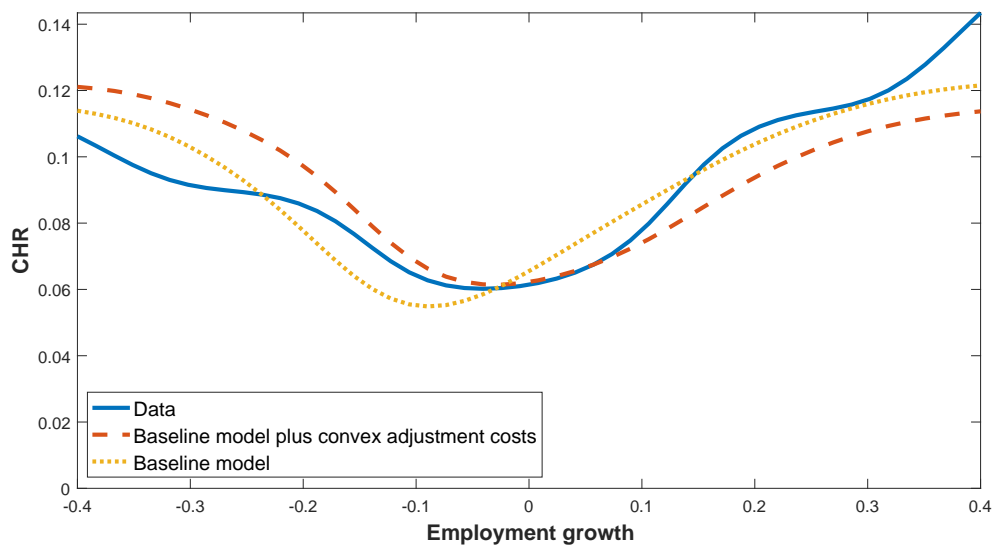
Differently from the baseline model, plants now pay also convex adjustment costs when actively adjusting their workforce, on top of the time-to-hire friction: $c_{it} = \psi \left(\Delta_{E_{it}}^a \right)^2$. The resulting dynamic programming problem becomes:

$$\begin{aligned} V(E_{-1}, \tilde{s}_{-1}) &= \max_{\Delta_E^a} \left\{ \mathbb{E}_{-1} z \{ E^\alpha - wE \} - \psi \left(\Delta_E^a \right)^2 + \frac{1}{1+r} \mathbb{E}_{-1} V(E, \tilde{s}) \right\} \quad (\text{A.13}) \\ E &= (1 - s)(E_{-1} + \Delta_{E_{it}}^a) \\ s &= \min\{\exp(\tilde{s}), 1\} \\ \tilde{s} &= (1 - \rho_s)\mu_s + \rho_s\tilde{s}_{-1} + \epsilon, \quad \epsilon \sim N(0, \sigma_s^2), \end{aligned}$$

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.