

Worker Flows Over the Business Cycle: the Role of Firm Quality[†]

Lisa B. Kahn
Yale School of Management and IZA

Erika McEntarfer
U.S. Census Bureau

April 1, 2013

Abstract

In this paper, we study net and gross worker flows over the business cycle as a function of firm quality. Using linked employer-employee data from the LEHD program at the U.S. Census, we measure employer quality as average pay, though our results are robust to a number of different measures. We first show that net job creation at high-quality firms is *more* responsive to the business cycle than that of low-quality firms; in recessions low-quality firms shrink less quickly, while in booms high-quality firms grow more quickly. We then show that gross hire and separation rates at high-quality firms are *less* responsive to the business cycle. While these gross flow rates decline in recessions, they decline by less in high-quality firms. Therefore the growth rate effect can be accounted for by a larger decline in job separations in low-, compared to high-, quality firms in recessions. We therefore find that jobs at low-quality firms become stickier in recessions, relative to jobs at high-quality firms. We conclude with a discussion of our results in the context of existing macroeconomic theories of the labor market. We provide further results based on implications from these models to determine the channel through which our results . For example, we find that our results are likely not being driven by differential downward wage rigidities over the business cycle, or differential business cycle sensitivity, but are more likely labor supply driven.

JEL codes: J63, J64

*The research program of the Center for Economic Studies (CES) produces a wide range of economic analyses to improve the statistical programs of the U.S. Census Bureau. Many of these analyses take the form of CES research papers. These papers have not undergone the review accorded Census Bureau publications and no endorsement should be inferred. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau or its staff. All results have been reviewed to ensure that no confidential information is disclosed. Republication in whole or part must be cleared with the authors. Correspondence: Lisa Kahn, Yale School of Management, 135 Prospect St, PO Box 208200, New Haven, CT 06510. Email: lisa.kahn@yale.edu. Erika McEntarfer, Center for Economic Studies, U.S. Census Bureau, 4600 Silver Hill Rd., Suitland, MD 20233. Email: erika.mcentarfer@census.gov.

[†]We are grateful to John Haltiwanger, Alina Lerman, Giuseppe Moscarini, Peter Schott, Jim Spletzer, Till von Wachter and seminar participants at the AEA 2013 meetings and McGill for helpful comments.

1 Introduction

Worker sorting across firms has long been thought to play a central role in labor market efficiency. Despite frictions that can inhibit this sorting process, such as search costs or imperfect learning, workers are thought to gradually move towards jobs of better overall- or match-specific quality.¹ At the same time, recessions may impede worker sorting. Several papers have noted a marked decline in worker churning and job-to-job mobility in recent recessions, with a particularly sharp downturn in job change during the Great Recession.² This suggests that workers’ ability to move on from poor job matches or bad jobs is curtailed in times of high unemployment. A natural question, then, is in what types of jobs are workers – at least temporarily – saddled? If the business cycle has differential employment impacts on jobs or firms of varying quality, the consequences of reduced mobility could be very different. In this paper, we ask how firm quality interacts with the business cycle. That is, we investigate whether the employment effects of the business cycle are heterogeneous across firms of differing quality.

If resources are reallocated to higher quality firms in recessions (the classic Schumpeter 1939 cleansing effect) then we might see a commensurate flow of workers to good firms. However, the cyclical upgrading literature (Okun 1973, Bils and McLaughlin 2001) suggests that high-quality jobs may be more sensitive to the business cycle, with opportunities to move into these jobs relatively more prevalent in expansions. Further, Barlevy (2002) shows that the decline in job-to-job transitions seen in recessions has a quantitatively important effect on overall match quality, terming this the “sullyng effect” of recessions. However, if job destruction in recessions occurs relatively less at lower quality firms, then we would have a further sullyng effect; jobs in recessions would be both lower quality matches and in lower absolute quality of firms. In this paper we directly analyze the differential impact of economic conditions on net and gross worker flows as a function of firm quality.

¹This idea goes at least as far back as the canonical work of Jovanovic (1979) and for empirical work on job mobility see Farber’s 1999 survey.

²See in particular Lazear and Spletzer (2012), Hyatt and McEntarfer (2012).

To identify worker flows over the business cycle we use data from the Longitudinal Employer Household Dynamics (LEHD) program; a U.S. employer-employee matched database drawn from the state unemployment insurance systems. This dataset allows us to match detailed worker job histories with a rich set of firm-level characteristics. To measure quality, we divide firms into quintiles based on average pay. Our results are robust to alternative measures of firm quality, including firm size and worker churn rate. Furthermore the LEHD data allow us to track gross, as well as net, worker flows across firms. We therefore analyze quarterly employment growth rates, as well as gross hire and separation rates, as a function of the unemployment rate and firm quality from 1998 to 2008. This time period allows us to capture the 2001 recession as well as some of the decline into the 2007-09 recession.

We find that net employment growth at high-quality firms is more responsive to the business cycle than that of low-quality firms. It is driven both by greater net job destruction among high-quality firms at times of high unemployment, as well as a greater net job creation in times of low unemployment. We can define unemployment using either national or state-level economic conditions, where in the latter case, we can control for date fixed effects. To explain these findings, we next look at gross worker flows. In contrast to their more-responsive growth rates, we find that at high-quality firms, gross worker flows are less responsive to the business cycle; separation and hire rates both decline by less in high-, compared to low-, quality firms when the unemployment rate increases. Therefore, the relatively higher growth rates seen in low-, compared to high-, quality firms in recessions is accounted for by a larger decline in separation rates. While the recession creates an adherent effect for most jobs, it does so relatively more at low-quality firms.

Our results are broadly consistent with a recent body of work looking at growth rates over the business cycle as a function of firm size. Moscarini and Postel-Vinay (2012a, hereafter MPV), show in a number of countries including the U.S. that differential growth rates of small-, compared to large-, firms are positively related to the unemployment rate. Fort, Haltiwanger, Jarmin and Miranda (2012) analyze firm growth over the business cycle as a

function of firm age and size, using U.S. data. They find that small, young, firms typically fare relatively better in cyclical contractions, although this relationship reversed in the 2007-09 recession. We contribute to this literature by showing the growth rate effect holds along a number of firm-quality characteristics, including and within firm size. Also, our added dimension of gross worker flows allows us to paint a richer picture of labor market dynamics over the business cycle, that is not possible in the other datasets used to study growth rates.

We use our body of evidence to disentangle macroeconomic models with predictions of worker mobility over the business cycle. We first use our wage data to test whether differential downward wage rigidities can explain our results. If high-quality firms have more trouble lowering their workers' pay, then they would be more likely to respond to a negative demand shock by laying off workers. However, following a similar methodology to (Dickens, et al., 2007) we find the opposite: high-quality firms are less likely to suffer from downward wage rigidities in recessions.

We next use Compustat data to test whether low-quality firms are in subsectors that are more sensitive to the business cycle. Specifically, we examine the relationship between changes in revenue and the business cycle. However we find that across firm quality categories, firms are similarly responsive to the business cycle in terms of revenue.

These tests point us toward labor supply driven explanations for our findings. In particular, the poaching story outlined in Moscarini and Postel-Vinay (2012b) is a plausible explanation: High-quality firms have an easier time attracting workers in booms, so they grow relative to low-quality firms and inflate in size. During the bust they must then shed some of these workers. At the same time, low-quality firms have an easier time retaining workers in a bust, since high-quality firms are less likely to poach workers then.

Our results are very much in the spirit of Barlevy (2002) where workers are stuck in low-quality matches in recessions. Beyond this, our results indicate that recessions cause workers to stay in worse overall quality firms. This has important implications for the long-lasting consequences of recessions on workers. A growing body of evidence suggests that recessions

have vastly differing impacts on workers over the long run, depending on what stage of their career the recession hits them in. First, labor market conditions at the beginning of a worker’s career have long-lasting scarring effects (Kahn 2010, Oreopoulos, von Wachter and Heisz 2012). Second, the consequences of job displacement have been shown to be much larger when displacement occurs in a recession (Davis and von Wachter 2011). It therefore seems that being forced to match to a firm during an economic downturn can be incredibly damaging to a worker’s career. Our finding that, relatively speaking, low-quality firms grow faster in recessions (or shrink less quickly) can potentially explain these findings. It suggests that matches occurring in downturns will be relatively stickier at low-quality firms than at high-quality firms.

The remainder of the paper proceeds as follows. Section 2 describes the data and presents some aggregate trends. Section 3 describes our methodology for studying firm-level net and gross flows. Section 3 presents our results. Section 4 discusses various models of the business cycle and the degree to which our results are consistent. Section 5 concludes.

2 Data

We analyze worker flows over the business cycle using data from the U.S. Census Bureau’s Longitudinal Employer-Household Dynamics (LEHD) program. The LEHD program maintains a variety of survey and administrative data from several state and federal agencies. For this paper, we chiefly use state unemployment insurance (UI) wage records and the Quarterly Census of Employment and Wages (QCEW) data. Both UI and QCEW data are available for states in partnership with the LEHD program, currently all 50 states and the District of Columbia. A thorough discussion of the LEHD data is provided in Abowd et. al. (2006); a brief description follows.

State-level unemployment insurance (UI) data contain quarterly earnings for employees covered by state unemployment insurance systems, over 96% of private sector employment. A

firm, as defined in this analysis, is a collection of workers who share a common unemployment insurance system identifier. Individual wage records can be linked across quarters to create individual work histories, worker flows, and earnings dynamics. The firm identifier on the UI records is used to link to information on the firm available in the QCEW data (we principally use employment size and industry). Worker demographics, namely sex and date of birth, are available from links to the Census administrative and survey data. For this paper we largely restrict attention to the 30 states that have UI and QCEW data for every quarter of our sample period 1998:Q1-2008:Q4.³

These data are advantageous in that they allow us to observe both gross and net worker flows for a substantial fraction of firms in the U.S. labor market. Furthermore, we can create a rich set of firm characteristics to measure employer quality. Finally, the time period over which we can exploit a balanced panel consisting of a large number of states allows us to capture one complete business cycle containing the 2001 recession, as well as some of the employment decline in the 2007-09 recession.

We focus on average wage as our measure of firm quality. Since one goal of this paper is to better understand the experiences of workers in recessions, we would ultimately like our quality measure to correlate with firms workers would like to be in. Higher paying firms nicely fit this description. Furthermore, Serafinelli (2012), for example, presents evidence using detailed administrative data in Italy that high paying firms are more productive and therefore probably less likely to close. Our exercise in this paper is to analyze how firms of different qualities are impacted by the business cycle. We therefore construct time-invariant firm quality measures by taking average wage within a firm (to be more precise, a firm-state) over our entire sample period. This avoids the well-known reclassification bias problem (discussed, for example, in MPV), though our results are robust to other measures.⁴

Our subsequent regression analysis is robust to measuring firm quality with average size

³For some analyses we reduce our sample to 25 states that have complete establishment-level worker flows data for the entire length of the panel.

⁴In particular, we have experimented with using a two-quarter moving average as in Fort et al. (2012) and we will also check robustness to using average quality in an initial period of measurement (as in MPV).

and excess churn. Larger firms have been shown to have higher pay, better working conditions, a greater degree of benefits provision, increased productivity, and increased probability of firm survival (Brown and Medoff 1989, Hurst and Pugsley 2011). Equation 1 defines the excess churn rate in a given period, t , at a firm, f , where A is hires, S is separations, and B and E are beginning and end of quarter employment, respectively. Thus we define churn as hires and separations in excess of the net employment change in the period ($E - B$), divided by average employment in the period. A firm with a high-churn rate has a high number of worker flows in excess of job flows. We take this definition, which is now standard in the literature, from Burgess, Lane, & Stevens (2000). Cambell et al. (2005) show that high churn is associated with lower productivity and lower survival rates for a select set of industries. Since both size and churn are correlated with pay, we will control for these characteristics in our analysis.

$$churn_{tf} = \frac{A_{tf} + S_{tf} - |E_{tf} - B_{tf}|}{.5 * (E_{tf} + B_{tf})} \quad (1)$$

Figure 1 shows employment-weighted kernel densities of each measure of firm quality. The top left panel shows the distributions of firm-level average churn rates; the top right panel shows the distribution of average monthly wages (for employees who work an entire quarter, in 2008 dollars); the bottom left panel shows the distribution of average firm size, which is the size of the state tax identity on the 12th day of the first month of the quarter, averaged over the life of the tax identification number. All of these distributions have long right tails; to avoid potential data disclosure issues in these graphs we cap churning at 2, average wages at \$12,000, and firm size at 15,000. As can be seen, we have substantial variation across firms over this time period in all measures.

In our subsequent analysis we divide firms into discrete categories based on these measures. For pay and churn, we use employment-weighted quintiles as dividing points. For most of our analysis, we use within industry cut points (measured at the two-digit NAICS level). For size, we use 5 categories: less than 20 employees, 20-50, 50-250, 250-500, and

greater than 500, following Fort et al. (2012).⁵ We use pay quintile as our proxy for firm quality. Therefore, high-quality firms will be those that pay more on average, *relative* to other firms in the same industry. In some specifications we will control for the size and churn distributions.

Our key dependent variables in this paper are net growth rates as well as gross flow rates. To calculate these rates, we aggregate our firm-level data to the year-quarter-industry-wage quintile category, by summing employment and worker flows in each cell.⁶ This level of aggregation allows us to control for industry, while still enabling us to capture employment dynamics driven from firm births and deaths.⁷

Table 1 presents employment-weighted summary statistics by firm category for our rates of interest, which we define next. The quarterly growth rate for a firm quality type, q , is defined in equation 2 as net employment change among all firms of type q (firms indexed 1 to F_q) divided by average employment over the quarter, t , among these firms. As can be seen in table 1, average growth rates range from 0.008 to 0.01, with a few differences across firm categories. Higher paying firms within an industry have slightly higher growth rates. Note from the churn and size distributions that these firms have lower churn rates and are larger.

$$\text{growth rate}_{tq} = \frac{\sum_{f=1}^{F_q} (E_{tf} - B_{tf})}{.5 * \sum_{f=1}^{F_q} (E_{tf} + B_{tf})} \quad (2)$$

⁵However, it is worth noting that our measure of size is the employment size of the state tax entity. Both Fort et al. and MPV use firm size data from the BDS, which contains information on both establishment-level employment and national employment. Our measure of firm size is correlated with the national size of the firm (0.75) but is not an exact match, more closely approximating the size of the firm in the state. BDS measures of firm size are newly available in the LEHD data but were not yet available at the time of this analysis.

⁶Our industry measure is the two-digit NAICS code, though in principle we could use much more disaggregated industry definitions.

⁷While in principle, we could conduct our analysis at the individual firm level, that would produce growth, hire and separations rates that are quite a bit noisier. These rates are misleadingly large in the period in which a firm starts or closes and outliers can be generated by seasonal employers or non-reporting events (or in principle mergers and acquisitions, though the LEHD use an algorithm to exclude these events). At the individual firm level, these outliers create problems for our estimation, so we prefer the somewhat aggregated analysis presented here.

Hire and separation rates are defined in equations 3 and 4, respectively, as the total number of hires or separations in quarter, t , at firms of quality, q , divided by average employment over the quarter. Our results are robust to an alternative denominator, total employment over the quarter, which is sometimes used in the literature. However, our definitions are convenient because the common denominator across three rates implies that the hire rate minus the separation rate must add up to the growth rate (see Lazear and Spletzer 2012).

$$\text{hire rate}_{tq} = \frac{\sum_{f=1}^{F_q} A_{tf}}{\sum_{f=1}^{F_q} (E_{tf} + B_{tf})} \quad (3)$$

$$\text{separation rate}_{tq} = \frac{\sum_{f=1}^{F_q} S_{tf}}{\sum_{f=1}^{F_q} (E_{tf} + B_{tf})} \quad (4)$$

These gross worker flows are not available in most datasets, even those containing measures of net employment growth, and herein lies much of our contribution. Table 1 indicates that hire and separation rates vary widely across firm category, from 0.14 to 0.33, and are highly correlated within firm category. This implies that most hiring is churn related, serving to replace workers who have separated. This is evident in the table which shows that lower paying firms are more likely to be high churn firms and these firms also have the highest hire and separation rates.

To gain a general sense of hiring over the business cycle, we first look at differential growth, hire and separation rates across our lowest and highest quality firm quintiles. We simply subtract the rate in the highest quality bucket from that in the lowest. MPV do a similar exercise comparing growth rates at large and small firms. These differential growth rates are plotted in figure 2 along with the national unemployment rate (dashed line). Both lines have been seasonally adjusted by residualizing on quarter dummies (therefore the levels are not that meaningful). We have also detrended each series using a Hodrick-Prescott filter,

following MPV.

The top left graph shows the differential growth rate across low- and high-wage firms. Though noisy, this differential growth rate very closely tracks the national unemployment rate. That is, when unemployment is high, low-wage firms grow relatively more quickly (or shrink less quickly) than high-wage firms, while when unemployment is low, low-wage firms grow relatively less quickly. Though this trend does seem to break down somewhat during the jobless recovery following the 2001 recession, the correlation between the two rates is positive and significant.⁸

The top right and bottom left panel show differential hire and separation rates, respectively. These gross flow rates exhibit very different patterns than the differential net growth rate: Differential separation and hire rates look roughly procyclical; when the unemployment rate is low, low-quality firms hire and separate at greater rates than high-quality firms, while the opposite is true in times of high unemployment. Further the correlations are negative and highly significant.

3 Methodology

In the next section we investigate the patterns exhibited in figure 1 in a regression framework where we can control for many potentially confounding factors. Specifically, we estimate regressions of the form specified in equation 5. We regress $rate_{tq}$, a growth, separation, or hire rate among firms of quality, q , in time period, t , on the national unemployment rate (nat_unemp_t), a vector of firm quality indicators (W_q) corresponding to wage quintiles and their interactions. We omit the lowest quality bucket. We additionally control for industry fixed effects at the two-digit NAICS level and X_t , a vector of controls for time period (including quarter dummies to control for seasonality and a time trend). We subsequently

⁸For ease of presentation, this graph excludes 2 outlier data points. This allows us to narrow in on the dynamics, which would have been obscured in a much larger scale on the y-axis. It is unclear whether those points might be errors in the data, but they do preserve the relationship between the differential growth rate and the unemployment rate, just on a much larger scale.

control for the average distribution of size and churn within a wage-quality bucket. As can be seen from table 1, firms across wage quintiles differ substantially in their churn and size distributions. We would like to know the extent to which these characteristics are driving our results. Finally, we will cluster our standard errors by firm quality-time period, since this is the level of variation underlying our key explanatory variables.⁹ All regressions are weighted by average employment over the quarter.

$$rate_{tq} = \alpha_0 + \alpha_1 nat_unemp_t + W_q \alpha_2 + [nat_unemp_t * W_q] \alpha_3 + X_t \alpha_4 + I^{industry} + \varepsilon_{tq} \quad (5)$$

We have chosen our firm quality categories as being relative to other firms in the same industry. We believe this is the correct measure for our purposes, as will become more evident below. However, we would also like to know whether these effects aggregate up to an important level in the national economy. Therefore we also test the robustness of our regression results to firm quality cut points that are overall, rather than within industry. We still weight by employment. Thus in this robustness check, a high quality firm is one who has higher pay relative to the distribution of workers in our sample. This alternative measure has advantages and disadvantages, which we discuss below.

Finally, we acknowledge that we only have a short panel of data, covering roughly one and a half business cycles. The shortness of our time series is evident in figure 2. This implies we only have limited business cycle variation as well as only a limited ability to control for other factors in our time series. We therefore also exploit cross-sectional variation in economic conditions. We disaggregate our data to the state-industry and also generate quintile cut points that are within state-industry. We then estimate regressions of the form specified in equation 6. These regressions additionally control for state and date fixed effects. We thus exploit within-state shocks, controlling for whatever is occurring in the national economy. This is useful because it also enables us to control for other time series events occurring at the national level. It also yields much more variation in business cycle conditions since

⁹The results presented here do not yet reflect clustering.

we have 30 states that could potentially bear the aggregate shocks differently. Here we will cluster our standard errors by state-date-wage quintile.

$$rate_{tsq} = \alpha_0 + \alpha_1 st_unemp_{ts} + W_{sq} \alpha_2 + [st_unemp_{st} * W_{sq}] \alpha_3 + I^{industry} + I^{state} + I^t + \varepsilon_{tsq} \quad (6)$$

4 Results

4.1 National Results

Table 2 presents our core set of regression results, summarize coefficients on the unemployment rate and its interactions with firm characteristics. Columns labeled I report the specification as described in equation 5, above, while columns labeled II additionally control for average churn and size distributions. Our wage quintile cut points here are within industry.

The first two columns summarize regressions where the dependent variable is the growth rate. The main effect of the unemployment rate, shown in the top row, is negative but is actually not statistically significant. This coefficient can be interpreted as the impact of the unemployment rate on the growth rate for the lowest quality firm bucket – the omitted category in each regression. Here we see that in low wage firms, a one percentage point increase in the unemployment rate corresponds to a modest 0.0026 point decline in the growth rate.

The interaction terms show the differential impact on the growth rate at higher quality firms. That they are all negative suggests that the unemployment rate has a larger, negative impact on the growth rate at higher quality firms. For example, a coefficient of -0.0084 among the highest quality firms, significant at the 1%-level, indicates that the unemployment rate at these firms has roughly four times the impact it had at low-quality firms. This effect is quite large considering the mean growth rate for this group was roughly 0.01. As would be expected, coefficients fall in magnitude for lower wage quintiles. The second column

shows that controlling for size and churn does not impact these coefficients. Therefore, these regressions show that higher quality firms fare worse in times of higher unemployment; their employment declines by more.

We next study gross separation and hire rates, to learn the extent to which movements in each are contributing to the relative declines in employment in high-quality firms. We start with the separation rate regressions, reported in the second set of columns in table 2. First, the main effect of the unemployment rate is negative and statistically significant at the 1%-level. Among the lowest quality firms, separation rates decline by 0.017 to 0.023, depending on the specification, when the unemployment rate increases by one percentage point. Though firms are more likely to make layoffs in a recession, our findings are consistent with a more-than-offsetting decline in voluntary quits (e.g., Shimer 2005, Hall 2005a). Interestingly, the positive, significant interaction effects across both specifications show that the response of separation rates to the business cycle is less negative in higher quality firms. These coefficients are all significant at the 1%-level and increase in magnitude for higher wage quintiles, as would be expected since these yield the sharpest contrast.

High-wage firms are on average larger and have lower excess churn rates (see the distributions summarized in table 1). We control for these distributions in the second column and find that the coefficients on the interaction terms fall somewhat. This suggests that larger firms and those with lower churn also experience smaller fluctuations in their separation rates with respect to the business cycle, and this is driving some of our effect. However the remaining effects are still sizeable, positive and statistically significant.

Finally, table 2 reports results from hire rate regressions in the last two columns. These are virtually identical to the separation rate results. The national unemployment rate significantly reduces hire rates, but by markedly less in high quality firms. Since the unemployment rate has a less depressive effect on both hire and separation rates at high quality firms, our growth rate effect must be accounted for by a larger differential impact on separation rates. Relatively speaking, low-quality firms grow during times of high unemployment because they

have a larger reduction in separations. This is true despite their larger reduction in hires.

Figure 3 exhibits these effects more clearly, plotting coefficients as a function of firm wage quintile, using the second specifications which control for size and churn distributions. We plot the main effect of the unemployment rate as the value for the lowest wage firm. We then add this coefficient to the interaction terms. Thus the graph represents the total impact of a one percentage point increase in the unemployment rate on growth, hire and separation rates at each wage quintile. We do this so that 0 is meaningful and highlight 0 in the graph.

As can be seen with the solid, blue line, the growth rate effect is negative and steadily declines (increases in magnitude) across higher wage quintiles. In contrast, the separation and hire rate impacts steadily increase across wage quintiles. The growth rate effect is the difference between the hire and separation rate effects. What is clear from this graph is that this differential becomes larger for higher wage quintiles. Therefore the growth rate effect is driven by the fact that the impact of the business cycle on hire rates does not decline in magnitude with firm quality as much as that of separation rates.

We therefore find that firms who pay less, relative to other firms in their industry fare relatively better in times of high unemployment in terms of their employment growth. An alternative firm categorization would be one that ranks firms by average pay overall, not only compared to other firms in their industry. We prefer the former because the latter will also incorporate other differences across industry that are correlated with pay and business cycle sensitivity. For example, differences in productivity, unionization, worker quality, product cycles, etc., could all generate differences in net and gross worker flows over the business cycle. Most of the theories that we are interested in examining in the next section do not involve these factors.

However, we would like to know whether the results reported above, that are within industry are economically meaningful in the aggregate. We clearly see large differences within industry in how firms of varying quality respond to the business cycle. But these effects could wash out when we average up to the economy as a whole. We therefore replicate our analysis

using a different classification of firm quality based on the overall pay distribution in our sample, still weighted by employment.

These results are illustrated in table 3 and figure 4, and are quite consistent with the previous analysis. The main effects on the unemployment rate are all larger in magnitude. Thus the lowest paying firms across our sample respond more negatively to the business cycle in terms of their growth, hire and separation rates, compared to the lowest paying firms within an industry. The interaction terms for the hire and separation rate analyses are also larger in magnitude and maintain their statistical significance.

The main difference arises in the growth rate analysis. We still find that higher paying firms experience more negative growth rates than do lower paying firms, but by not as much. Also the effects for the highest quintile are only marginally significant. This is probably because we lose some precision in this analysis. Before, all industries had the same share of workers in high- and low-paying firms. In this analysis, since our cut points are overall, that need not be the case. Thus some industries may not have full support over the quality buckets.

However, these results are broadly consistent with what we have seen before. Figure 4 shows that the growth rate effect becomes more negative as firm quality increases, while the hire and separation rate effects become less negative. In fact, separation rates in the highest paying firms increase with the unemployment rate. The declining growth rate effect appears shallower, compared to figure 3, since here the impacts on hire and separation rates are larger in magnitude. Also evident on the graph is that the gap between separation and hire rate effects increases with firm quality. Thus again, we see the larger declines in growth rates in higher quality firms can be accounted for by smaller declines (or increases) in separation rates, and hold despite smaller declines in hire rates.

Thus we believe we have identified an economically important phenomenon at a macro-level. Low-quality firms fare better than high-quality firms in recessions but the comparison group is important: we find stronger effects when ranking firms within industry. This will

inform our discussion in the next section.

4.2 State results

As noted above, we have only a limited time series of data to identify business cycle effects. We therefore also exploit cross-sectional variation in economic conditions using state unemployment rates. This is useful because it yields more variation in business cycle conditions and it allows us to control for other events that coincide with our time series. In fact, our inclusion of date fixed effects allows us to control completely nonparametrically for anything occurring at the national level over this time period. Our results are reported in table 4 and figure 5.

These results are consistent with the national analysis. The main effects of the state unemployment rate are significant and in the expected direction, though they are smaller in magnitude compared to the national effects. This is because we are already controlling for the national economy with our date fixed effects. We might expect the residual local labor market shocks we identify off of to have smaller impacts on firms. The interaction effects are also all of the expected sign and significant at the 1%-level. We do gain substantial precision exploiting this cross-sectional variation.

The effects are easiest to see in figure 5. While we still find a negative growth rate effect (the solid blue line), which increases in magnitude among higher quality firms, the latter effect is shallower compared to figure 3. This is because while we see negative hire and separation rate effects (dashed red and green lines, respectively) that decline in magnitude among higher quality firms, we do not see as large a widening gap between the two effects. That is, separation rates respond less negatively to the business cycle than do hire rates for all firms in our sample, even low quality firms. However, we still find that this gap is largest among firms in the highest wage quintile. Therefore these results reinforce our national findings.

4.3 Elasticities

Our regression specifications yield the impact of the unemployment rate on growth, separation and hire rates in terms of levels. On the one hand, this is useful for our purposes because we can say something about the overall impact on the economy. However, from a comparative standpoint, it is unclear whether changes in levels is the right unit of analysis. Table 1 indicates that firms differ substantially in their average gross flow rates; high-quality firms have much lower hire and separation rates than do low-quality firms. Therefore these firms might be impacted by the business cycle, but on a different scale. Perhaps this is why we find only small impacts on hire and separation rates in high-quality firms. An alternative analysis to address this concern would be in percent changes, since they are unit free and would therefore not be impacted by differences in base levels. We therefore here convert our estimates into elasticities.

Table 5 presents these elasticities for both the national analysis (using within industry cut points) and the state analysis. As in the graphs, we also here use the coefficient on the unemployment rate as the effect at the lowest quality firms. We then add this coefficient to the interaction terms for the other firm quality categories. Thus these elasticities are the total impact of the unemployment rate at each firm quality bucket. For simplicity, we only show elasticities corresponding to the second specifications including the full set of controls.

The first column reports elasticities of the growth rate using the national analysis. We find that in response to a 1% increase in the national unemployment rate, the growth rate declines by 1.6% in the lowest quality firms. The effects increase in magnitude among higher quality firms. For example, growth rates decline by 5.5% in the highest quality firms. The second column, using the state unemployment rate shows a similar trend, though actually the effect in high wage firms is roughly seven times that in low wage firms. These growth rate elasticities mirror our regressions, which is unsurprising given the average growth rates are similar across wage quintiles.

We next analyze separation rate elasticities where we again confirm the regression results.

For a 1% increase in the national unemployment rate, the lowest quality firms respond with a 0.25% decline in their separation rate. This effect falls in magnitude among higher quality firms, to a 0.06% decline among the highest quality firms. Again, the state analysis exhibits similar effects. Therefore, despite the fact that high quality firms have overall lower separation rates, we still find their separation rates are less responsive to the business cycle in percentage terms.

The hire rate effects are less strong. Here we find that the lowest paying firms see a 0.22% decline in their separation rates for a 1% increase in the national unemployment rate. This effect is roughly flat across firm quality bucket. Therefore, in percentage terms, firms see very similar impacts of the unemployment rate on hire rates. Our overall conclusion remains the same: low quality firms are less impacted by the business cycle in terms of their growth rates and this is because they see larger declines in separation rates in times of high unemployment. In percentage terms, they do not experience an accompanying decline in their hire rates.

5 Discussion

We have shown that high-quality firms are more sensitive to the business cycle in terms of their employment growth rates. In contrast, they experience smaller fluctuations in both separations and hires. Therefore, the more negative impact on growth rates at high-quality firms in recessions is being driven by a smaller decline in separations, while hires go in the opposite direction – high-quality firms experience a smaller reduction in hires. Further, we find larger growth rate effects when categorizing firms within industry and smaller growth rate effects (in levels) in response to local labor market conditions. We now attempt to interpret this rich set of patterns, in light of several existing theories of labor market dynamics.

First, since at least as far back as Schumpeter (1939), economists advanced the notion that recessions serve a “cleansing” mechanism, reallocating resources from least to most

productive.¹⁰ Our results on growth rates are strongly counter to this prediction, since, relatively speaking, resources flow to low-quality firms. This relative ability of low quality firms to retain their workforce in recessions could be labor supply driven if the decline in voluntary quits seen in recessions has a larger impact in low quality firms. Alternatively it could be labor demand driven if high quality firms have a greater need to layoff workers. In our data we unfortunately cannot measure whether a separation was voluntary or involuntary. However, we can still provide indirect evidence on both of these channels.

5.1 Labor demand explanations

Why would it be that high quality firms have a greater need to layoff workers? It could be that they experience more cyclical demand in the product market and therefore bear a disproportionate share of aggregate demand shocks. Alternatively it could be that firms bear similar costs of recessions, but high quality firms have more difficulty cutting costs because of rigidities. We explore each of these hypotheses next.¹¹

Differential Business Cycle Sensitivities The cyclical upgrading literature (e.g., Okun 1973 or Bils and McLaughlin 2001) finds that high-paying industries have more cyclical employment; in expansions, workers move from low-wage to high-wage industries, working their way up a quality ladder. It is worth noting that this literature is sparse on the underlying mechanisms driving this result.¹² Instead, it could be that at the time periods studied, the

¹⁰Many theoretical papers seek to explain this pattern by exploiting a friction that inhibits resources from being allocated optimally. In recessions, productivity falls for all firms, thus making the least productive ventures no longer viable. These resources can then be reallocated to more productive ventures. See for example Hall (1991), Mortensen and Pissarides (1994), Caballero and Hammour (1994, 1996) and Gomes, Greenwood and Rebelo (2001).

¹¹Another set of hypotheses are drawn from the literature on internal labor markets and managerial discretion. For example, the “pit stop” view that in booms managers are focused on growth and in busts they are focused on efficiency (see for example Koenders and Rogerson 2005) might predict our finding if low-quality firms are always closer to the margin of survival and must therefore always focus on efficiency. This would result in a relatively greater need for high-quality firms to separate workers in recessions. We unfortunately cannot measure managerial efficiency. However, in future work we hope to explore other accounting variables that might help us proxy for this phenomenon.

¹²In some periods, the effects were driven by the cyclical employment in durable goods manufacturing, which was also higher paying than non-durables. Okun (1973) proposes a model where all sectors benefit

particular aggregate shocks relatively favored low wage industries (for example non-durable construction), and that need not be the case over time. We would also like to point out that our analysis is within two-digit industry, and results are stronger when we rate firms on a within-industry quality spectrum. However, it is certainly possible that the same dynamics could occur within our aggregated industry categories.¹³ For example, consider an inexpensive chain restaurant and a five-star restaurant. The latter most likely pays their employees more and also probably has more cyclical demand.

To investigate this issue, we need information about firm performance beyond employment. We turn to Compustat North America by Standard & Poors, the most complete database of U.S. accounting data.¹⁴ Using Compustat data has the disadvantage that we can only measure balance sheet data for publicly traded companies and we cannot link data to individual firms in our LEHD sample. However, the advantage of Compustat is that the balance sheet data is extremely high-quality, since publicly-traded firms must report these variables to comply with federal regulations. Furthermore, we can link Compustat to LEHD via a disaggregated industry measure, even if we cannot link individual firms. In this analysis, we aggregate Compustat data to the three-digit NAICS industry level and merge it with the LEHD. We can therefore determine whether our wage quintiles are made up of firms in subsectors which typically experience more or less business cycle volatility.

The variable we explore here is percent change in quarterly revenue. Profit maximizing firms will set employment such that marginal cost equals marginal revenue product. Here we take changes in average revenue as a proxy for a firm's incentive to hire. Presumably firms with more cyclical product demand will experience accompanying revenue declines. In each quarter, we take the average percent change in quarterly revenue in the three-digit NAICS

from positive productivity shocks but employment growth in the high-quality sector puts upward pressure on wages in the low quality sector. The high-quality sector does not face wage pressure if it has a queue of workers waiting to fill vacancies. The low-quality sector therefore cannot expand as much. This model does well in predicting patterns during expansions but less well in predicting the mirror image in contractions.

¹³Holmes and Stevens (2012), for example, argue that within manufacturing, small firms are less impacted by trade-driven competition in the product market since they produce to a more niche, local market.

¹⁴We obtain these data via Wharton Research Data Services.

level. We then average these to the wage quintile, weighting by employment.

Figure 6 plots these percent revenue changes for the average firm (solid blue line), firms in the lowest wage quintile (dashed red line), and firms in the highest wage quintile (dashed green line). We also include recession bars. Reassuringly, revenue has a strongly cyclical pattern, falling in recessions and rising in booms. However, the graph shows little differences across low- and high-quality firms over the business cycle. If anything, low-quality firms experienced a larger decline in revenue during the 2001 recession. We can also plot a differential revenue change rate across low- and high-quality firms, analogous to our figure 2, and find no systematic relationship there with the unemployment rate. This figure is therefore inconsistent with the notion that high-quality firms are more sensitive to the business cycle.

Differential Wage Rigidities A long-standing literature (see for example Shimer 2004 or Hall 2005b among many others) points to nominal wage rigidities as an explanation for reduced labor demand in recessions. After experiencing an aggregate negative productivity shock, firms cannot afford to hire or keep workers if they cannot lower wages by a commensurate amount. If higher quality firms face larger downward wage rigidities then they would be forced to make layoffs when they need to cut labor costs. The degree to which nominal wages are downwardly rigid remains a completely open empirical question with evidence on all sides (see Pissarides 2009 for a survey). Our findings that high-quality firms shrink more in recessions and are relatively more likely to make separations, but are also relatively more likely to make hires, is consistent with the literature finding that starting wages are more procyclical than incumbent wages (Martins, Solon and Thomas 2010).

To our knowledge no one has examined wage rigidity as a function of firm quality. In our data, we can measure quarterly earnings for each worker in a firm. We can therefore estimate the degree to which there are downward nominal rigidities in *quarterly* earnings. This measure has both advantages and disadvantages. On the one hand, we cannot measure

whether there exist nominal reductions in pay *rates*, the variable most discussed in the literature. However, our measure incorporates a number of dimensions along which a firm can adjust labor costs besides lowering the base rate of pay, for example hours or overtime and bonuses. This is certainly the more relevant measure for our purposes. Our administrative data is also useful since much of the time the hours data typically used to measure hourly pay rates is often reported with error.

It is unclear *ex ante*, whether we should find a differential strength of earnings rigidities across firms of differing quality. Low-quality firms probably pay more of their workforce on an hourly basis and therefore have the ability to adjust pay by adjusting hours worked, even if they cannot adjust pay rates.¹⁵ However, high-quality firms might have more pay tied to bonuses, which are easy to adjust.

To test whether the strength of downward pay rigidities vary with firm quality, we follow a similar methodology to Dickens et al. (2007). We measure nominal annual pay changes in earnings, Δp_{it} , for job stayers.¹⁶ For a firm, f , in time period t , we then estimate the nominal pay rigidity as per equation 7. That is, for a firm with N workers who have a valid pay change measure, we take the number whose annual pay change was equal to 0 and divide that by the number whose pay change was less than or equal to 0. In practice, we define a pay change to be equal to 0 if it is within $\pm\$1$, to allow for some noise, and results are

¹⁵Another foundation for higher wage rigidities in higher quality firms is found in the queuing literature. Suppose high-quality firms build up a queue of workers who wish to work there, driven for example by an efficiency wage (Akerlof and Yellen 1985), imperfect information (Weiss 1980) or explicit personnel policies (Okun 1973). They would then find it easier to adjust the size of their workforce without adjusting wages. A corroborating piece of evidence comes from the cyclical upgrading literature which finds that wages are more cyclical in low-paying industries (Bils and McLaughlin 2001). This would nicely explain why high-quality firms need not increase wages in expansions, however, is it somewhat less compelling for explaining why a firm would not lower wages in a bust.

¹⁶A worker must have 10 continuous quarters of earnings to be included in the sample. At the quarterly level, issues arise such as differences in the number of pay cycles within a quarter that vary across firm and across calendar year. To avoid additional noisiness, we therefore measure annual pay changes. We also trim the distribution of earnings changes to those who had more than $\pm 50\%$ changes, since these presumably represent errors in reporting.

robust to larger bounds.

$$\text{nominal pay rigidity}_{ft} = \frac{\sum_{i=1}^N \mathbf{1}(\Delta p_{it} = 0)}{\sum_{i=1}^N \mathbf{1}(\Delta p_{it} \leq 0)} \quad (7)$$

This measure proxies for the following: Among workers who were at risk for receiving a nominal pay decrease, what share did not receive one? We find that on average over our time period, this share is roughly 0.25. We then average these within our firm-quality buckets, weighting by average employment, to gain a sense of whether firms of varying quality experience differential pay rigidities.

Figure 7 plots these estimates over time for the average firm (solid blue line), firms in the lowest wage quintile (dashed red line), and firms in the highest wage quintile (dashed green line). We also include recession bars. As can be seen, pay rigidities to have a cyclical pattern, falling in recessions and rising in booms. However, the graph very clearly shows that high-wage firms have a much larger drop in rigidities in recessions. It seems as though high-quality firms can reduce their labor costs by adjusting downwards worker pay.

This measure of downward nominal pay rigidities is far from perfect. For example, we can only estimate the rigidity among stayers at a firm, while we have already shown that high- and low-quality firms differ in their gross separation rates over the business cycle. However, this evidence accompanied with the fact that high-quality firms experience relatively more gross separations in times of higher unemployment is suggestive that high-quality firms are able to adjust labor costs in recessions; they do so using both pay reductions and increased separations, relative to low-quality firms. Therefore, we do not believe differential wage rigidities can be driving our results.

5.2 Labor supply explanations

A labor supply driven explanation is very much in the spirit of Barlevy (2002) who shows that declining job-to-job transitions in recessions has a quantitatively important "sullyng" effect on match quality. A small empirical literature further supports the idea that match quality declines in recessions (e.g., Bowlus 1993, Davis, Haltiwanger and Schuh 1996). The focus here is on match quality, rather than overall firm quality, but we view our analysis as still very much in line with this literature.

To directly address the issue of overall firm quality, Moscarini and Postel-Vinay (2012b) develop a search model in which firms compete for worker talent. High-quality firms can offer more generous contracts and are therefore more successful at attracting workers. In fact, a key element of this model is the "poaching" of workers away from low-quality firms that high-quality firms will engage in during boomtimes. High-quality firms thus spend the boom inflating in size. In busts, they must then "trim the fat" and make layoffs. Low-quality firms, on the other hand, can finally retain their workforce since in the bust high-quality firms will not be poaching.¹⁷

Though we do not test this theory directly, our finding that low-wage firms can grow thanks to experiencing relatively fewer separations in recessions is highly consistent with this theory. Furthermore, we find stronger effects when we rate firms on quality within their industry. If firms are competing for workers it is likely that two firms in the same industry are closer competitors. Thus we view this result as also consistent with the model. Finally, since labor demand explanations were unsuccessful in accounting for our results, we view a search theoretic explanation such as the Moscarini Postel-Vinay model is the most plausible driver of our findings.

¹⁷Along these lines a compensating differentials framework can also yield this result. High-quality firms are better places to work, thus in equilibrium, they might also be more volatile, in order for the marginal worker to be indifferent between working there and a low-quality firm. We find this plausible, and the Moscarini Postel-Vinay model provides a candidate mechanisms driving this increased volatility.

6 Conclusion

In this paper, we use employer-employee matched U.S. data to study net and gross worker flows over the business cycle as a function of firm quality. We find that low-quality firms fare relatively better in the recession; their growth rates shrink by less. This is because separation rates at low-wage firms fall by more. It looks as though high-quality firms are more likely to make layoffs in an economic downturn, while still keeping up a modest amount of hiring. This set of results is consistent with the need for low-quality firms to continually replenish their stock of workers in boomtimes when they lose their workforce to high-quality firms, while in busts they can grow, relative to high-quality firms. In contrast, high-quality firms grow relatively faster in boomtimes and experience relatively more separation in busts. As we have said, these findings are consistent with the Moscarini Postel-Vinay poaching model described above, while we provide ancillary evidence that labor demand explanations cannot be driving our results.

Furthermore, this set of facts is suggestive of two important implications for workers matching in recessions. First, low-quality firms may have an easier time attracting and retaining high-quality workers in a recession. We might therefore see that among workers matching in recessions, workers will be overqualified, relative to the firms that hire them. Second, relatively speaking, low-quality firms have an easier time retaining workers in recessions, since, as we have shown, they shrink less quickly. Therefore a worker matching to a low-quality firm in a recession is likely to stay there for longer; he or she will have less of an opportunity to make a job-to-job transition to a high-quality firm. In our data, we can look at both of these effects directly and we do so in Kahn and McEntarfer (2013).

While previous research has emphasized match quality may decline in recessions due to a lack of workforce reallocation (Barlevy 2002), our evidence here suggests an additional sullyng effect. The types of jobs workers get stuck in are more likely to be low-quality. This is evident in our finding that, relatively speaking, low-quality firms have an easier time growing in the bust, while high-quality firms want to reduce the size of their workforce. One

interpretation of our results is that the reduced ability to move on to better matches caused by a recession has a greater impact on workers in low-quality firms compared to those in high-quality firms. These results have implications then for the costs of recessions, both in the short- and long-run. These results also have important implications for the literatures on the differential impact of recessions of workers. For example, that entering the labor market in a recession (Kahn 2010, Oreopoulos, von Wachter and Heisz 2010) or being displaced from a long-term job in a recession (Davis and von Wachter 2010) has particularly long-lasting, negative wage impacts, could potentially be explained by these workers spending more time in low-quality firms.

References

- [1] Abowd, John, Bryce Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock. “The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators” LEHD Technical Paper No. 2006-01.
- [2] Akerlof and Yellen (1985), “A Near Rational Model of the Business Cycle, with Wage and Price Inertia,” *Quarterly Journal of Economics*, 100, suppl., pp. 823-38.
- [3] Barlevy, Gadi (2002), “The Sullyng Effect of Recessions,” *The Review of Economic Studies*, 69(1), pp. 65-96.
- [4] Bils, Mark and Kenneth J. McLaughlin (2001), “Inter-Industry Mobility and the Cyclical Upgrading of Labor,” *Journal of Labor Economics*, 19(1), pp. 94-135.
- [5] Bowlus, Audra. (1993), “Matching Workers and Jobs: Cyclical Fluctuations in Match Quality,” *Journal of Labor Economics*, 13(2), pp. 335-350.
- [6] Burgess, Simon, Julia Lane and David Stevens (2000), “Job Flows, Worker Flows, and Churning,” *Journal of Labor Economics*, 18(3): pp. 473-502.

- [7] Caballero, R. and Hammour, M. (1994), "The Cleansing Effect of Recessions, *American Economic Review*, 84, pp. 1350-1368.
- [8] Caballero, R. and Hammour, M. (1996), "On the Timing and Efficiency of Creative Destruction," *Quarterly Journal of Economics*, 111, pp. 805-852.
- [9] Campbell, B., H. Chiang, J. Haltiwanger, L. Hunter, R. Jarmin, N. Nestoriak, T. Park, and K. Sandusky (2005), "Firm Performance, Workforce Quality, and Workforce Churn," mimeo.
- [10] Chari, V. V., L. Christiano and P. Kehoe (2007), "The Gertler-Gilchrist Evidence on Small and Large Firm Sales," mimeo.
- [11] Davis, Stephen J., John Haltiwanger and Scott Schuh (1996), *Job Creation and Destruction*, MIT Press.
- [12] Davis, Stephen and Till von Wachter (2011), "Recessions and the Cost of Job Loss," *Brookings Papers on Economic Activity*, Fall.
- [13] Dickens, W., L. Goette, E. Groshen, S. Holden, J. Messina, M. Schweitzer, J. Turunen, M. Ward, (2007), "How Wages Change: Micro Evidence from the International Wage Flexibility Program," *Federal Reserve Bank of New York Staff Reports*, staff report no. 275.
- [14] Farber, Henry (1999), "Mobility and Stability: the Dynamics of Job Change in Labor Markets," in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, vol. 3, North Holland, pp. 2439-2483.
- [15] Fort, Teresa, John Haltiwanger, Ron Jarmin and Javier Miranda (2012), "How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size," mimeo.
- [16] Gertler and Gilchrist (1994), "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *The Quarterly Journal of Economics*, 109 (2), 309-340.

- [17] Gomes, J., J. Greenwood, and S. Rebelo (2001), "Equilibrium Unemployment," *Journal of Monetary Economics*, 48(1), pp. 109-152.
- [18] Hall, Robert (1991), "Labor Demand, Labor Supply and Employment Volatility," in *NBER Macroeconomics Annual* (Cambridge: MIT Press)
- [19] Hall, Robert (2005a), "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years," *NBER Macroeconomics Annual 2005*, pp. 101-137.
- [20] Hall, Robert (2005b), "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95(1), pp. 50-65.
- [21] Haltiwanger, John, Ron Jarmin and Javier Miranda (2010), "Who Creates Jobs? Small vs. Large vs. Young," *Review of Economics and Statistics*, forthcoming.
- [22] Holmes, Thomas and John Stevens (2012), "An Alternative Theory of the Plant Size Distribution, with Geographic and Intra- and International Trade," Federal Reserve Bank of Minneapolis Economic Policy Paper 11-4.
- [23] Hurst, Erik and Benjamin Pugsley (2011), "What Do Small Businesses Do?," *Brookings Papers on Economic Activity*, Fall.
- [24] Hyatt, Henry and Erika McEntarfer (2012) "Job-to-Job Flows and the Business Cycle" *American Economic Review Papers and Proceedings*
- [25] Jovanovic, Boyan (1979), "Job Matching and the Theory of Turnover," *The Journal of Political Economy*, 87(October), pp. 972-90.
- [26] Kahn, Lisa (2010), "The Long-Term Consequences of Graduating from College in a Recession," *Labour Economics*, 17(2): pp. 303-316.
- [27] Kahn, Lisa and Erika McEntarfer (2013), "The Causes and Consequences of Being Hired in a Recession," mimeo.

