Large Employers Are More Cyclically Sensitive

Giuseppe Moscarini
&
Fabien Postel-Vinay

February 2009

Discussion Paper No. 09/609

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN

Large Employers Are More Cyclically Sensitive*

Giuseppe Moscarini†
Yale University
and
NBER

Fabien Postel-Vinay‡
University of Bristol
and
Paris School of Economics

February 2009

Abstract

We provide new evidence that large firms or establishments are more sensitive than small ones to business cycle conditions. Larger employers shed proportionally more jobs in recessions and create more of their new jobs late in expansions, both in gross and net terms. The differential growth rate of employment between large and small firms varies by about 5% over the business cycle. Omitting cyclical indicators may lead to conclude that, on average, these cyclical effects wash out and size does not predict subsequent growth (Gibrat's law).

We employ a variety of measures of relative employment growth, employer size and classification by size. We revisit two statistical fallacies, the Regression and Reclassification biases, that can affect our results, and we show empirically that they are quantitatively modest given our focus on relative cyclical behavior. We exploit a variety of (mostly novel) U.S. datasets, both repeated cross-sections and job flows with employer longitudinal information, starting in the mid 1970's and now spanning four business cycles. The pattern that we uncover is robust to different treatments of entry and exit of firms and establishments, and occurs within, not across broad industries, regions and states. Evidence on worker flows suggests that the pattern is driven at least in part by excess layoffs by large employers in and just after recessions, and by excess poaching by large employers late in expansions. We find the same pattern in similar datasets in four other countries, including full longitudinal censuses of employers from Denmark and Brazil. Finally, we sketch a simple firm-ladder model of turnover that can shed light on these facts, and that we analyze in detail in companion papers.

*We would like to thank Jesper Bagger, Carlos Corseuil, Jason Faberman and David Talan for providing data. The usual disclaimer applies.

†Address: Department of Economics, Yale University, PO Box 208268, New Haven CT 06520-8268. Tel. +1-203-432-3596. E-mail `giuseppe.moscarini@yale.edu`. Web `www.econ.yale.edu/faculty1/moscarini.htm`

‡Address: Department of Economics, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, UK. Tel: +44 117 928 8431. E-mail `Fabien.Postel-Vinay@bristol.ac.uk`. Internet `www.efm.bris.ac.uk/economics/staff/vinay/`. Postel-Vinay is also affiliated with CEPR (London) and IZA (Bonn).

1 Introduction

We present evidence that recessions are times when large employers are hit particularly hard and lose a disproportionate share of their employment, while small employers, as a group, hold up better. This pattern continues well into the subsequent expansion, as long as unemployment remains above trend, and slowly reverses as the unemployment rate declines. Late in aggregate expansions, large employers expand especially fast, both in absolute and relative terms.

Specifically, we establish five facts:

1. Large employers destroy proportionally more jobs during and after recessions and create proportionally more jobs late in expansions (relative to small employers), both in gross and net terms.

2. Employers that are initially larger have a much more cyclical one-year ahead growth rate of employment.

3. The higher cyclical sensitivity of large employers holds principally within industries and within States, not between those units.

4. Reclassification of employers into larger size classes during an aggregate expansion plays a quantitatively negligible role in explaining the higher cyclical sensitivity of large employers.

5. This phenomenon is not unique to the U.S., and we observe it in several countries of different sizes and stages of development.

We also find evidence, more limited in its time scope, that these patterns are related to excess layoffs by large employers in recessions, and to excess job-to-job quits towards large employers in tight labor markets.

In order to establish these facts, we exploit a variety of datasets on employment stocks and flows by size of the employer: repeated cross-sections (distribution of employment among firm size classes); semi-aggregate statistics containing limited longitudinal information, such tabulations of gross and net job flows by initial or end-of-period employer size, divided by initial, average or end-of-period employment; employer panels with full longitudinal information. The data often break down the separate contributions of continuing, entering, and closing establishments and firms, firms of different age, and are broken down by industry and location. Particularly useful proved the new Census Bureau's Business Dynamic Statistics (BDS), as well as matched employer-employee datasets from Denmark and Brazil. The different datasets allow us to address, and relieve concerns about, the effects of two potential sources of bias. The Regression Bias is a well-known fallacy that creates the illusion of a negative size/growth relationship. We are not interested in this sign but in how it changes over the business cycle, and we show that this bias is not relevant in this context. The Reclassification Bias generates the illusion of our Fact #2, as employers are reclassified into larger size bins as the economy grows. Longitudinal data allow us to assess and circumvent this problem, which we find to be quantitatively negligible (Fact #4).

While this is, by and large, a 'facts' paper, ours is by no means a theory-free exercise. In fact, our approach to the data is guided and motivated by our companion theoretical work in Moscarini and Postel Vinay (2008), [MPV08], and Moscarini and Postel-Vinay (2009), [MPV09]. After laying out the facts, we sketch a simple model of firm dynamics, entirely based on hiring and employment turnover frictions, which can simply and parsimoniously explain the new facts. Our findings in this and in our companion papers suggest the following view of how business cycles propagate. Early in an expansion, firms hire mostly and cheaply from unemployment. As the reservoir of unemployment dries out, more productive firms find it profitable to start raising wages to raid workers from less productive competitors. Workers quit mostly from small, less productive, low-paying firms to large, high-paying firms. The growth in the employment of large firms is fueled by the stock of employment at small firms, which takes some time to replenish after a recession. Hence, employment at small firms grows faster and peaks earlier than at large firms.

We now discuss several implications of our findings. The size/growth relationship of employers is the subject of a vast literature originating from Gibrat's (1931) seminal contribution. Firm size is measured by either employment (as in Gibrat's original work) or, more often, assets, capital, or sales. This literature, however, typically ignores business cycle effects. Our findings suggest that firm employment size may predict its growth, but the sign of this relationship may flip depending on cyclical conditions. In particular, the relationship tends to be negative (or less positive) when unemployment is high and positive (or less negative) when it is low. Omitting cyclical indicators may lead to conclude that, on average, these cyclical effects wash out and size does not predict subsequent growth, which is Gibrat's law. This is typically captured by a zero estimate in a growth-size regression. Adding an interaction term between the detrended unemployment rate and initial size in the regression is likely to invalidate this conclusion. Verifying this conjecture is a task on our current agenda.

The new business cycle fact that we uncover is reminiscent of Okun's (1973) idea of *Cyclical Upgrading* (see Bils and McLaughlin, 2001 for a recent new interpretation), a cross-industry pattern whereby labor reallocates itself from low- to high-paying industries in booms, and vice versa in recessions. Instead, the phenomena that we emphasize in this section hold within industries and not across. This is surely worth noticing, although it does not pose a particular problem for our proposed interpretation, which seems to apply equally well to many industries. It is in fact natural to expect that, if workers have any significant attachment to an industry, for example because of their industry-specific occupation, then they should upgrade within industries more than across.

Our facts appear to plainly contradict a well-established set of facts regarding the sensitivity of small firms to cyclical conditions and monetary shocks. In a very influential paper, Gertler and Gilchrist (1994) present evidence that small firms, which they argue are more credit constrained, are more sensitive to Romer and Romer (1989) monetary policy shocks. Gertler and Gilchrist use the Quarterly Financial Report for Manufacturing Corporations (QFR), since 1958, which is confined to manufacturing. They define size in terms of nominal sales, raising the issue of industry-specific price indexes within manufacturing. The data lack longitudinal links, leading them to make an ingenious yet ad hoc correction to avoid the Reclassification Bias. Even then, their conclusion that small firms are more cyclically sensitive holds, if any, for Romer shocks, which are notoriously controversial, and not for NBER-dated recessions. Of the six recessions in their sample period, only in 1970 one sees a clear collapse in the growth rate of sales at small firms relative to large ones (their Fig. I), and the opposite occurs in 1982. The other four episodes appear fairly neutral.[1] Chari, Christiano and Kehoe (2007) also notice the distinction between Romer dates and NBER dates in the Gertler and Gilchrist QFR time series, and extend them to the early 2000s. They focus on the growth rate of firm sales by size of a firm's assets, and do not find a differential behavior of large and small firms around NBER recessions, but rather a much higher sensitivity of small firms to Romer shocks. Our measure of performance and size is employment, not sales and capital, because we are interested in job creation and inputs more generally, and we have employment growth data that are immune from Reclassification Bias. All in all, we believe that the picture of the whole economy over NBER-dated cycles for the post-war period, after controlling for Reclassification, is likely to be much closer to our conclusions that to Gertler and Gilchrist's. In the period we cover, since the mid-1970s, our facts about firm size and growth are very firmly established.

The conventional wisdom and oft-heard slogan that "small businesses are the engine of job creation" finds some empirical support in our data for periods of high unemployment, recessions and their aftermaths, which is presumably when jobs are more needed. This

[1]Sharpe (1994) replicates their findings for employment growth, by initial size defined in terms of net capital. He uses the NBER Manufacturing Panel from Compustat. As we will show later, in the full Compustat panel comprising all industries, over a longer time period, our pattern of differential growth rate by initial size still emerges at NBER-dated business cycles.

statement clearly fails in tight labor markets, when job creation is taken over by large employers. The Regression Bias is not a severe problem because we focus on cyclical, rather than average, patterns of job creation by firm size.

Finally, our findings may explain at least in part the discrepancy between household-based (Current Population Survey, CPS) and payroll-based (Current Employment Statistics, CES) measures of aggregate employment at monthly frequency. The former surveys about 50,000 households with 140,000 individuals, and generates the unemployment rate, the latter about 160,000 firms with 400,000 establishments, and generates official employment figures. Both are run by the BLS. The CES/CPS ratio of employment measures is strongly procyclical, very much like the share of large employers or the relative growth rate of large vs. small employers.[2] This discrepancy in levels and especially cyclical behavior between the two surveys of employment is still a mystery and has defied all attempts for an explanation. None of these explanations is based on the evolution of the firm size distribution. If the CES sample is skewed towards large employers and misses small firms, that the CPS household survey is able to capture, then CES employment will reflect the cyclical patterns of the employment of large employers. When unemployment is low, large employers, which are better captured by CES, appear to grow very fast. When unemployment is high, the reduction in CES employment is very severe, in part because it misses small employers that we have shown to be less cyclically sensitive, as a group. Given the well-know difficulties that both the BLS and the Census Bureau encounter in tracking small employers, our new finding that large employers are more cyclically sensitive appears to offer a promising avenue to resolve the puzzle.[3] More work needs to be done in this direction.

In Section 2 we present our definitions and methodology, and discuss two potential biases that can affect our results, and how we cope with them. In Section 3 we present evidence drawn from business micro data that contain a longitudinal dimension and are immune from the Reclassification Bias, and we document our main finding. In Section 4 we present auxiliary evidence from repeated cross-sections of employment distribution by size of the employer. In Section 5 we shed some light on the new patterns by introducing evidence on worker flows, specifically job-to-job quits and layoffs, by size of the firms involved in the transition. Section 6 presents evidence from several other countries. Section 7 illustrates a simple model of firm size dynamics. Concluding remarks follow.

[2] Bowler and Morisi (2006).

[3] We thank Bob Hall for suggesting a possible connection.

2 Definitions and Methodological Issues

The main purpose of this paper is to show that *large employers are more cyclically sensitive*. In this section we lay out the definitions and methodology that we adopt to establish this fact. In the following sections we apply them to various datasets.

Our notion of size is employment, not capital, assets, or sales. This choice is motivated by our previous theoretical work (MPV08, MPV09), which identifies in a firm's productivity and employment level the two main determinants of an optimally posted wage. Since productivity is hard to measure, and proposed measures are highly correlated with employment size, in this paper we focus on the latter. By "employers" we mean either firms or establishments, depending on the dataset at hand.

In order to measure relative cyclical sensitivity of employers of different sizes, we compute the difference in employment growth rates between large and small employers, each taken as a group. Employers of different sizes may add systematically more or less jobs, an issue of great conceptual confusion and political importance. By taking the difference in growth rates and focusing on its fluctuations, rather than on its level, we sidestep this issue of relative contributions of small businesses to job creation. We will provide conceptual details and empirical evidence on this point.

Our measure of the economy's business cycle conditions is the detrended civilian unemployment rate. Again, this is motivated by our theoretical work. We correlate this statistics with the differential employment growth rate by employer size described above, within the relevant geographical unit. We also plot the two series with NBER-dated cyclical peaks and troughs, episodes which are very concentrated in time and fail to capture the slack in labor markets that persists long after a recession. Nonetheless, we document that the last three NBER-dated recessions, and probably also the current one, are clearly times when large employers fare worse than small employers.

As we have learned from the literatures on economic growth across countries and on firm size growth, specifically Gibrat's Law (Sutton, 1997), the size/growth relationship is rife with statistical fallacies. Great care must be taken in defining the timing of observation of size and growth for individual employers. Does size refer to the period before observing growth, after growth, an average of the two? Our ideal measure of employer size is one that is not contaminated by subsequent growth. That is, we would like to interpret employer size as a predictor of its subsequent growth, and in particular how this predictive power depends on aggregate economic conditions. One statistical fallacy is sample selection, in that the sample is not representative of the entire population of interest but is skewed towards employers that are particularly large or small at the end of the sample. This is not an issue for us as we will mostly exploit either censuses or representative samples of employers. A more serious issue is mean reversion in size.

2.1 The Regression Bias

Let L_{it} denote the number of employees working for employer i at (discrete observation) time t, and define a weighted-average size between $t-1$ and t:

$$L_{it-1}^{(\alpha)} = \alpha\left(L_{it}, L_{it-1}\right) \cdot L_{it} + \left[1 - \alpha\left(L_{it}, L_{it-1}\right)\right] L_{it-1}$$

Here $\alpha : \mathbb{N}^2 \to [0,1]$ is the weight on size "after the fact", and $1 - \alpha\left(L_{it}, L_{it-1}\right)$ on initial size. This is a weighting function that can depend on both numbers.

Let the weighted growth rate

$$g_{it}^{(\alpha)} = \frac{L_{it} - L_{it-1}}{L_{it-1}^{(\alpha)}}.$$

If employer size is mean-reverting, then small employers will tend to grow more than large ones, as they all converge back to a long run middle ground. This generates the illusion of a negative size/growth relationship if one uses the conventional measure of growth rate $g_{it}^{(0)} = L_{it}/L_{it-1} - 1$.

An alternative is to divide growth by eventual size and use $g_{it}^{(1)} = 1 - L_{it-1}/L_{it}$, but here the contamination of growth with size is quite strong. As we will see, for our purposes $\alpha = 1$ is probably much more problematic than $\alpha = 0$, as it gives rise to another issue, the Reclassification Bias, on which we elaborate below.

A middle ground, (see e.g. Davis et. al. 1996) is to use $g_{it}^{(1/2)}$. That is, the base for the growth rate between times $t-1$ and t is the average of employment at $t-1$ and t. This reduces the mean reversion and Reclassification problems, and has the added value of being well defined both for entrants, who have $L_{it-1} = 0 < L_{it}$, and closing employers, who have $L_{it-1} > 0 = L_{it}$.

Let $\overline{L} > \underline{L} > 0$ two integers that define "large" employers ($L_i^{(\alpha)} \geq \overline{L}$) and "small" employers ($L_i^{(\alpha)} \leq \underline{L}$). Let the growth rate between $t-1$ and t of employment at all employers that are classified as large at time $s \leq t$

$$g_{s,t,LARGE}^{(\alpha',\alpha)} = \frac{\sum\limits_{i:L_{is}^{(\alpha')} \geq \overline{L}} \left(L_{it} - L_{it-1}\right)}{\sum\limits_{i:L_{is}^{(\alpha)} \geq \overline{L}} L_{it-1}^{(\alpha)}}$$

Notice that, using the weighting α', we can choose the size class over which to compute net job creation (the numerator) by either initial, average, or eventual size observed at any given date $s < t$. This initial date could be fixed once and for all with longitudinal data, or re-assigned every period, either before growth occurs ($s = t - 1$) or after ($s = t$). Also, we can assign the base by which to divide job creation, the denominator, independently of the numerator, with the weighting function α, which can differ from α'. Again, this allows employment to be either initial, average, or eventual size, at the same time s. Similarly for every size class, including the smallest $g^{(\alpha',\alpha)}_{s,t,SMALL}$ for $L^{(\alpha)}_i \leq \underline{L}$.

We are interested in the relative growth rate by size class, defined as the difference in growth rates between large and small employers

$$\Delta g^{(\alpha',\alpha)}_{s,t} = g^{(\alpha',\alpha)}_{s,t,LARGE} - g^{(\alpha',\alpha)}_{s,t,SMALL}$$

and in particular how it correlates with the business cycle conditions at $t - 1$.

Whether mean reversion affects also the *cyclicality* of this *relative* growth rate $\Delta g^{(\alpha',\alpha)}_{s,t}$ is questionable, and depends on the specific statistical model of firm growth implied by the underlying structural model. At any rate, to err on the side of safety, in some datasets we use both $\Delta g^{(0,0)}_{t-1,t}$, which is differential growth between initial size classes, and $\Delta g^{(0,1/2)}_{t-1,t}$, differential growth between employers classified by their initial size but where the growth rates are computed dividing net JC by average employment over the period. We will show that the cyclical behavior of these two measures is essentially the same, suggesting that mean reversion is not a problem for our purposes. With longitudinal data, we compute $\Delta g^{(0,0)}_{t_0,t}$, where t_0 is the time when the dataset begins, for many years t after t_0, when the effects of mean reversion would have presumably washed out. As we will see, our empirical results are robust.

2.2 The Reclassification Bias

Consider $\Delta g^{(1,\alpha)}_{t-1,t}$, which is relative growth between $t - 1$ and t of employers classified by their size at t. If the economy grows, and all employers with it, while the size cutoffs $\overline{L} > \underline{L}$ remain time-invariant, employers tend to grow in size with the economy and to jump into higher and higher bins. It then appears that more and more job creation is attributed to larger size classes, and this differential growth rate is more likely to be positive. Conversely when the economy shrinks. This is a bias that produces the illusion of procyclical relative growth by size, precisely the fact that we aim to document.

In one of the datasets that we exploit, the Census Bureau's Business Dynamic Statistics, we know employment growth for a size class j by either initial size, $\sum_{i:L_{it-1}\in[L_j,L_{j+1})} (L_{it} - L_{it-1})$, or final size. We also observe employment shares of size classes, whose growth rates we can compare over time. Employment shares suffer from an even more obvious reclassification problem. When using the three methods, we find that the results are similar, suggesting that neither bias changes our results qualitatively.

An alternative method of attribution of growth to size class is followed by the Bureau of Labor Statistics in computing quarterly Job Creation (JC) and Job Destruction (JD) rates presented in the Business Employment Dynamics dataset, that we exploit and illustrate below. Recall that JC is the addition of employment positions at all units that expand, and vice versa, so JC−JD is net job creation. The BLS "Dynamic allocation" method even changes class assignments at infra-quarterly frequency for firms crossing the line between two size classes. For example, if firm i has $L_{it-1} = 7$, $L_{it} = 15$, then, of the 8 jobs created on net, 2 are attributed to the size class [5,9] and 6 to the size class [10,19]. So the weighting function α' is $\tilde{\alpha} = \left(L_{it} - L_{j(i),t}\right) / \left(L_{it} - L_{it-1}\right)$ where $L_{j(i),t}$ is the size class cutoff that falls in (L_{it-1}, L_{it}), if indeed the firm jumps size class, otherwise it is just $\alpha' = 0$. The denominator in the published BED job flows rate is the average $L_{it-1}^{(1/2)}$, so the relative growth rate obtained by subtracting their JD rate from their JC rate for large and small classes is approximately $\Delta g_{t-1,t}^{(\tilde{\alpha},1/2)}$.

This method is also potentially vulnerable to the Reclassification Bias, both in the numerator (dynamic allocation) and the denominator (average employment). We obtained from the BLS also employment by size in each period, so we can correct the latter problem, potentially incurring into mean reversion. Once again, we find little difference in the results.

We also use fully longitudinal business micro data, Compustat for the U.S. public companies (a selected sample) and censuses of Danish and Brazilian employers, where we can fix size class for each employer once and for all at the beginning of the sample t_0 and compute $\Delta g_{t_0,t}^{(0,0)}$. This takes care not only of mean reversion but, even more strongly, of Reclassification, as employers never change size class, even after one or two decades. We find that initial size predicts growth many many years later, in a way that depends on cyclical conditions.

2.3 Employment Shares by Size Class

The easiest data to obtain are repeated cross-sections of employment shares of size classes. That is, if we only know

$$e_{jt} = \frac{\sum\limits_{i:L_{it} \in [L_j, L_{j+1})} L_{it}}{\sum\limits_{i} L_{it}}$$

which is the conventional definition of the share of employment at time t working at employers of size in class j, then the change in e_{jt} is an estimate of employer growth for size class j. When small firms grow faster, their share rises. However, as small employers gain size, they are reclassified into larger size classes, so repeated cross-sections are subject to

Reclassification Bias. Specifically, consider the following decomposition:

$$e_{jt} = \frac{\sum\limits_{i:L_{it}\in[L_j,L_{j+1})} L_{it}}{\sum\limits_i L_{it}} = \frac{\sum\limits_i L_{it-1}}{\sum\limits_i L_{it}} \cdot \frac{\sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it-1}}{\sum\limits_i L_{it-1}} \cdot \frac{\sum\limits_{i:L_{it}\in[L_j,L_{j+1})} L_{it}}{\sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it-1}}.$$

Defining the growth rate of aggregate employment as $g_t = \sum\limits_i L_{it}/\sum\limits_i L_{it-1} - 1$, the latter equality can be recast as

$$\frac{e_{jt}}{e_{jt-1}} = \frac{1}{1+g_t}\left(\frac{\sum\limits_{i:L_{it}\in[L_j,L_{j+1})} L_{it} - \sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it}}{\sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it-1}} \right.$$
$$\left. +1 + \frac{\sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it} - \sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it-1}}{\sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it-1}}\right).$$

Since all these numbers are all well below 1 at typical sampling frequencies, to a first order we can write:

$$\frac{e_{jt}}{e_{jt-1}} \simeq 1 + g_{jt}^{(0)} - g_t + \frac{\sum\limits_{i:L_{it}\in[L_j,L_{j+1})} L_{it} - \sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it}}{\sum\limits_{i:L_{it-1}\in[L_j,L_{j+1})} L_{it-1}}. \tag{1}$$

The growth rate of the employment share of a given size class j has thus been decomposed into the sum of three terms: $g_{jt}^{(0)}$, reflecting growth of firms that were initially (i.e. at date $t-1$) in size class j, minus a normalization by aggregate employment growth g_t, which drops out when taking differences in growth rates across size classes, plus a third term reflecting the contribution of employers that either enter or exit size class j between dates $t-1$ and t. Omission of that third term, which amounts to ignoring changes in the composition of size classes, is the source of the Reclassification Bias: that term would only equal zero if the identity of employers in a given size class did not change, so all growth and decline occurred without leaving a size class.

Equation (1) clearly shows that the growth rate of employment shares only approximates $g_{jt}^{(0)}$ up to the Reclassification Bias. Nevertheless, we will also present evidence based on employment shares for two reasons. First, data on employment shares by size classes are relatively easy to come by in many countries and over reasonably long time periods. Second, richer (and scarcer) data sets that allow to fix the composition of size classes over two consecutive periods also allow to compute growth rates of employment shares, which provides a way to gauge the magnitude of the Reclassification Bias. Our results suggest that it is small.

An example. Four firms, A through D, are observed at two consecutive dates, $t-1$ and t. Their employment counts are:

	Firm			
	A	B	C	D
$t-1$	36	44	353	1690
t	48	1566	18	723

Suppose the data are presented and organized in three size classes, $[1-49]$, $[50-999]$, $[1000,\infty)$ employees. The first class is "small", the last is "large". Then the growth rate of firms classified by initial size is $(48-36+1566-44)\,/\,(36+44) = 1917\%$ for the small class, $(723-1690)\,/1690 = -57.2\%$ for the large, the large minus small difference is -1974%. If we classify them by average size, the growth rate is (48-36)/36=33.3% for the small class, as only firm A has average size in the small category, again -57.2% for the large class, difference -90.5%. If we apply BLS dynamic size allocation, the growth rate of the small class is $2\,(48-36+49-44+18-50)\,/\,(48+36+44+49+18+50) = -11.7\%$ and for the large class it is $2\,(1566-1000+999-1690)\,/\,(1566+1000+1690+999) = -4.75\%$, difference 6.9%. Finally, the employment shares of the small and large class are (resp.) 3% and 79% at $t-1$ and 2.8% and 66.5% at t, so they both decrease, but the share of large firms declines by much more. Qualitatively, all methods except dynamic allocation suggest that small firms are doing relatively better, a pattern that we uncover in recessions. Quantitatively, the results vary widely, but the example is intentionally extreme.

2.4 Entry and Exit: Firms vs. Establishments

Most firms enter at the bottom of the size distribution in the dataset, and the contribution of entry to net JC typically declines by size. For establishments these patterns are similar but weaker, because existing firms open establishments of all sizes. A natural question to ask is whether the higher cyclical sensitivity of large employers is due entirely or in part to entry.

The datasets at our disposal are organized at different levels. If the dataset contains information on job flows and size either at the firm or at the establishment level, then one natural question is whether we want to include entering and exiting employers from the computation of relative growth rates.

There is no compelling reason to exclude exiting employers if we know their initial size class, as their job destruction should definitely be imputed to that class. But we will provide some evidence with and without exiting employers, to check whether the relative growth performance over business cycles is driven by different patterns of exit (extensive margin) or net JC at continuing employers (intensive margin) across size classes.

For entrants, the main issue is the attribution of their first growth to initial size, which is zero by definition. Even in longitudinal datasets, where we are able to classify JC and JD by initial size, for entrants this attribution is done necessarily by their eventual size class (the one they have when first observed), ruling out a size class "0". So the Reclassification Bias is unavoidable for such entrants, and may work in the opposite direction: if entry is very procyclical and most entrants are small when first observed, then the contribution of small employers to employment growth will appear to be more procyclical, going against our stylized fact. Therefore, we present evidence with and without entry. If we fix employer sizes once and for all at some initial date, then necessarily we exclude subsequent entrants and we focus on a sample of continuing employers.

A special issue is raised by the Census BDS dataset, where size refers to the parent company, but the entry/exit/continuing status by which job flows are broken down is that of each establishment the firm owns. For example, we find JD between years $t-1$ and t of all closing establishments owned by firms of either initial (at $t-1$) or final (at t) size in the same employment size class $[500-999]$. It is possible that an existing firm decides to open or close establishments, while, by definition, an entering firm owns only entering establishments and a firm that is shutting down is closing all of its establishments.

Here we define "firm count" to be the situation where we subtract from both net JC and initial employment of an initial-size ($L_{it-1}^{(0)}$) class the job growth of firms that were of age 0 at $t-1$, that is, firms that entered between $t-1$ and t and are reported as having an "initial size" $L_{it-1}^{(0)}$ equal to their post-entry size $L_{it}^{(0)}$. The reason is that the "initial size" of these entrants in BDS is in fact their final one, that they reach after entering. This convention introduces a Reclassification Bias both in the numerator (net JC) and the denominator (initial employment) of $\Delta g_{s,t}^{(0,0)}$. We denote the firm-count adjusted growth rate differential by $\Delta \widetilde{g}_{s,t}^{(0,0)}$. In the "firm count", the addition of new jobs by a firm of size, say, 700, which opens a new establishment is entirely attributed to size class $[500-999]$, because the firm existed before in that class, although the establishment did not. We define "establishment count" the situation where we subtract from both numerator and denominator the net JC of new establishments. This excludes both new firms *and* new establishments of existing firms. The first, firm count is our main focus, and establishment count is just to check robustness and investigate how much of the firm count pattern is driven just by entry of new establishments.

2.5 Large and Small Employers

Our definition of a "large" and "small" employer is necessarily arbitrary, but guided by a few considerations. First, if the size classes in the data refer to establishments, then we choose what appears to be fairly reasonable cutoffs, either 20 or 50 employees as upper bounds for small establishments and 500 employees as lower bound for large ones, which is also often the largest cutoff available in the data.

If the size classes in the data are organized at the firm level (total employment of the firm), sometimes the highest cutoff in the data partition is as large as 10,000 employees. Whenever possible we proceed as follows. Consider firm size class j characterized by cutoffs $[L_j, L_{j+1})$, for example [1,4] employees, [500,999] etc. Let $\mathfrak{s}_j$ denote the average number of employees per establishment among all establishments that are owned by firms in size class $[L_j, L_{j+1}-1]$. Clearly, $\mathfrak{s}_j/L_{j+1} < 1$, because by definition no firm in size class j has more than $L_{j+1}-1$ workers, and no firm can be smaller than any establishment it owns. If all (or most) firms in a size class are mono-establishment, then $\mathfrak{s}_j/L_j \geq 1$, and we call these firms "small". If instead most firms in a size class have multiple establishments, then the firm's labor force is distributed among many locations, and the average size of these establishments is smaller than that of the parent firm, which must be the case if in particular $\mathfrak{s}_j/L_j < 1$. We call "large" the firms that belong to size classes that satisfy the last inequality. Presumably, most single-establishment firms do not have a Human Resource department and do most of their hiring in house, while multi-establishment large firms have a HR department. So it is plausible that the former have a harder time than the latter plucking employees from other firms, and have to rely more on the unemployment pool. This differential source of hiring, unemployment for smaller firms and poaching for larger firms, is the core of our theoretical explanation for the patterns that we document in this paper (MPV08, MPV09). Hence, we try to present the data in a way that is congruent with our theory, distinguishing between small, single-establishment firms that hire mostly from unemployment and large firms that own multiple large establishments and rely more on poaching (we will also present direct evidence on that).

Finally, when we use less aggregate data and compute relative employment growth, say by sector, we compute sector-specific cutoffs by considering, among the size classes available in the data partition, the bins that are closest to the first and fourth quartiles of the average distribution of employment over a long time period.

2.6 Cyclical Indicators and Detrending

As mentioned, our preferred cyclical indicator is the civilian unemployment rate, that we detrend over the post-war period and then correlate with the relative firm-count-adjusted growth rate by size one period ahead. So our main statistics is corr $\left(\Delta\widetilde{g}_{s,t}^{(0,0)}, u_{t-1}^{\text{detr}}\right)$. As we showed in MPV08, using the employment/population ratio instead makes little difference, as participation has no clear cyclical pattern.

From datasets of repeated cross-sections, the employment distribution by size of the employer, we detrend the growth rates of employment shares, as we find they have significant long term movements. This raises the issue of detrending method. We use a Hodrick Prescott filter. Following Shimer (2005) we use a high smoothing parameter for unemployment (10^5 at quarterly frequency). For shares' growth rates, we find that a high smoothing parameter is also necessary so that the trend loses obvious cyclical fluctuations. Fitting a linear trend makes little difference in this case.[4]

3 The U.S. Employers' Growth/Size Relationship over the Business Cycle

3.1 Evidence from the Business Dynamic Statistics (BDS)

3.1.1 The BDS Dataset

The primary source of information on the identity, location, employment, sales, payroll, industry, of all U.S. businesses is the Census Bureau's Business Register (BR), a.k.a. Standard Statistical Establishment List (SSEL). For establishments in the BR, information is updated annually with the Company Organization Survey. The list of establishments in the BR is updated at every quinquennial Economic Census.[5] Establishments in the BR can be linked up to the firm level using the Employer Identification Number (EIN) assigned by Internal Revenue Service. County Business Patterns (CBP) are semi-aggregate statistics from the BR publicly available annually in 1977-2006. They report number of employees in the week that includes March 12 of each year, payroll in the calendar year to date and over the entire previous 12 months, and number of establishments, all by size of the parent company (in number of employees, one of nine size classes), by industry, state and county.

Starting in December 2008, the Census Bureau has made publicly available a set of semi-aggregate statistics from the Business Register, under the name of Business Dynamic Statistics. BDS covers approximately 98% of U.S. employment, so it is a fairly comprehensive picture, not a survey. BDS contains information on establishment-level employment stocks and job flows, for continuing, entering, and exiting establishments, at annual frequency for

[4] The Census publishes County Business Patterns aggregate statistics that include the number of firms, but only starting in 1989. We find that, in 1989-2006, employment has been migrating towards larger companies. The average size of the U.S. firm has risen from 17.7 to 19.91 employees. In the BLS Business Employment Dynamics, average firm size increased from 20.7 to 22.1 employees in 1992:III-2008:I. An inspection of the entire time series reveals that these are not purely cyclical effects. To the best of our knowledge, this observation is new. Because larger firms are less typically volatile, this shifting composition of the universe of U.S. employers might in part explain the simultaneous decline in business volatility, documented by Davis et al. (2008a) without controlling for size composition.

[5] Visit `www.census.gov/epcd/susb/introusb.htm` for more details.

the 1977-2006 period, broken down by location and industry of the establishment, and by age and size of the parent *firm*. To obtain this size, the Census aggregates employment at all establishments (the original unit of observation) owned by the same firm, according to various criteria, mainly Employer Identification Number. Two notions of firm size are available in BDS: the size of the firm at the end of the observation year, after the job flows have taken place, and the size of the firm *before* the flows are measured, the "initial firm size", which is zero for entrants. These two allow us to avoid the Reclassification Bias and to calculate employment growth by initial size of the employer, as well as to assess whether the Regression Bias matters for cyclical patterns.[6]

More specifically, we calculate the growth rate of employment in a size class as the ratio between net job creation — namely gross Job Creation (JC) minus gross Job Destruction (JD) — over the period [March of year $t-1$, March of year t], and a measure of employment. Everything is classified by a measure of initial firm size, as of March of year $t-1$. After computing growth, firms are reclassified into their new size classes, and their new size becomes their initial size for the following period, March of year t to March of year $t+1$. As discussed in Section 2, we consider the following two alternative definitions of employment: either initial employment at $t-1$ or the arithmetic average of employment between $t-1$ and t. When using all establishments, the two notions of employment yield almost identical results, so we will report them only for the former. When distinguishing between continuing and other establishments, in the next section, the choice of the latter definition becomes mandatory. Notice that gross JC excludes, by definition, exiting establishments, because it is the sum of all employment gains at establishments that grow in size, and an exiting one declines in size.

As explained earlier, in the firm count our notion of gross (and net) JC refers only to pre-existing firms, either continuing or exiting. These are not to be confused with continuing and exiting *establishments*, whose gross job flows are provided directly in the BDS, but are not the appropriate measures for our classification by size of the firm. That is, we include in net JC that of new establishments of existing firms (age>0), because the initial size of those pre-existing firms was well defined and correctly attributed. Obviously, we also include in JD both continuing and exiting establishments, because they existed at time $t-1$ and had a correct firm size attached.

To define small and large firms, we choose employment size classes 1-49 and 1000+. Besides being a reasonable choice, we do so after looking at Table 1, where we report the ratio

[6] In previous work (MPV08) we used another set of semi-aggregated statistics from the Census Bureau which contains some longitudinal information, the Business Information Tracking Series. This is now by and large also subsumed in time and scope by BDS.

$\mathfrak{s}_j/L_j$, discussed earlier, by size classes j. We find that the ratio is greater than 1, suggesting a prevalence of mono-establishment firms, for firms of sizes up to 49, and less than 1 for larger firms. The ratio nearly vanishes for very large firms, implying that these are almost exclusively firms that own very many establishments. Indeed, the average establishment size stabilizes around 60 employees for all firms of total size above 500.

Firm size category L_j to $L_{j+1}-1$	Mean establishment size $\mathfrak{s}_j$	Size ratio $\mathfrak{s}_j/L_j$
1 to 4	2.2	2.265
5 to 9	6.5	1.300
10 to 19	12.5	1.251
20 to 49	24.1	1.203
50 to 99	39.2	.784
100 to 249	49.0	.490
250 to 499	53.9	.215
500 to 999	57.9	.115
1000 to 2499	62.1	.062
2500 to 4999	58.6	.023
5000 to 9999	55.8	.011
10000+	62.5	.006

Table 1: Ratio between average establishment size and smallest possible size of its parent company, by size class of the parent company

3.1.2 The Aggregate Picture

In Figure 1 we plot the employment growth rates of firms that have initially less than 50 and more than 1000 employees, starting in 1979, with NBER peaks and troughs. Large firms suffer much more from the 1983 and 2001 recessions. The growth rate of large firms is smaller than that of small firms in the first few years after each of the three NBER troughs, including 1992-1995, and larger in the few years preceding each of the three NBER peaks. The 1991 recession appears to differ, but this is mostly a time aggregation effect, as large firms do fare worse in 1992 (and beyond). The NBER trough was in March 1991, which is exactly the month to which BDS observations refer for the year. In fact, if we compute two-year growth rates 1990-1992, the difference disappears, and reverses in 1990-1993, when large firms once again grow less. This picture corroborates only in part the common wisdom that small businesses are the engine of (net) job creation. As mentioned earlier, this statement has been amply criticized (see for example Davis et al., 1996) as subject to Galton's regression fallacy, which is known to generate a negative size-growth relationship. While this fallacy is almost certain to affect the relative *levels* of net JC by initial size also in our BDS data, it is less likely to have a clear impact on their relative *cyclical patterns*, as this would require that idiosyncratic, firm-level shocks become much more mean-reverting in aggregate slumps, which is fairly implausible. At any rate, small firms appear to create more jobs as a fraction of their employment only when unemployment is high (which is, arguably, when jobs are most needed). In terms of absolute number of jobs added, large firms dominate at nearly all times, because they employ a larger fraction of employment to begin with.

Figure 2 decomposes the differential (initially large minus small firms) net JC rates of Figure 1 into differential gross JC and JD rates. The action is all on the JD side in 1983, namely, large employers that contracted shed a much larger proportion of their payroll than contracting small firms. Less dramatic but qualitatively similar is the pattern after the 1991 recession, with a lag of a year or so. In 2001 both JC and JD contribute to the worse performance of large employers. In between recessions, it is a surge in gross JC by large firms late in expansions to account for their better performance in those phases, including again in 2005-2006. Overall, these figures indicate the following:

> **Fact 1.** *Large employers destroy proportionally more jobs during and after recessions and create proportionally more jobs late in expansions (relative to small employers), both in gross and net terms.*

To better visualize the cyclical effect, we take the difference in growth rates between (initially) large and small firms. If this differential growth rate is procyclical, as we will find, then large firms are more cyclically sensitive: they shed proportionally more of their employment in recessions and gain more in booms. Because our underlying theoretical hypothesis emphasizes the role of unemployment and labor market tightness as a cyclical indicator, and because NBER dates may be marking cyclical slumps and recoveries too narrowly, Figure 3 reports the relative growth (net JC) rate against the detrended one-year lagged civilian unemployment rate. We lag unemployment as we envision it as predetermined and contributing to cause growth of employment between $t-1$ and t. Here is one central finding of this paper: as visually clear and confirmed by the -0.6 correlation,

> **Fact 2.** *Employers that are initially larger have a much more cyclical one-year ahead growth rate of employment.*

An alternative way of cutting the data and focusing on recession is an "episode analysis". For each of the three complete business cycles in the time span, we renormalize time at 0 in the corresponding NBER trough year, and plot it in a worm graph for several years before and after that trough, without any prior detrending. We center the 1989-1995 episode in 1992, although the trough was in March 1991, because our 1991 datapoint refers to employment

Solid = Net job creation rate, large firms, Dash = Net job creation rate, small firms.
Categories defined each year as <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BDS and authors' calculations.

Fig. 1: Net job creation rates

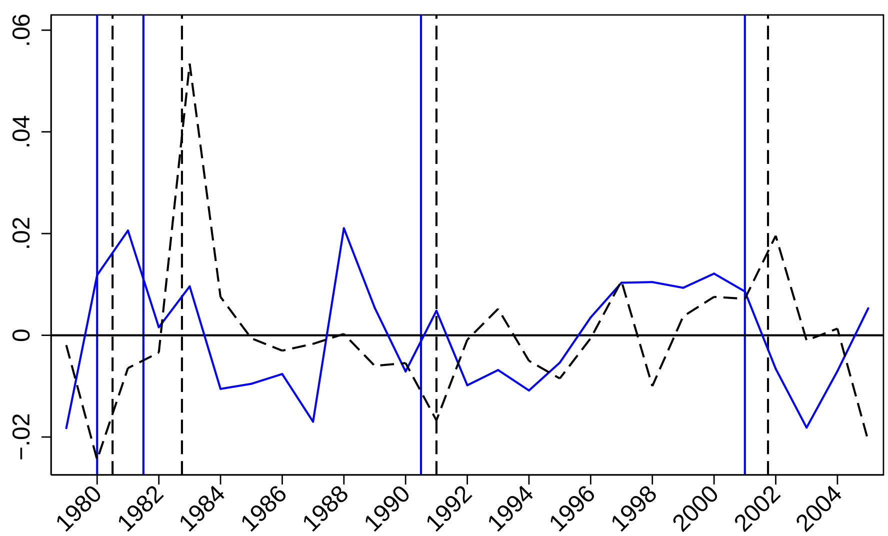

Solid = Differential gross job creation, Dash = Differential gross job destruction.
Categories defined each year as <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BDS and authors' calculations.

Fig. 2: Differential gross job flows

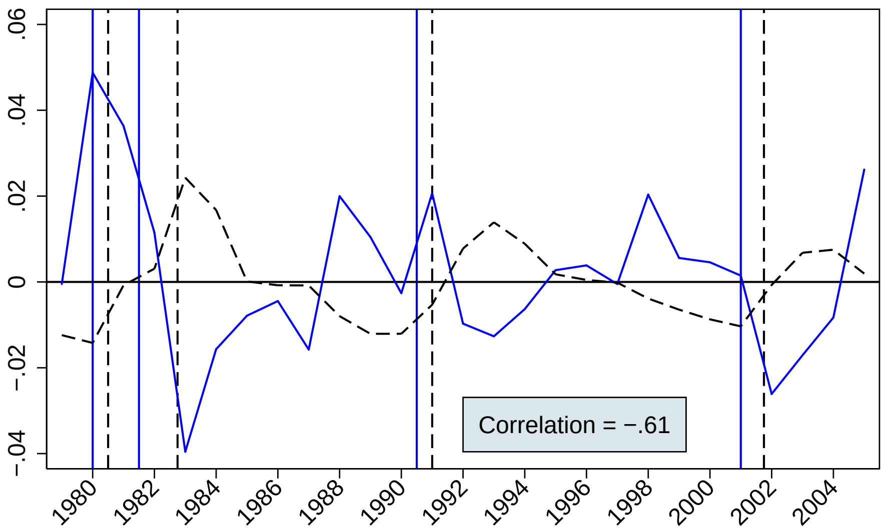

Solid = Differential net job creation, Dash = HP−detrended unemployment rate.
Categories defined each year as <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BDS and authors' calculations.

Fig. 3: Differential growth

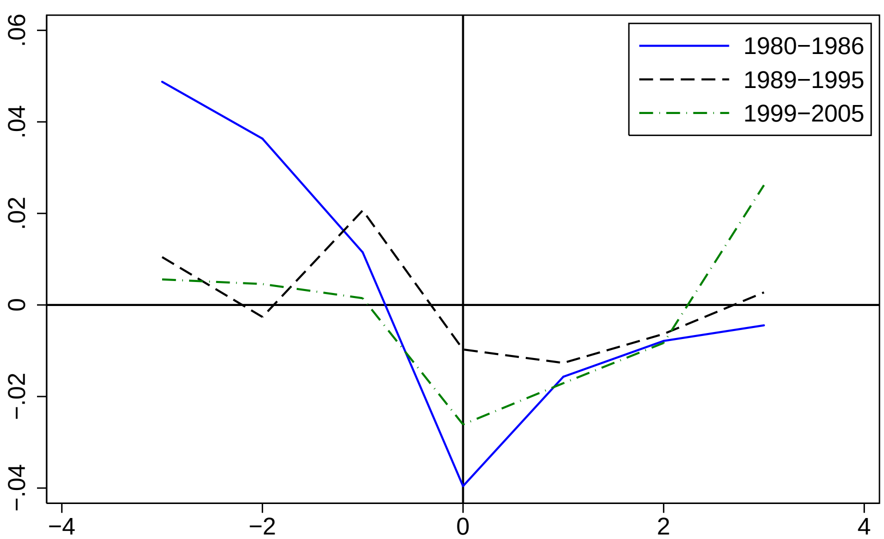

Differential average firm size growth, large minus small, around NBER troughs.
Size categories defined each year as <50 and >1000.
Date zero in 1983, 1992 and 2002.
Source: BDS and authors' calculations.

Fig. 4: Differential growth, around three consecutive NBER troughs

growth from March 1990 to March 1991. It is quite clear from Figure 4 that recessions are times when initially large employers suffer a (proportionally) much larger loss of employment.

Before we move on to a more thorough analysis of the higher cyclical sensitivity of large employers (the basic fact listed as Fact 2 above), we should point out the following two issues about the BDS data. First, while BDS begins in 1977 (that is, with job flows between 1976 and 1977), it shows some troubling discrepancies with comparable, well-proven datasets in the first two years, 1977 and 1978. In particular, net JC in BDS in 1977-1978 is 2% higher in BDS than in the Census Bureau County Business Pattern (CBP) data.[7] Afterwards, the two series (BDS and CBP) more or less coincide. This anomaly is reflected in Figure 5, where we plot the differential employment growth rate (between initially large and small firms) against the unemployment rate, both detrended. The behavior of the growth rate is quite unusual in 1977-1978, suggesting serious data issues of unknown nature. Because CBP is a true and tested dataset while BDS is a new resource that is currently undergoing some revisions, but is immune from reclassification, we begin our analysis in 1979, since when CBP and BDS coincide. Fortunately, 1977 and 1978 are not recession years, and we lose little in terms of cyclical variability in the economy.

Second, all the evidence presented in this subsection was constructed from BDS data based on "firm counts", as defined in Subsection 2.4, meaning that entering *firms* were excluded from the JC count, while entering *establishments* created within *continuing firms* were kept in. An alternate option is to use data based on "establishment counts" (again, see Subsection 2.4), which exclude all new establishments from the JC count. Differences between the two options are generally minor, as exemplified in Figure 6, which replicates Figure 3 using the establishment count.

We now dig deeper and check whether the basic Fact 2 that we uncover, that initially larger firms have more cyclically sensitive employment, holds within or across geographical locations, industries and age of the firm. One important proviso is that in BDS the state and industry refer to the establishment, and the initial size to the parent company. It is of course impossible to attribute a state and industry to most large firms that have establishments in many states and industries. These firms, though, will weigh more because their many and diverse establishments will be all attributed to the large firm size category, wherever they are and whatever they produce. Unfortunately, cross-tabulations by two or three of such criteria, such as net JC by initial firm size, within each industry and each state, is not usable because too many observations are suppressed for confidentiality reasons. We can, however, perform the analysis within each of these categories, one at a time, in addition to initial firm size. We begin with industries and continue with geographical units.

[7]Although CBP is based on the same basic source data as BDS, differences in how the source data are processed lead to differences in the published statistics. For details, see the BDS technical note available at `www.ces.census.gov/docs/bds/BDS_Technical%20Note_102008-1.doc`, in which the discrepancies between BDS and comparable data sets are also discussed.

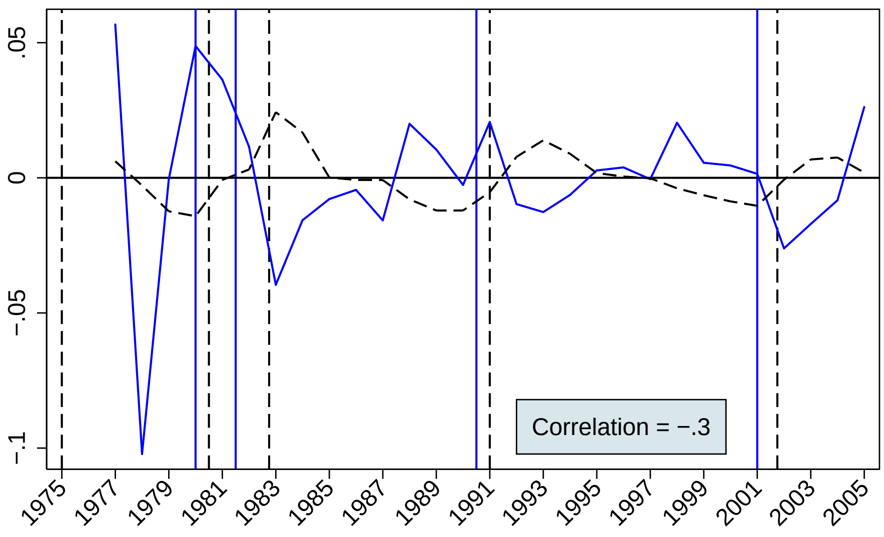

Solid = Differential net job creation, Dash = HP−detrended unemployment rate.
Categories defined each year as <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BDS and authors' calculations.

Fig. 5: Differential growth, since 1977

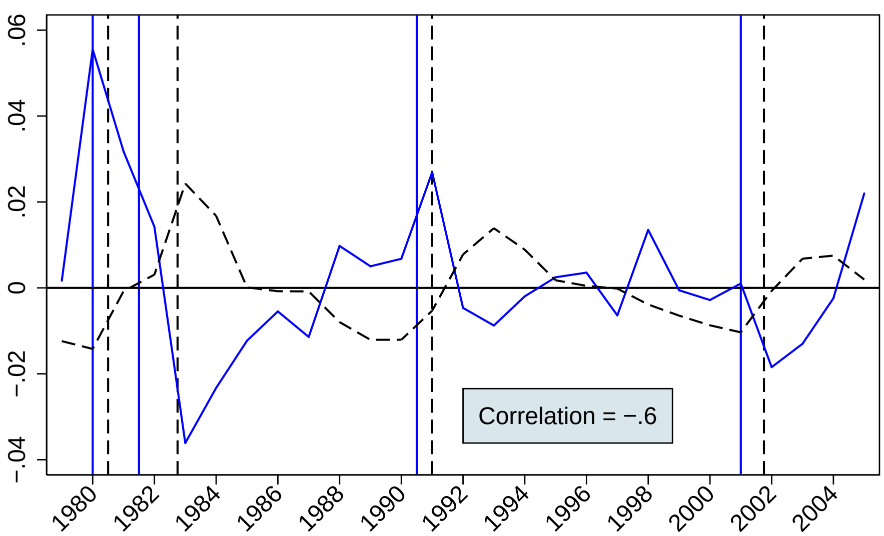

Solid = Differential net job creation, Dash = HP−detrended unemployment rate.
Categories defined each year as <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BDS and authors' calculations.

Fig. 6: Differential growth, based on establishment counts

3.1.3 Industry Patterns

One natural question is whether the more pronounced cyclical sensitivity of large employers just reflects a larger cyclical sensitivity of sectors, like manufacturing, that have above-average firm and establishment size. That is, our main finding may be due to a composition effect. We would like to emphasize that, even if true, this fact would not diminish the novelty and interest of our finding, although it may appear less surprising in light of what we already know about the firm size and cyclical sensitivity of different sectors. But it turns out that the higher cyclical sensitivity of large employers is, by and large, a phenomenon that occurs within, and not between, industries.

Industries			
Construction	−.589	TCPU	−.369
FIRE	−.566	Services	−.36
Retail Trade	−.534	Agric., Forestry, Fishing	−.262
Manufacturing	−.41	Mining	−.229
Wholesale Trade	−.387	**All**	−.61

Table 2: Industry-level correlations between average unemployment over past year and differential firm growth, based on firm counts

We maintain our classification in nine broad sectors, which is reasonably consistent through the changeover from SIC to NAICS in 1998.[8] In Figure 7 we plot again the differential growth rate of (initially) large minus small firms in each broad sector against the civilian unemployment rate, both detrended.[9] The within-industry variability of the growth rate differential is now much larger and swamps visually that of the unemployment rate. Nonetheless, the pattern emerges also within most sectors, including the larger ones. In Table 2, we report the correlations, which are all negative.

[8]One slightly puzzling piece of information on industry coverage in the Economic Census, which is as mentioned the source of the BR, CBP and BDS data every five years, is in the following document: `www.census.gov/prod/ec02/ec02-00r-hist.pdf#page=171`,. “The 1992 Economic Census covered retail trade; wholesale trade; service industries; transportation, communications, and utilities; finance, insurance, and real estate; construction industries, manufactures, and mineral industries; and the TIU.S.. [...] By including coverage of financial, insurance, real estate industries, communications, and utilities, the 1992 census marked the most significant expansion of the census in half a century. [...] Altogether, the 1992 Economic Census covered approximately 95 new industries, expanding coverage to approximately 98 percent of the nations economic activity. [Coverage of the 1987 Economic Census was approximately 75 percent of the nation’s economic activity.]”

[9]The small and large firm cutoffs are set here to <50 and >500 employees for all industries. Similar plots using the first and third quantiles of the average (over the entire period) firm size distribution within each sector as cutoffs — to allow differentiation of sectors that have exceedingly high average establishment size, such as manufacturing, from much of the rest of the economy — are available on request, and draw a very similar picture.

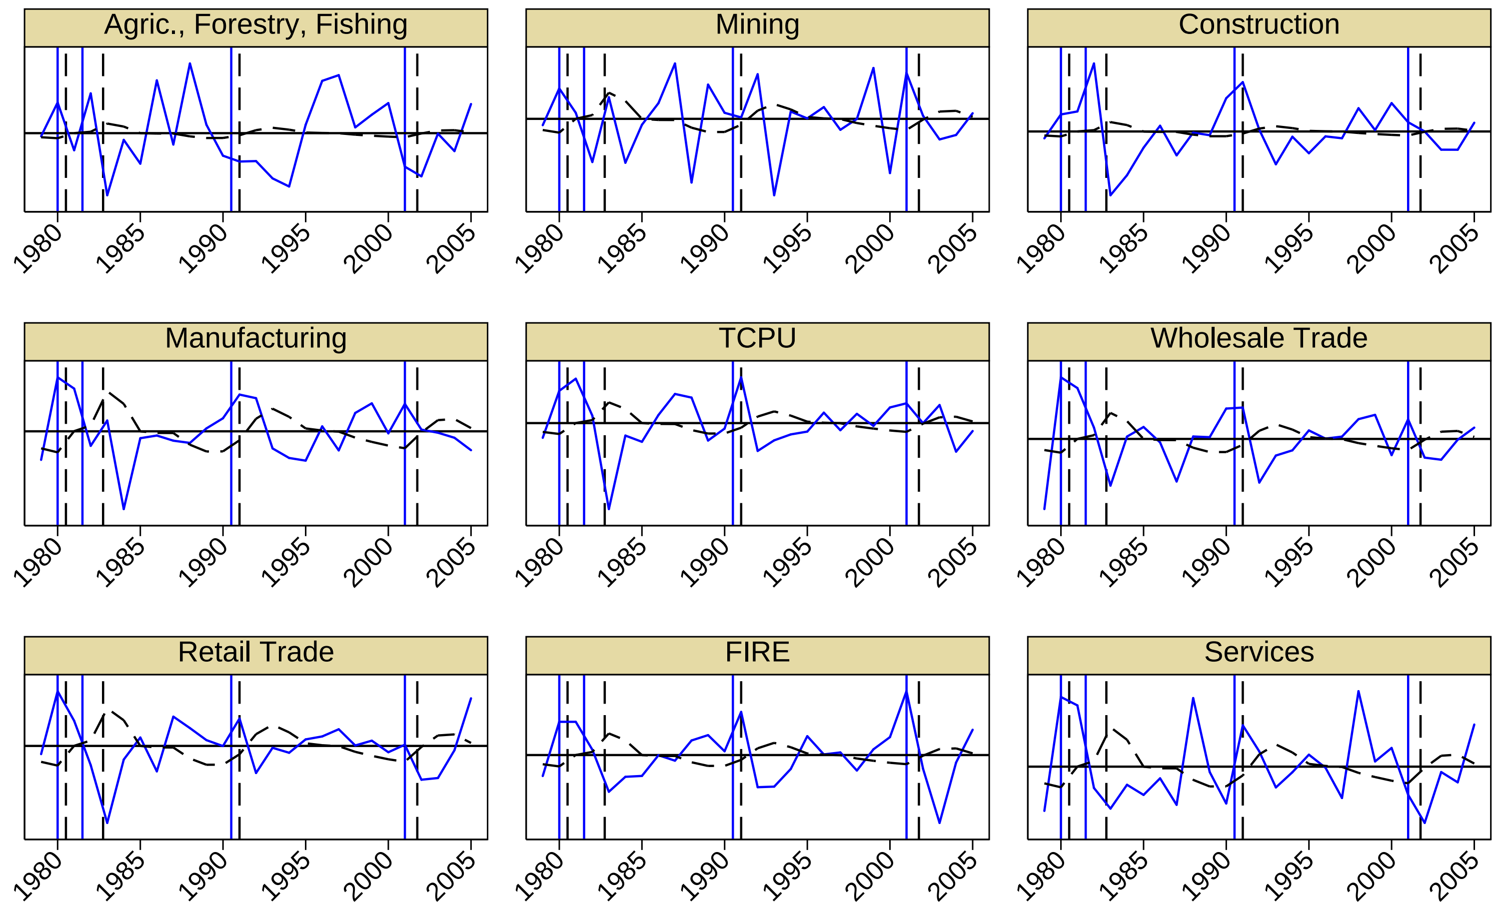

Solid = Differential net job creation, Dash = HP−detrended unemployment rate.
Categories defined each year <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BDS and authors' calculations.

Fig. 7: Differential growth, by broad industry

In Figure 8 we show the results of the reverse exercise. We classify and rank sectors by mean firm size over the period. Then we calculate the employment growth of the largest three sectors and that of the smallest three, and take the difference. This is a between-industry measure of employment reallocation from sectors that have on average larger or smaller employers. The cyclical pattern disappears. The relative growth differential has an essentially zero correlation with the unemployment rate.

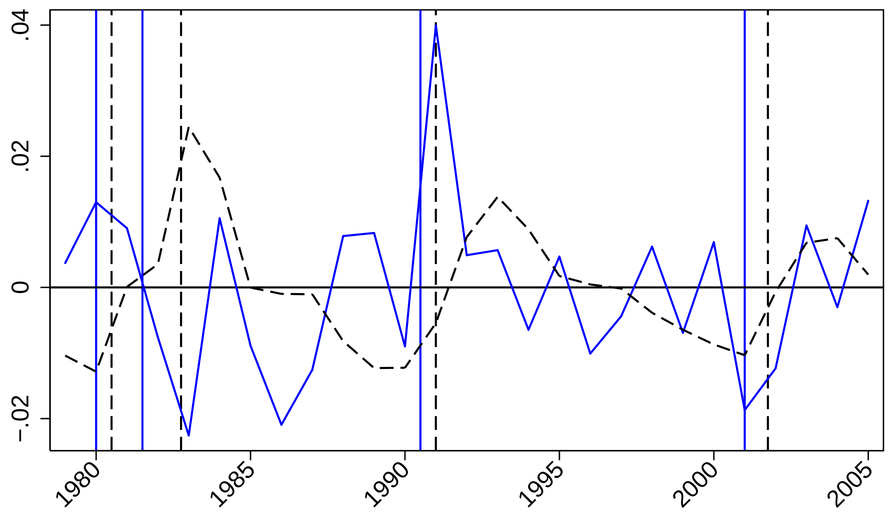

Solid = Between−industry differential net job creation, Dash = HP−detrended unemployment rate.
Categories defined each year based on cumulated employment shares of industries, ranked by mean firm size.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BDS and authors' calculations.

Fig. 8: Differential growth, between industries

3.1.4 Geographical Patterns

We associate to each Census division and U.S. state the corresponding local civilian unemployment rate from the BLS, and correlate the differential firm (large minus small) growth rate and the detrended local unemployment rate. Figure 9 reports the time series of firm size differential growth rates and detrended unemployment rates within each U.S. Census division. The correlations are reported in Table 3. Table 4 reports the same correlations for all U.S. states. Both the figure and the tables show clearly that the phenomenon that we identify takes place within states and regions, and it is not driven by employment moving from locations with small firms to locations with large firms in a boom, and vice versa in recessions.

At this point we can summarize the findings of this and the previous Subsections in the following statement:

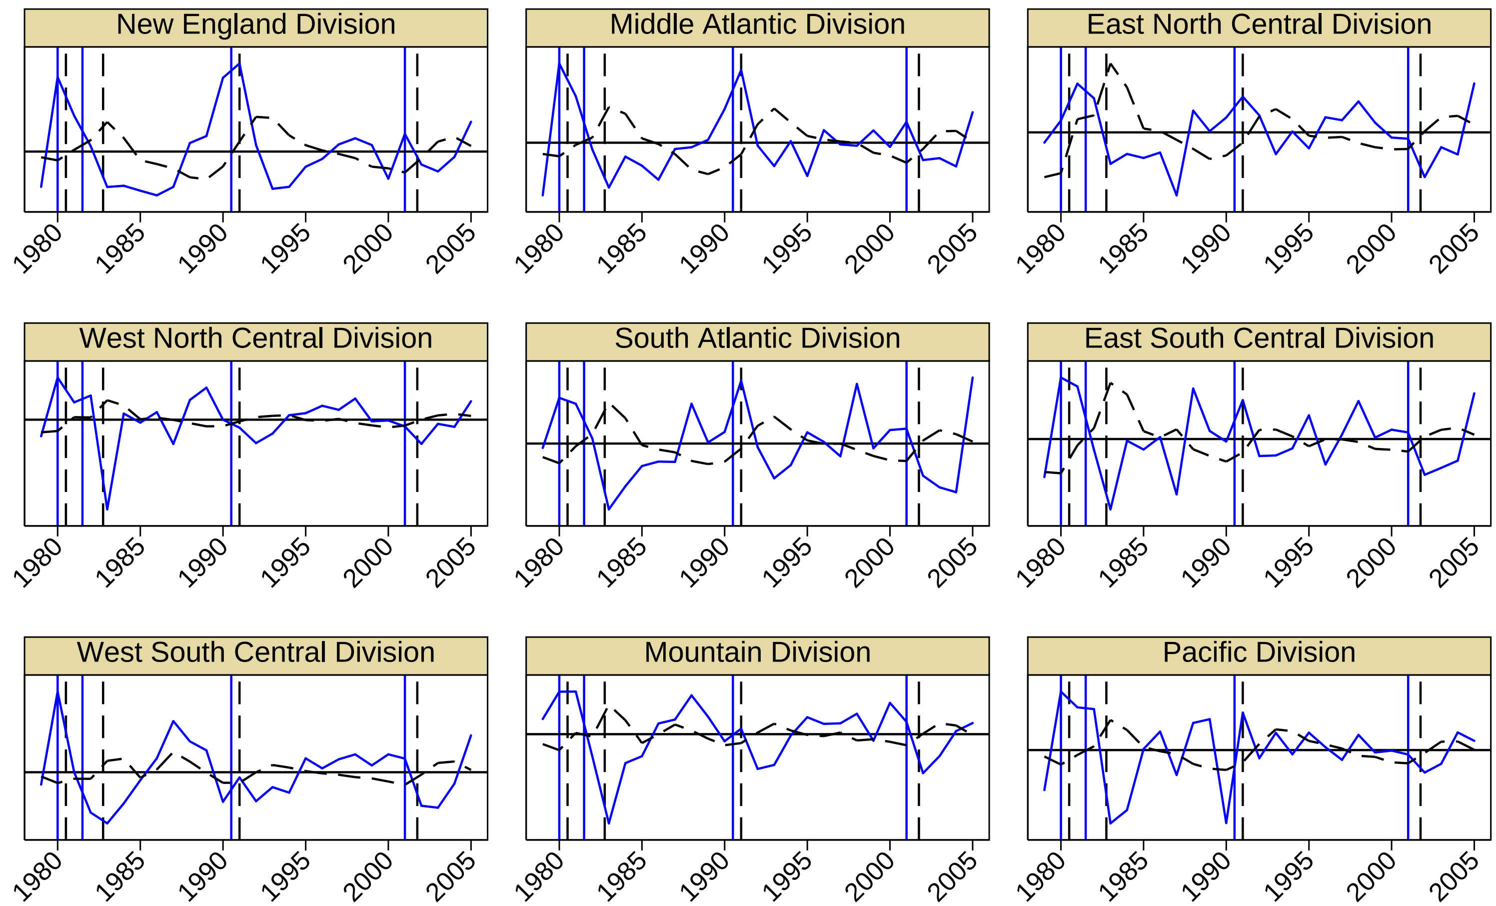

Solid = Differential net job creation, Dash = HP−detrended unemployment rate.
Categories defined each year as <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BDS and authors' calculations.

Fig. 9: Differential growth, by Census division

Census Divisions			
Mountain	−.595	Pacific	−.274
South Atlantic	−.574	West North Central	−.269
East South Central	−.49	New England	−.087
Middle Atlantic	−.354	West South Central	−.043
East North Central	−.298	**All**	−.61

Table 3: Division-level correlations between average unemployment over past year and differential firm growth

North Carolina	−.601	New York	−.276
Kansas	−.566	South Dakota	−.252
Georgia	−.554	Missouri	−.205
Colorado	−.532	Nevada	−.198
Idaho	−.511	Washington	−.182
Virginia	−.478	Oregon	−.174
South Carolina	−.453	Indiana	−.163
Arkansas	−.449	Vermont	−.158
Iowa	−.437	Massachusetts	−.117
Kentucky	−.404	District of Columbia	−.111
Florida	−.397	Hawaii	−.105
Pennsylvania	−.392	New Hampshire	−.084
Alabama	−.388	Maine	−.074
Mississippi	−.372	Nebraska	−.056
New Jersey	−.365	Texas	−.056
Utah	−.356	New Mexico	−.041
Wisconsin	−.353	Montana	−.033
Illinois	−.348	Arizona	−.03
North Dakota	−.339	Michigan	−.003
California	−.332	Ohio	.009
Minnesota	−.32	Louisiana	.028
Connecticut	−.311	Delaware	.037
Maryland	−.31	West Virginia	.056
Tennessee	−.303	Rhode Island	.07
Alaska	−.301	Oklahoma	.095
		Wyoming	.339

Table 4: State-level correlations between average unemployment over past year and differential firm growth

Fact 3. *The higher cyclical sensitivity of large employers holds principally within industries and within States, not between those units.*

3.1.5 Firm Age Rather than Size

Recent work on business micro data (e.g. Foster et al., 2008) as well established theories of firm dynamics (Jovanovic, 1982) point to age as a major predictor of firm behavior. Older firms are much less volatile, productive, and growing, conditional on survival, than younger firms.

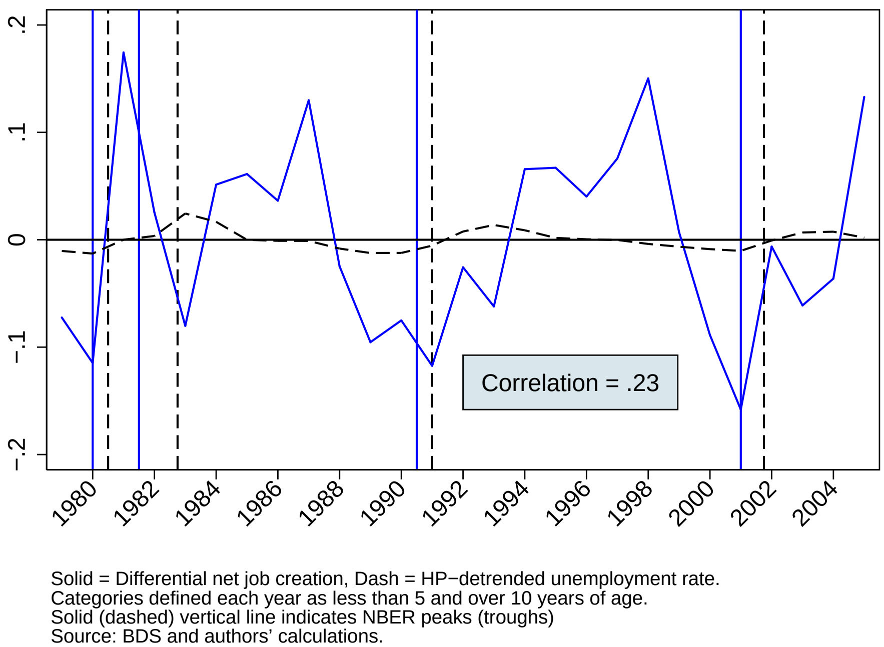

Fig. 10: Differential growth, based on firm age

Figure 10 plots the differential growth rate of initially (the year before) old and young firms (old minus young) against the detrended unemployment rate. Two aspects of the data emerge clearly. First, heterogeneity by age over the cycle is very large, as unemployment fluctuations are now barely visible. Second, older firms are hit hardest by recessions. Interestingly, the relative growth rate by initial firm age leads the business cycle.

3.2 Business Employment Dynamics (BED)

A different source of information on the distribution of employment by employer size is available from the Business Employment Dynamics program at the Bureau of Labor Statistics. This program collects information accruing from the States' unemployment insurance programs. The dataset only begins in 1992, but its frequency is quarterly. As the name suggests, the BED is primarily a dataset of job flows. Although the classification is dynamic, thus still subject to the Reclassification Bias (see Section 2), the small bias found in BDS (documented below in Section 4) suggests that valuable information is contained also in BED's flows by reclassified size. The BED has the dual advantage over the BDS of a quarterly frequency and also of including the first quarter of 2008, thus a new cyclical trough in 2007:Q4.

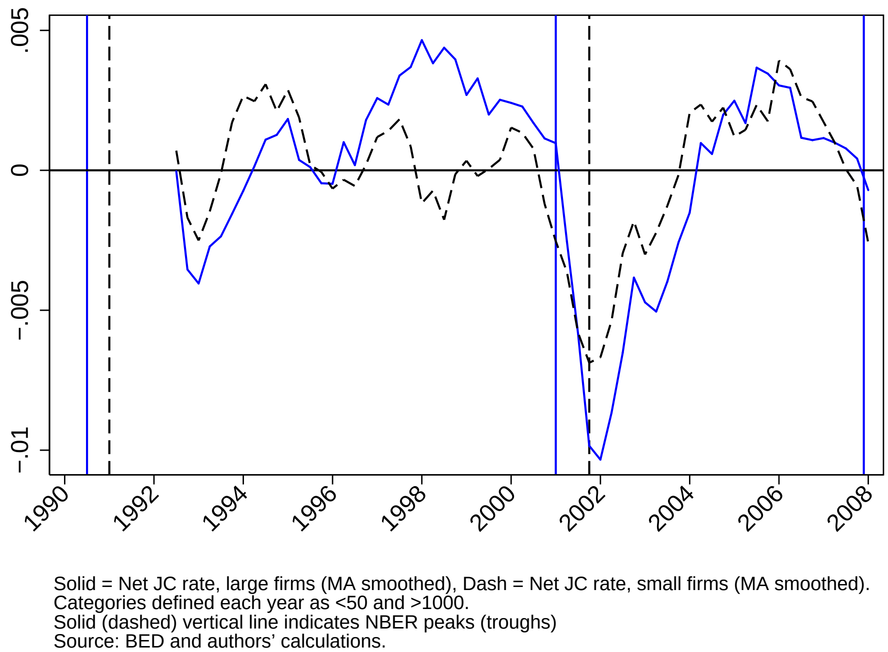

Fig. 11: Net job creation rates

Figure 11 plots the net JC rate of large and small firms, according to the usual size cutoffs, and NBER dates.[10] The higher cyclical volatility of large firms jumps out. In Figure 12 we take the difference in these growth rates and plot it against the detrended unemployment rate. The two series mirror each other, correlation -0.76. Figure 13 decomposes the net flows of Figure 11 into differential gross JC and differential JD, large minus small firms. Clearly, excess JC by large employers accounts for their faster net growth in the late 1990s, while a spike in JD by large employers explains their slump in the 2001 recession. The message from these BLS data is the same that we obtained in the Census BDS data.

3.3 Compustat

Longitudinal links in BDS are valid only for one year, on a rolling basis. Proper longitudinal business data allow to fix a firm's identity once and for all at a common date and follow its

[10]Recall that the denominator in the published BED job flow rates is the average $L_{it-1}^{(1/2)}$ in the notation of Section 2, i.e. average employment in the size class between quarters $t-1$ and t. This is slightly different from initial employment in the size class, $L_{it-1}^{(0)}$, which is used in BDS.

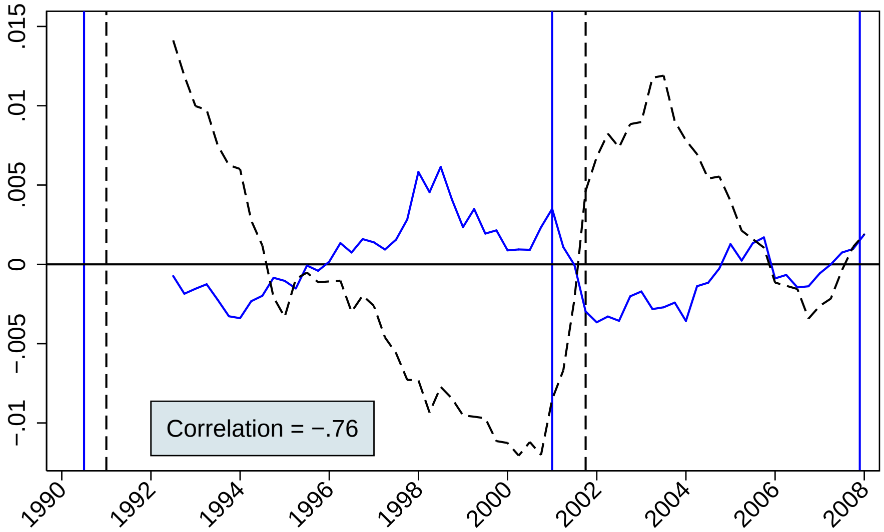

Solid = Differential net job creation (MA−smoothed), Dash = HP−detrended unemployment rate.
Categories defined each year as <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BED and authors' calculations.

Fig. 12: Differential growth

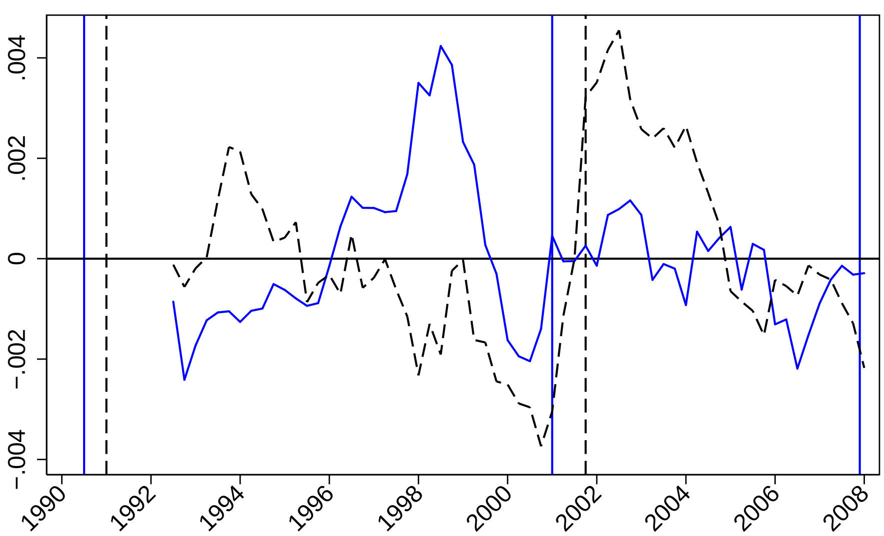

Solid = Differential gross JC (MA smoothed), Dash = Differential gross JD (MA smoothed).
Categories defined each year as <50 and >1000.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: BED and authors' calculations.

Fig. 13: Differential gross job flows

growth over a long period of time including several business cycles, as a function of its initial size. Since all firms in existence start from size 0, the initial date cannot be too far back in time, but certainly more than the one year of BDS can be helpful.

The only fully longitudinal U.S. business dataset that we have been able to access is Compustat (see below for other countries.) This comprises only public companies, so it is not a representative sample, like BDS. Yet, it contains useful information. Figure 14 is reported from MPV08. We fix firm identities in 1975 and classify them once and for all in size bins (by employment), large above 5,000 employees, and small otherwise. The reason for the large size cutoff is that public companies are very large. Then, for each year from 1976 to 2005 (these are starting years, so the same 1977-2006 period covered by the BDS where these are ending years), we calculate the growth rate of employment over the past year at (initially, in 1975) large firms and subtract the growth rate of the other, initially small firms in the sample. We plot this difference in growth rates, in a way that mimics Figure 3. Consistent with the patterns uncovered in BDS and BED, over four consecutive business cycles this difference in growth rates is procyclical, and crosses zero when the labor market turns tight, except for an outlier in 1987.

This analysis highlights one difficult issue that unavoidably arises with fully longitudinal datasets. As explained above, the advantage is that we can fix the identities of a given group of firms and then track them over a long time period. This allows us to prevent reclassification and that new entrants change the composition of the sample. But the effect of exit cannot be eliminated. Figure 14 only includes public companies that survived through 1976-2005, so there is a strong survivorship bias. If we tracked all companies that were listed in 1976, we would then be losing some of them along the way to delisting and bankruptcy, so we would allow only for job destruction from exit and not for job creation from entry. This fact works in favor of using short longitudinal links as in BDS.[11]

4 The Distribution of Employment by Employer Size over the Business Cycle

We continue our data analysis with the cyclical behavior of the distribution of employment across size classes of their employers, firms or establishments. As discussed in Subsection 2.3, the growth rate of the employment share of a given size class only approximates the growth rate of employment in the set of firms *initially* in that size class, which is our main object of interest, up to the Reclassification Bias. We find it useful, however, to report evidence

[11]In the NBER Manufacturing Panel from Compustat, fixing firm sizes in 1975, the difference in growth rates of employment between firms in the top quartile and firms below the median has the now familiar pattern throughout 1975-1995: falling in recessions and gradually climbing in expansions. This is true whether we define size in terms of initial employment, sales, capital or assets. We thank David Berger for these calculations. Results are available upon request.

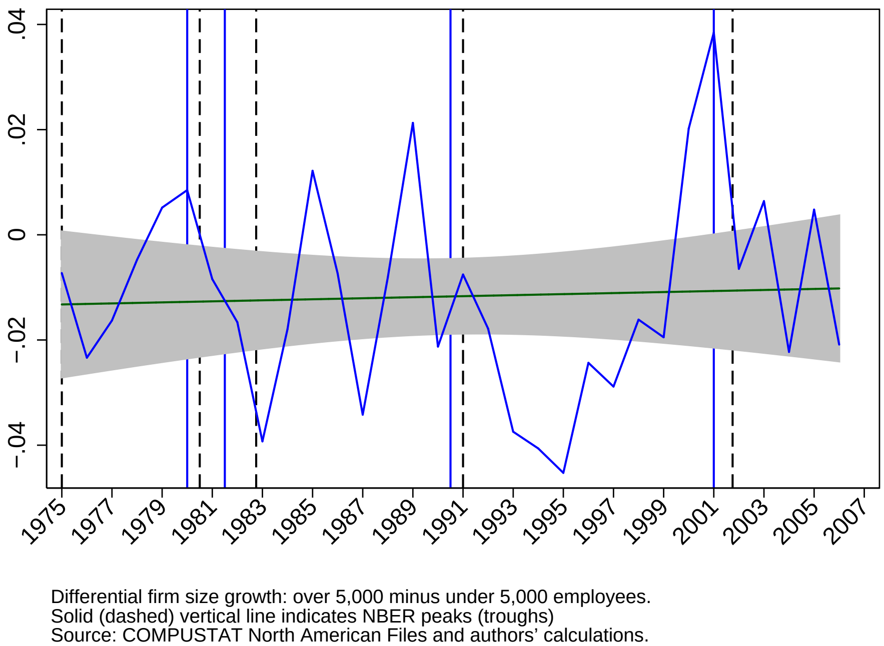

Fig. 14: Differential growth, fixed sample of publicly traded U.S. companies classified by size in 1975.

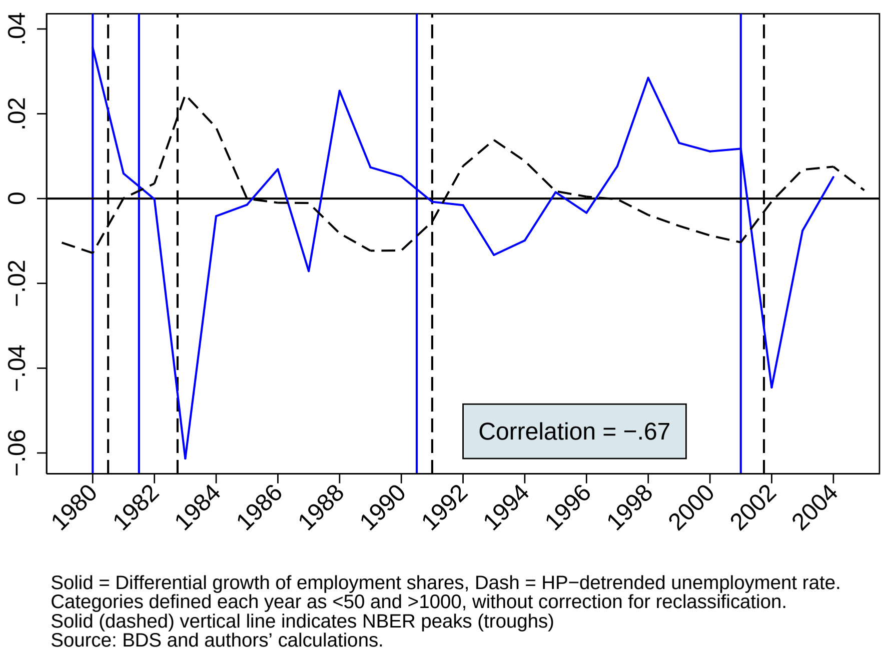

Fig. 15: Differential growth of employment shares, BDS data

on changes in the employment distribution from repeated cross-sections, both to gauge the magnitude of the Reclassification Bias and because the data required to construct employment shares are much more widely available across countries than longitudinal data such as BDS that make it possible to correct for reclassification. We provide some international evidence towards the end of the paper.

4.1 Using BDS Data to Quantitatively Assess the Reclassification Bias

The BDS dataset reports job flows and employment growth in each year t by initial (year-$t-1$) size. It is therefore immune to the Reclassification Bias as it conveys longitudinal information over two year-periods $[t-1, t]$ about size classes, the composition of which is *fixed* over the entire period. However the same BDS dataset also allows construction of employment shares of size classes each year. We can thus repeat the exercise of subsection 3.1.2 only using differences in the growth rates of employment shares across size classes. Although this methodology is incorrect (because subject to the Reclassification Bias — see Subsection 2.3), a comparison with the correct one illustrated in Subsection 3.1.2 can help us gauge the quantitative relevance of the Reclassification Bias, as formally measured by the third term in equation (1).

Figure 15 parallels the aggregate picture of the cyclical behavior of differential firm growth across size classes drawn earlier in Figure 3, and shows that the message is essentially unchanged when one neglects the reclassification issue. Although the correlation is slightly more negative in Figure 15 than in Figure 3, as we would expect because of reclassification which pushes more firms in, hence imputes more growth to, the large firm size categories when the economy grows and unemployment is low in booms, and vice versa in recessions, the bias appears to be quantitatively very modest.

> **Fact 4.** *Reclassification plays a quantitatively negligible part in explaining the higher cyclical sensitivity of large employers.*

4.2 BED

While the BED data contains only job flows, in levels and rate (as a fraction of average size over the quarter, $L_{it-1}^{(1/2)}$), also by firm size, we have obtained from the BLS an unpublished tabulation of the time series of employment, establishment and firm counts by initial size of the firm, for the whole U.S. economy (no breakdown by location or industry), in 1992:III-2008:I. As for BDS we detrend employment share growth rates and correlate the deviations from trend with those of the unemployment rate, at the same quarterly frequency. The results, depicted on Figure 16, corroborate our previous findings from BDS: the negative correlation is again very strong, at -0.77. Although that correlation is likely to overstate the true difference in cyclical sensitivity of large and small employers, due to the Reclassification Bias, our previous BDS findings suggest that the bias is likely to be small.

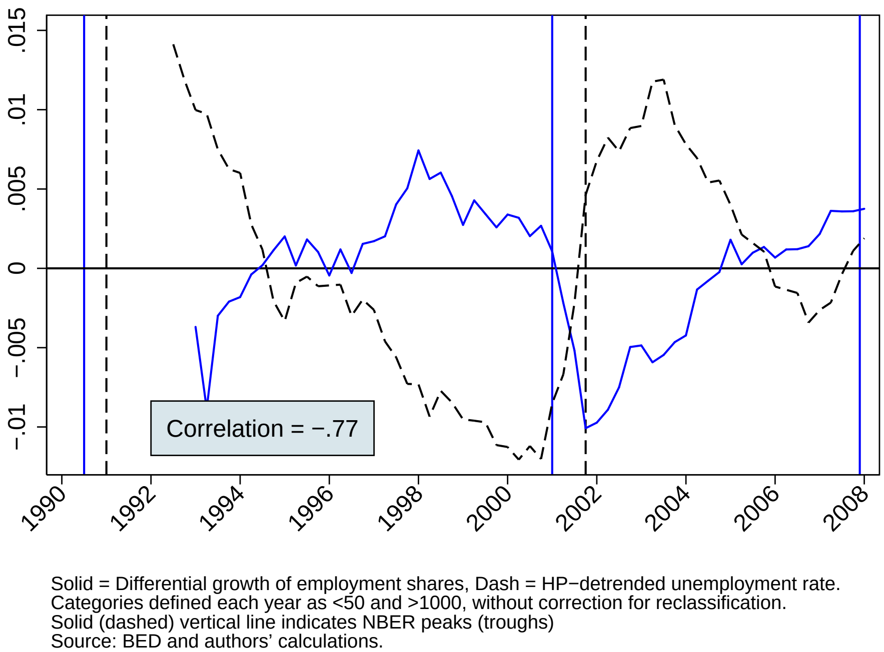

Fig. 16: Differential growth of employment shares, BED data

5 Worker Flows by Employer Size over the Business Cycle

In order to shed more light on the underlying sources of the cyclical movement in employment shares and growth rates by firm size, we now turn to worker flows. A group of firms of similar initial size can add employment by either adding more jobs (higher JC) or shedding fewer (lower JD). Both are changes on the extensive margin in the number of firms in the group that gain and lose employment. In turn, the intensity of each individual firm JC or JD can vary. Furthermore, the status of a firm as a net creator or destroyer of jobs can change in different ways. In particular, the firm can hire more or fewer workers, for given attrition, or separations might increase, for given inflow. Finally, among separations, some originate from outside opportunities that the firm cannot counter, such as outside offers to its workers and voluntary quits to non-employment. Other separations such as layoffs, are originated by the firm due to poor business conditions. In this section, we aim to shed some light on which of these mechanisms are at play behind the cyclical behavior of the growth/size relationship that we uncovered. The need for information on worker flows by employer size at high frequency restricts the available time span in existing datasets. However we find that, in the tight labor market of the late 1990s, large employers grew faster in part because they poached more employees than their smaller competitors, hence had an additional source of hires (beyond unemployment) they could draw from. Conversely, in the 2001 recession layoffs at large employers, as measured by mass layoffs, spiked much more than layoffs at large in the economy.

5.1 Sources of New Hires over the Business Cycle

We report from MPV08 evidence on where do new workers join employers of different sizes at different stages of a business cycles. In that paper, we argue that the reason for the faster growth of small firms in loose labor market and of large firms in a tight one is related to the ease with which new hires can be found among the unemployed. When these abound, all employers hire and grow at a rate that depends on their sampling weight in job search. When unemployment grows scarce, large employers, which are typically more productive (in revenue terms) and higher paying, can more easily poach employees from smaller competitors. So they can keep growing their employment through that channel, and smaller employers are out of luck and their growth is curbed, in relative terms.

We corroborate this hypothesis with evidence from publicly available data from the Census' Survey of Income and Program Participation (SIPP).[12] SIPP contains information about the workforce size of an individual's current employer (employing establishment), as well as individual job histories, with 4-year long worker longitudinal links and weekly information on employment status. This allows a crude analysis of the poaching activity of establishments as a function of their size. Figure 17 plots a measure of the fraction of new hires coming from another employer (i.e. following an employer-to-employer transition) for three categories of hiring establishment size.[13] In other words, it plots a measure of the importance of poaching in the recruitment activity of establishments, by size of the hiring establishment. Changes in the design of the SIPP and other data limitations restrict the period over which that indicator can be constructed to the years shown on Figure 17. While this admittedly constitutes very limited evidence, we still notice the following two points. First, poaching was more intense in the latter half of the 1990s expansion than in the immediate aftermath of the 2001 recession. This is true for all three categories of establishment size. Second, larger establishments almost always poach more than smaller ones. This difference in "poaching

[12] Information about the SIPP, as well as data files are available at `www.bls.census.gov/sipp/`.

[13] Specifically, it is constructed as the fraction of workers who have changed employers in the previous year and are now employed at an establishment in size category X *without work interruption* among all workers having changed employers in the previous year and now employed at an establishment in size category X.

intensity", however, is more pronounced in 1997-1999, when the labor market turns tight, than in 2002-2004, when it is slack.

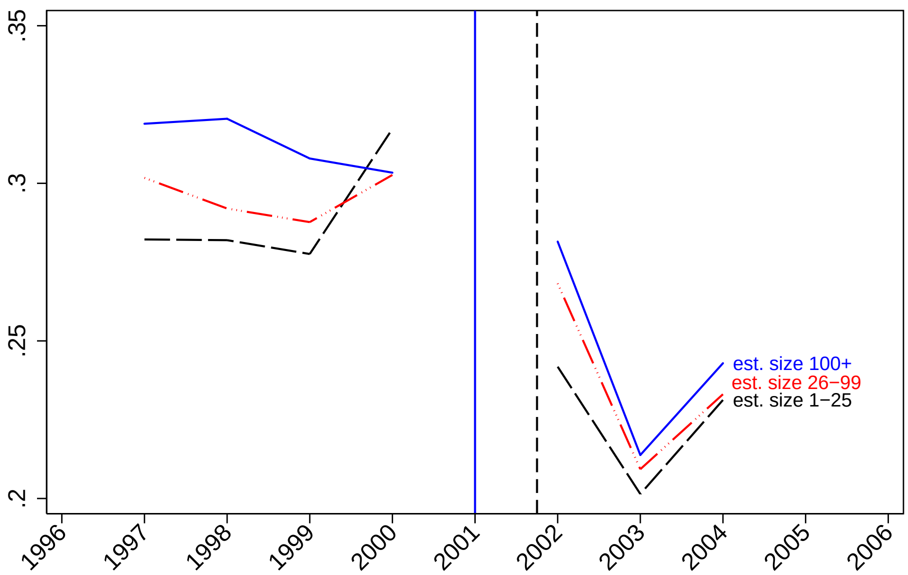

Fraction of new hires coming from other employers, by establishment size category.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: SIPP and authors' calculations.

Fig. 17: Poaching and establishment size

5.2 Layoffs

We do not have access to micro business data on gross separations by size of the employer. One important component of separations is layoffs. The BLS provides two sources of information that we can use to assess whether large employers lay off (proportionally) more workers in recessions.

The Mass Layoff Statistics (MLS) available starting in 1995, record the number of episodes and the total number of workers involved each month in layoffs involving more than 50 workers simultaneously. This is broken down by state and 3-digit industry.[14]

The Job Openings and Labor Turnover Survey (JOLTS), starting in December 2000, records (among others) layoffs for a sample of firms, also broken down by region and broad industry groups.

[14] "A mass layoff occurs when at least 50 initial UI claims are filed against an establishment during a consecutive 5-week period. An extended mass layoff occurs when at least 50 initial claims are filed against an establishment during a consecutive 5-week period and at least 50 workers have been separated from jobs for more than 30 days."

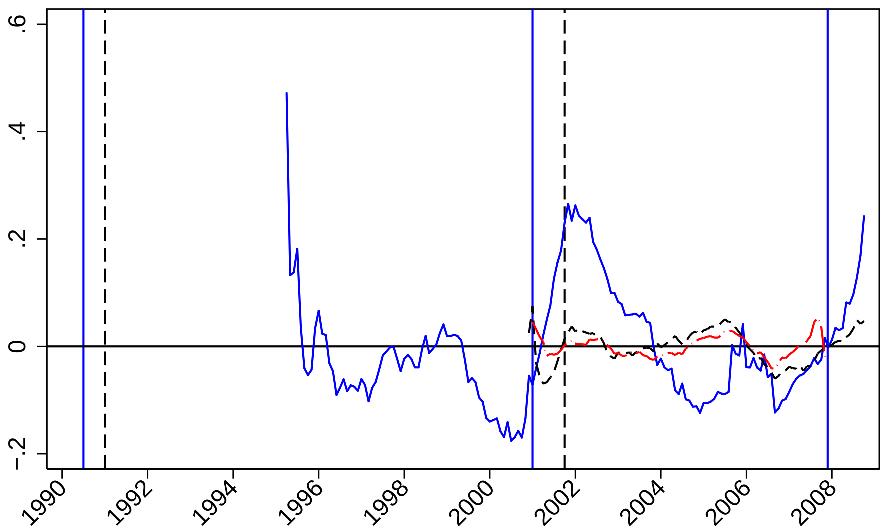

Solid = Mass layoffs, Dash = Total layoffs, Dash−dot = Total layoffs, adjusted.
All series seasonally adjusted, in percentage deviation from trend.
Solid (dashed) vertical line indicates NBER peaks (troughs)
Source: JOLTS, MLS, Davis et al. (2008), and authors' calculations.

Fig. 18: Mass layoffs vs. all layoffs

By comparing layoffs at large establishments (MLS), which can only belong to medium-to-large firms, with layoffs at all employers (JOLTS), we can check whether in the 2001 recession and again in 2008 MLS layoffs rise faster than layoffs in general. Figure 18 leaves no doubt that mass layoffs are wildly more volatile and concentrated in recessions than layoffs at large.[15] Large employers shrink more in recessions in part because they lay off a larger proportion of their payroll, not (only) because they hire fewer or lose more workers to quits.

6 International Evidence

Business micro or semi-aggregated data (by employer size) on employment are available in many countries. We now present evidence of the following, final stylized fact:

> **Fact 5.** *The higher cyclical sensitivity of large employers is a phenomenon that is not unique to the U.S.. It is common to various countries, of different sizes and stages of development.*

[15] JOLTS has been shown to have problems in that it severely undersamples small and young establishments, thus generating too small turnover rates (Davis et al., 2008). This makes the JOLTS sample biased towards large employers. For us it is not a problem, because if JOLTS and MLS show a difference than a fortiori the real universe of layoffs must be even more different than MLS. At any rate, Figure 18 plots both the original JOLTS data and the amended data from Davis et al. (2008b).

A large empirical literature has tested in several countries the well-known hypothesis that the distribution of firm size can be well approximated by a Pareto distribution, that is, the log rank of an employer in the overall size distribution is linearly related to the log of its size. The employment size distribution counts the shares of employers (firms or establishments) in each size class, rather than the shares of employment. But there is a strict relationship between the two, as employment in a size class is the average size in that class times the share of employers. While most of this literature focuses on cross-sectional patterns, either at one point in time or on average over some time period, a few articles do present and comment on time series of the cross-sectional evidence. In particular, a power law is fitted to cross-sectional data year by year, which allows to study the cyclical behavior of the intercept and slope of the power law. We found that the growth rate of the employment shares of large size classes is procyclical relative to small classes, so we should expect the size distribution of employers to be more heavily weighted on large ones, hence to be more unequal, when unemployment is low, and vice versa. This implies that the Pareto slope should be countercyclical, yielding a flatter power law and a more unequal size distribution in booms, and a steeper law in recessions. This is exactly what the literature has found (e.g., Marsili (2006) in the Dutch manufacturing sector, Delli Gatti et al. (2006) in a large sample of Italian firms), although it has failed to draw its further implications, that is, to make the connection and take the further step to the cyclicality of the size/growth relationship.

6.1 Employer Growth by Initial Size over the Business Cycle

For two countries (viz. Denmark and Brazil), we have been able to access full longitudinal business microdata on employment from censuses for an extended period of time allowing us to replicate the exercises that we performed for the U.S. with BDS and Compustat. In both cases we compute the differential growth rate between initially large and small firms, where "initially" refers either to the year before or to a given year, fixed once and for all at the beginning of the sample. In both cases, the business micro data that we exploit have been matched to information on the employees, an important dimension that we will exploit in future research.

6.1.1 Denmark

The Danish register-based matched employer-employee dataset IDA (*Integreret Database for Arbejdsmarkedsforskning* — Integrated Database for Labor Market Research) contains basic socio-economic information collected annually in the last week of November on workers, and some background information on employers (including employer identifiers). It covers

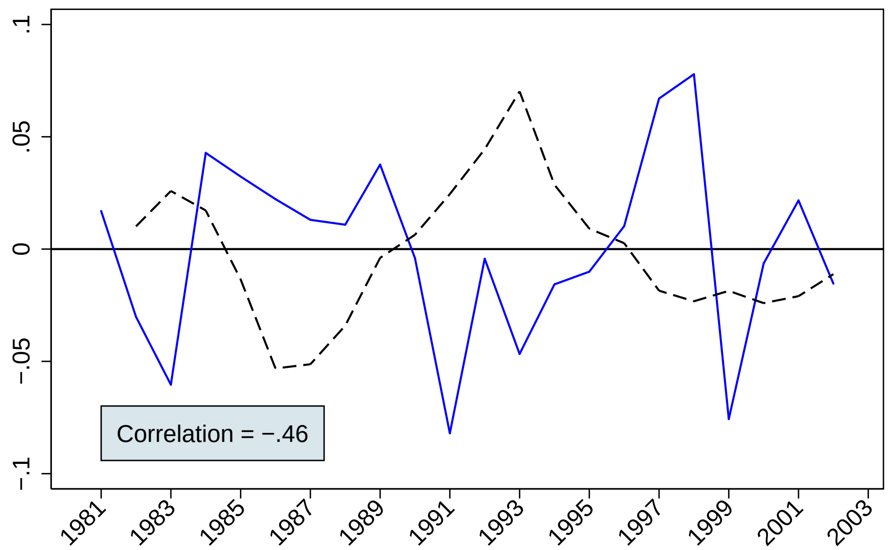

Solid = Differential net job creation, Dash = HP−detrended unemployment rate.
Categories defined each year as <20 and >500.
Source: IDA, OECD and authors' calculations.

Fig. 19: Differential growth, Denmark

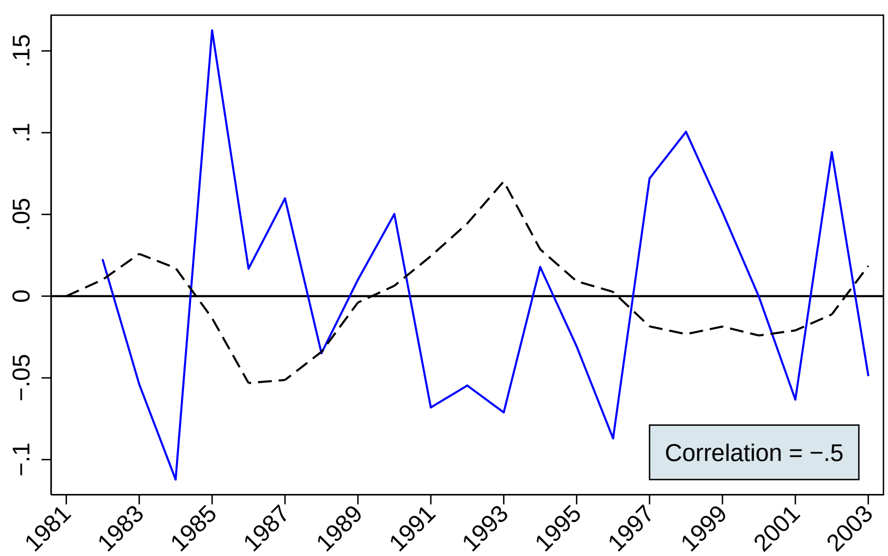

Solid = Differential net job creation, Dash = HP−detrended unemployment rate.
Categories defined each year as <20 and >500.
Source: IDA, OECD and authors' calculations.

Fig. 20: Differential growth, Danish private companies classified by size in 1981.

the entire Danish population aged 16 to 69.[16] As a part of the IDA programme, Statistics Denmark maintains an employer-level panel which contains all the basic information on employers (essentially defined by a tax identifier), including workforce size in the last week of November. The current panel length is 22 years, from 1980 to 2002, and the sample that we use excludes public-sector and not-for-profit employers.

Being a panel of employers, the IDA firm file enables us to assign any particular employer to a fixed size class for as many years as we like. It is possible, in particular, to replicate the structure of the U.S. BDS data, where firms are assigned to a fixed size class over rolling two-year windows. This is what we do in Figure 19, which plots for 1981-2003 the (detrended) OECD unemployment rate and growth rate of large minus small firms, where size is fixed a year in advance, as in the BDS. The familiar pattern emerges: initially large employers grow faster when unemployment is unusually low, and vice versa. The correlation between the two series is $-.46$.

As an alternative, the Danish employer panel also enables us to fix the composition of size classes once and for all, as we did for the U.S. with the Compustat sample (see Subsection 3.3). In Figure 20 we allocate firms to the two size classes in 1983 and track the relative growth rate of these two groups over the following 20 years, without ever reclassifying them. While the two series are overall negatively correlated again (-0.5), the pattern is clearest in 1983-1993, and then appears more blurred in the second half, 1993-2003. This is not too surprising, and we do not take it as contradictory evidence, as firm sizes are not perennial and significant reshuffling occurs over decades.

Our underlying theoretical framework considers size mostly as an indirect measure of productivity, based on the amply documented positive relationship between employment size and revenue-based measures of productivity. Another robust relationship is between average wages paid by a firm and size. For this Danish dataset, both relationships have been confirmed by Lentz and Mortensen (2008). Since IDA contains wage information, we also consider initial wage, rather than initial size, as a measure of undelying productivity. It is plausible that wages reflect productivity better than size, especially for young firms that are still in their initial growth phase. In Figure 21 we allocate firms once and for all to low- and high-paying bins according to mean wage earned by their employees in 1981. We then compute, detrend and plot with the detrended unemployment rate the growth rate of employment at initially high- and low-paying firms, where the former are on average larger. Now the pattern holds strikingly well through the 1980s and 1990s, until 2003. Indeed, except for the very first two years, the two series mirror each other almost perfectly (correlation -0.76).

[16]See Bagger et al. (2009) for a detailed description of the IDA data set.

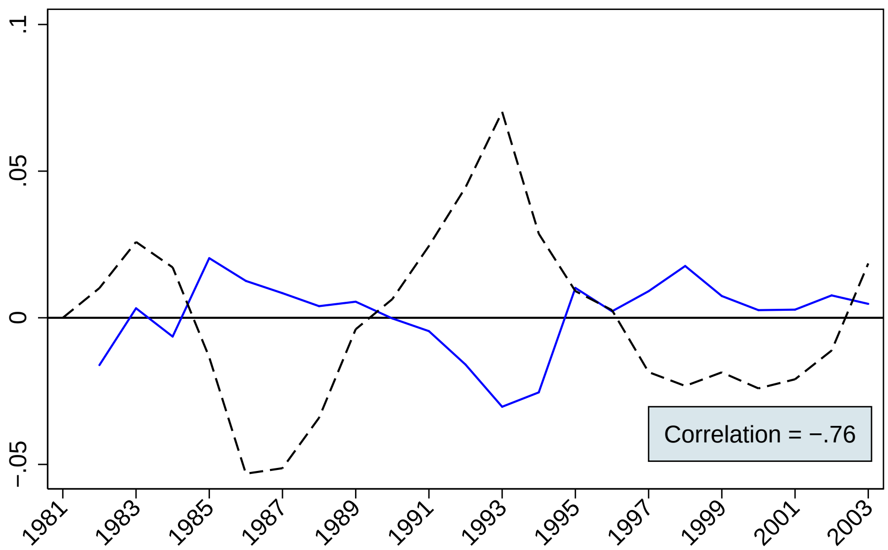

Solid = Differential net job creation, Dash = HP−detrended unemployment rate.
Categories based on mean wage in 1981 and defined as <DKK 58,000/year and >DKK 83,000/year.
Source: IDA, OECD and authors' calculations.

Fig. 21: Differential growth, Danish private companies classified by mean wage in 1981.

6.1.2 Brazil

We exploit longitudinal information on Brazilian employers in the formal sector from the labor market census RAIS (Relação Annual de Informações Sociais). This is an administrative dataset collected annually by the Brazilian labor ministry, which includes all firms in the Brazilian formal sector and provides information for all their workers. The ministry collects demographic information for workers, such as age, education and sex, some information about establishments, such as sector and location, and provides information about the job, such as the average wage earned during that year, the wage in December, the average number of hours worked, occupation, dates of admission and separation, type of contract, causes for separation.[17]

[17] Access to RAIS is restricted to authorized researchers. Carlos Corseuil of IPEA performed the data analysis for us. We take the occasion to thank him for his time and expertise.

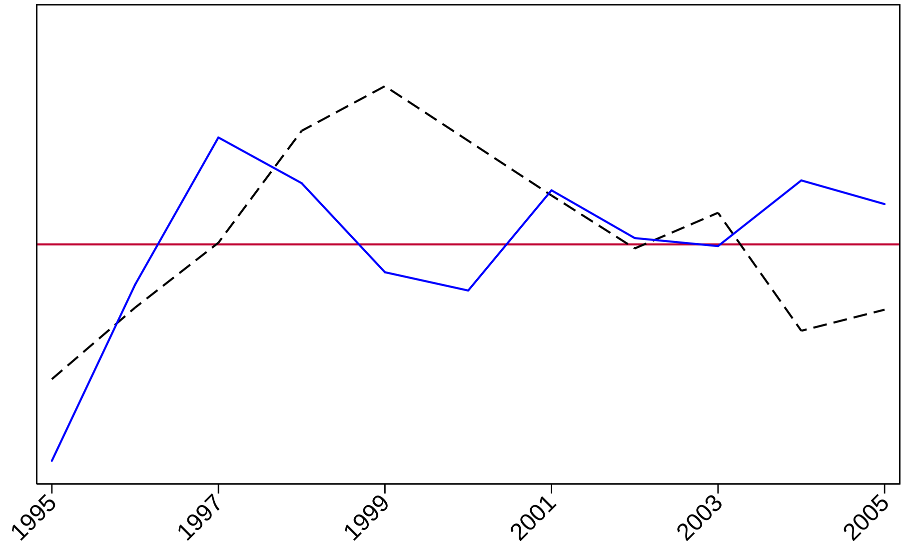

Fig. 22: Differential growth by firm size fixed in 1995, RAIS census of Brazilian employers.

We allocate firms to the familiar size bins in 1995 and track the growth rate of employment in the two groups over the subsequent 10 years. As done before for the U.S. and Denmark, in Figure 22 we plot the detrended differential growth rate of employment at initially (in 1995) large minus small firms and the detrended unemployment rate. The correlation is clearly negative after 1997. Most vividly, the harsh blow suffered by the Brazilian economy in 1998 as a consequence of the Asian crisis hits the larger firms hardest: it generates a sharp rise in the unemployment rate and a corresponding drop in the relative employment growth of firms that started out larger. As the shock is re-absorbed and the unemployment rate declines to more normal levels, the differential growth of employment in large firms nicely recovers. The initial 1995-1996 increase in the unemployment rate, which works against Fact 5, reflects in part structural, rather than cyclical, factors. Estimates of the proportion of Brazilian workers in the informal sector are declining before the Asian crisis, while unemployment rises. After the crisis, this decline in the informal sector grinds to a halt.

6.2 The Distribution of Employment by Employer Size over the Business Cycle

The type of data of easiest access is the distribution of employment among employers of different sizes. Although this is only indirect evidence, as shown in previous sections the cyclical behavior of employment shares by size classes in the U.S. is informative of the

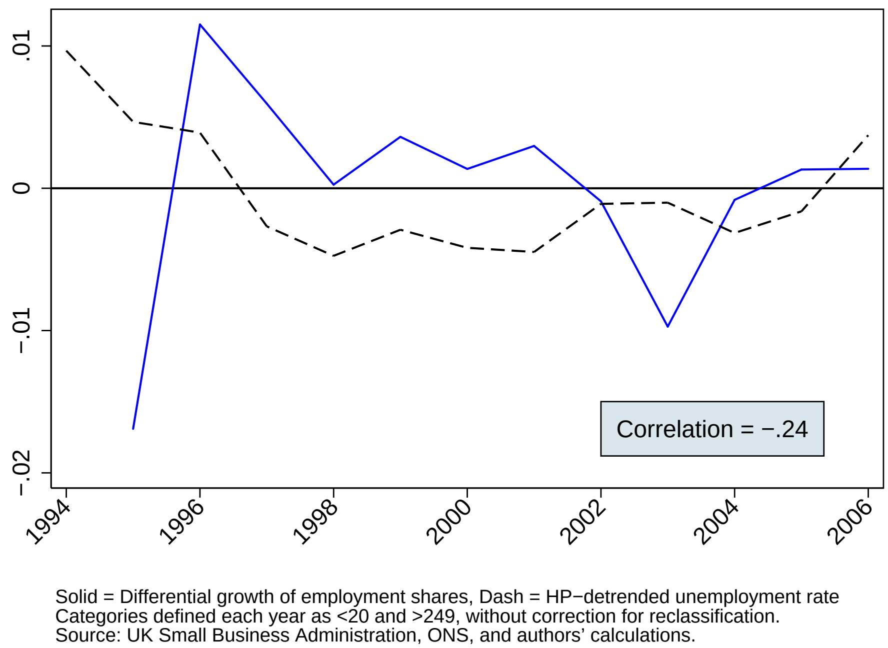

Fig. 23: Differential growth of employment shares, United Kingdom

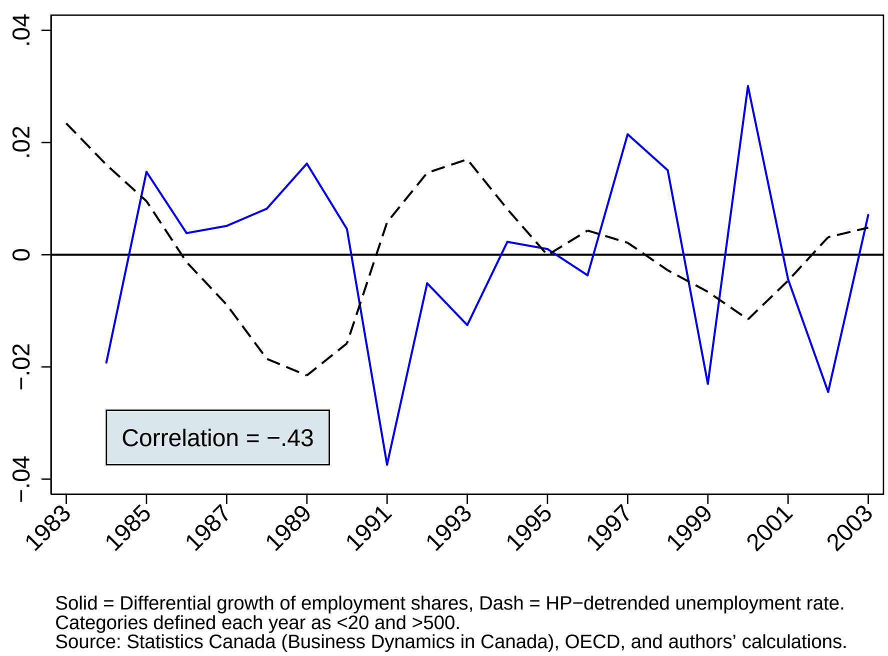

Fig. 24: Differential growth of employment shares, Canada

underlying pattern of growth by initial size. We present evidence from other countries that the growth rate of the employment share of large employers is indeed procyclical. This exercise can probably be replicated in many other countries using publicly available data.

6.2.1 United Kingdom

The UK Small Business Administration publishes a table of employment shares by classes of firm size at annual frequency. We found on the web data for 1994-2006. The size cutoffs are <20 and >249 employees. We compute and detrend employment share growth rates and plot the cross-size-class difference thereof against the detrended UK unemployment rate (from the UK Office of National Statistics). The correlation between those two series over the 1994-2006 observation period is -0.24, smaller in absolute value than what we found for other countries but still clearly negative, as visually clear on Figure 23.

6.2.2 Canada

As a part of its report on "Business Dynamics in Canada" (see Kanagarajah, 2003), Statistics Canada has compiled annual employment shares by firm size categories over the two decades 1983-2003.[18] The largest size category available is 500. From these data, we can compute employment growth rates, with reclassification and including entrants. In Figure 24 we plot the differential growth rate (>500 employees minus <20 employees) against detrended unemployment (from the OECD), and we find the usual negative correlation, in this case equal to -0.43.

7 A Firm-Ladder Model of Employer Size Dynamics

We now sketch a simple model of firm employment dynamics that can help to shed some light on the empirical patterns that we have documented. Intentionally very stylized, this model is meant to illustrate what we have argued in other work is an important mechanism in labor markets. We refer to our theoretical work (MPV08, MPV09) for more details.

Let u_t denote the unemployment rate at (continuous) time t. Let δ_t denote the separation rate and λ_{0t} the job-finding rate from unemployment. Assume constant labor force participation. The equation

$$\dot{u}_t = \delta_t (1 - u_t) - \lambda_{0t} u_t$$

is called a "bathtub" model. As in a bathtub the inflow of water from the faucet and the outflow from the drain determine the level of the water line, so the inflow into unemployment, equal to the separation rate δ_t times the employment rate $1-u_t$, and the outflow, equal to the

[18]Visit www.statcan.gc.ca/bsolc/olc-cel/olc-cel?catno=61-534-XWE&lang=eng for details.

job-finding rate λ_{0t} times the unemployment rate, determine the change in the unemployment rate.

We propose a "firm ladder" model. Workers can search randomly on and off the job and receive offers at rates λ_{1t} and λ_{0t} respectively. When unemployed, a worker accepts any offers. When employed and confronted with an outside offer, a worker always chooses one of the two firms, according to a ranking that all workers agree upon. That is, if any worker prefers working for any firm A over any firm B, so do all other workers, and thus any A employee will reject offers from B and any B employee will accept offers from A. Additionally, some worker lose jobs from all rungs of the firm ladder and become unemployed, as in the bathtub model, so this is a slippery ladder, while the unemployed find firms randomly.

Let $\theta \in [0,1]$ denote a firm's rank in the unanimous worker ordering, and $L_t(\theta)$ its employment size, that we treat as a continuous variable. Normalize the measure of active firms to 1. Assume that, conditional on making a contact with another firm, the probability of sampling a firm with rank below θ, either from employment or unemployment, is $Q(\theta)$, a proper cdf with density q. Applying a law of large number at the individual firm level, the size of firm $L_t(\theta)$ evolves according to

$$\dot{L}_t(\theta) = -\left\{\delta + \lambda_{1t}\left[1 - Q(\theta)\right]\right\} L_t(\theta) + q(\theta)\left[\lambda_{0t} u_t + \lambda_{1t}\int_0^\theta L_t(x)\,dx.\right]$$

The outflow from the firm occurs at rate δ to unemployment and at rate λ_{1t}, the contact rate on the job, times $1-Q(\theta)$, the probability of sampling a higher-ranked firm. The inflow occurs at rate λ_{0t} from unemployment and at rate λ_{1t} from other firms, and a measure $\int_0^\theta L_t(x)\,dx$ are employed by lower-ranked firms and move to θ. In either type of inflow, workers contact a firm of rank θ, conditional on making a contact, with "chance" $q(\theta)$.

This is a Partial Differential Equation in time and rank, with a given initial condition $L_0(\theta)$ which describes the initial distribution of firm size by preference ranking. Notice that this pins down also initial unemployment $u_0 = 1 - \int L_0(\theta)\,d\theta$. If the arrival rates λ_{it} are constant over time, as shown in MPV08, this PDE has the solution

$$\begin{aligned} L_t(\theta) = e^{-\{\delta+\lambda_1[1-Q(\theta)]\}t}&\left[L_0(\theta) + \lambda_1 t q(\theta)\int_0^\theta L_0(x)\,dx\right. \\ &\left.+\lambda_0 q(\theta)\int_0^t \left[1 + \lambda_1(t-s)Q(\theta)\right] u_s e^{\{\delta+\lambda_1[1-Q(\theta)]\}s} ds\right] \end{aligned}$$

which converges to a steady state distribution

$$L_\infty(\theta) = \delta\lambda_0 \frac{\delta+\lambda_1}{\delta+\lambda_0}\frac{q(\theta)}{\left\{\delta+\lambda_1\left[1-Q(\theta)\right]\right\}^2}.$$

Now suppose that $L_0(\theta) < L_\infty(\theta)$ for all θ. That is, the economy starts "depressed" and all firms grow in size, while unemployment falls. At what relative pace do firms of different ranks grow? In other words, is the growth rate of employment by given rank $\dot{L}_t(\theta)/L_t(\theta)$ increasing or decreasing in rank?

The natural application to our setting is as follows. Let θ denote a firm-specific productivity parameter, fixed over time, and assume that $L_0(\theta)$ is increasing in θ. That is, more productive firms start out bigger. Then, given the assumed worker preference ordering over employers, this initial ranking is never reversed. In MPV08 we calibrate the parameters of this PDE ($\lambda_0, \lambda_1, \delta, L_0$ and Q) to match stylized facts about worker turnover (to/from unemployment and job-to-job) and the distribution of employment by size coming out of a recession (L_0), but not the evolution of firm sizes. When we compute the resulting path of $L_t(\theta)$, we find that initially larger, preferred firms converge more slowly to their steady state size. As we observe in the data, the relative growth rate of initially large minus small firms increases from a minimum after a trough to the following peak.

The intuition is simple. All firms have access to an inflow from unemployment that is independent of firm size. This inflow declines over time with unemployment. As the size of small firms grows, the pool of workers that can be poached also rises. While initially small, lower ranked firms cannot poach much and soon lose their hiring pool of unemployed, larger, higher ranked firms compensate this shortfall with an increasing inflow of job-to-job quits from lower firms. See MPV08 for quantitative results.

In MPV08, MPV09 we rationalize these dynamics in the equilibrium of a dynamic large game where firms differ by productivity θ and post employment contracts, which can be summarized by a time-varying continuation utility promised and delivered to the worker. If more productive firms start out larger, they have two incentives to make better offers: they have more to lose from not producing, and they have more workers to lose to competitors. Therefore, more productive firms always offer more and attract workers from less productive ones, which rationalizes our preference ordering over employers. This Rank-Preserving Equilibrium is generic under weak assumptions, and generates the dynamic described above. In our quantitative exercise, wages are also increasing in productivity, thus in size. Hence, ranking firms by initial productivity or size makes little difference to predict the timing of growth, as we indeed find in the Danish data. In MPV09, simulations from a stochastic model with aggregate productivity shocks generate patterns that are qualitatively consistent with our empirical findings both in booms and slumps. More productive, larger firms are more cyclically sensitive.

8 Concluding Remarks

In this paper, we show evidence that *large employers are more cyclically sensitive*. This pattern is robust to a variety of measures of relative employment growth, employer size and classification by size, treatments of entry and exit of firms and establishments, industry, geographical and firm age breakdowns. Evidence on gross job flows and on worker flows by employer size depicts a coherent picture of labor market dynamics. Very similar patterns are observed in other, quite diverse countries.

We conclude by illustrating our research agenda. Although the datasets that we have explored are of high quality in terms of contents, coverage, accuracy, and duration, there is always scope for improvement. First, there is an issue of time span. The main datasets, BDS and BED, cannot be extended back in time given the availability of underlying micro data. Repeated cross-sections like CBP are available in archival form for the post-war period, but not every year. Alternative sources that extend over longer periods cover only specific sectors. Second, finer industry-level analysis appears infeasible because of the 1997 transition from SIC to NAICS. Third, micro business data underlying BDS (the Longitudinal Business Database at the Census Bureau) and BED (at the Bureau of Labor Statistics) would let us replicate our analysis of Compustat, Denmark, and Brazil, allocate firms to size classes once and for all and track their growth over decades. We discussed the pros and cons of this approach. Additionally, information at the establishment level on firm wages, assets, sales, intermediate inputs would allow us to verify whether employer size is really capturing some measure of productivity. Finally, matched employer-employee longitudinal datasets like IDA from Denmark can help us better gauge the firm-specific component of wages that attracts workers to specific employers, as well as connect directly worker and job flows.

From a conceptual viewpoint, the data require a theoretical framework to make sense of the patterns that we uncover. Indeed, even our measurement is strongly influenced by our theoretical work in MPV08 and MPV09. There, we identify a 'firm' with a wage policy, and we base our explanation of the facts on competition for workers among heterogeneous firms in frictional labor markets. Alternative definitions of a firm, based on technology (scale of operation, capital adjustment costs), span of control, borrowing constraints, and others, can be similarly embedded in an equilibrium framework to produce predictions that can be confronted with our facts.

References

[1] Bagger, J., H. Bunzel, B. J. Christensen and N. Pedersen (2009), "Danish Matched Employer-Employee Data: Patterns and Puzzles", Aarhus University *mimeo*.

[2] Bils, M. and K. McLaughlin (2001), "Inter-Industry Mobility and the Cyclical Upgrading of Labor", *Journal of Labor Economics*, 19(1), 94-135.

[3] Bowler, M. and T. Morisi (2006), "Understanding the employment measures from the CPS and CES survey", *Monthly Labor Review*, February, 23-38.

[4] Chari, V.V., L. Christiano and P. Kehoe (2007), "The Gertler-Gilchrist Evidence on Small and Large Firm Sales", *mimeo*.

[5] Davis, S. J., R. J. Faberman, J. C. Haltiwanger and I. Rucker (2008a), "Adjusted Estimates of Worker Flows and Job Openings in JOLTS", NBER working paper No.14137.

[6] Davis, S. J., R. J. Faberman, J. C. Haltiwanger, R. Jarmin, J. Miranda (2008b), "Business Volatility, Job Destruction, and Unemployment", NBER working paper No.14300.

[7] Davis, S. J., J. C. Haltiwanger and S. Schuh (1996), *Job Creation and Detruction*, Cambridge, MA: The MIT Press.

[8] Delli Gatti, D., C. Di Guilmi, E. Gaffeo, G. Giulioni, M. Gallegati and A. Palestrini (2004), "Business Cycle Fluctuations and Firms' Size Distribution Dynamics", *Advances in Complex Systems*, **7**(2), 1-18.

[9] Foster, L., J. Haltiwanger and C. Syverson (2008), "Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability?", *American Economic Review*.

[10] Gertler, M. and S. Gilchrist (1994), "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms", *Quarterly Journal of Economics*, 109(2), 309-340.

[11] Gibrat, R. (1931), *Les inegalités économiques; applications aux inegalités des richesses, à la concentration des entreprises, aux populations des villes, aux statistiques des familles, etc., d'une loi nouvelle, la loi de l'effet proportionnel*. Paris: Librairie du Recueil Sirey.

[12] Jovanovic, B. (1982), "Selection and the Evolution of Industry", *Econometrica* 50(3), 649-670.

[13] Kanagarajah, S. (2003), "Business Dynamics in Canada", Statistics Canada: Catalogue No. 61-534-XIE.

[14] Lentz, R. and D. T. Mortensen (2008), "An Empirical Model of Growth through Product Innovation", *Econometrica* 76(6), 1317-73.

[15] Marsili, O. (2006), "Stability and Turbulence in the Size Distribution of Firms: Evidence from Dutch Manufacturing", *International Review of Applied Economics*, **20**(2), 255-72.

[16] Moscarini, G. and F. Postel-Vinay (2008), "The Timing of Labor Market Expansions: New Facts and a New Hypothesis". Forthcoming in *NBER Macroeconomics Annual*, Edited by D. Acemoglu, K. Rogoff and M. Woodford. [MPV08]

[17] Moscarini, G. and F. Postel-Vinay (2009), "Nonstationary Search Equilibrium" Yale University and University of Bristol *mimeo*. [MPV09]

[18] Okun, A. (1973), "Upward Mobility in a High Pressure Economy", *Brooking Papers on Economic Activity*, 207-52.

[19] Romer, C. and D. Romer (1989), "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz", *NBER Macroeconomics Annual*, IV, 121-70.

[20] Sharpe, S. (1994), "Financial Market Imperfections, Firm Leverage, and the Cyclicality of Employment", *American Economic Review*, 84(4), 1060-1074.

[21] Shimer, R. (2005), "Reassessing the Ins and Outs of Unemployment", University of Chicago *mimeo*.

[22] Sutton, J. (1997), "Gibrat's Legacy," *Journal of Economic Literature*, 35(1), 40-59.