



Discussion paper No. 09/06

Large Employers Are More Cyclically Sensitive

Giuseppe Moscarini
Yale University
and
NBER

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

October 2009

Large Employers Are More Cyclically Sensitive*

Giuseppe Moscarini[†]

Yale University

and

NBER

Fabien Postel-Vinay[‡]

University of Bristol

and

Paris School of Economics

September 2009

Abstract

We present new empirical evidence that large firms or establishments are more sensitive than small ones to business cycle conditions. Larger employers shrink faster, or expand more slowly, during and after a typical recession, and create more of their new jobs late in the following expansion, both in gross and net terms. The differential growth rate of employment between large and small firms is strongly negatively correlated with the unemployment rate, and 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 and size classifications. 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 (partly novel) U.S. datasets, both repeated cross-sections and job flows with employer longitudinal information, starting in the late 1970s 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, sectors and states. We find the same pattern in several other countries, including in longitudinal censuses of employers from Denmark, France and Brazil. Finally, we sketch a simple firm-ladder model of turnover that can shed light on these facts.

*We thank Bob Hall for helpful discussions of this paper and our related work, participants to many seminars and conferences for comments, Jesper Bagger, Carlos Corseuil, Jason Faberman, David Talan and Linas Tarasonis for providing data. Postel-Vinay gratefully acknowledges financial support from the ESRC (grant reference RES-063-27-0090). 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 new empirical evidence that large firms or establishments are more sensitive than small ones to business cycle conditions. The differential growth rate of employment between large and small firms is strongly negatively correlated with the unemployment rate, and varies by about 5% over the business cycle. Specifically, we establish five facts:

1. Large employers destroy proportionally more jobs during and right after recessions, and create proportionally more jobs late in expansions, relative to small employers, both in gross and net terms. Therefore, employers that are initially larger have a much more cyclical one-year ahead growth rate of employment.
2. The higher cyclical sensitivity of large employers is not due (only) to their different entry and exit patterns, but holds also for continuing firms and establishments, as well as for older, established firms.
3. The higher cyclical sensitivity of large employers holds principally within, not across, sectors and States.
4. The higher cyclical sensitivity of large employers is not due to reclassification of employers into larger (smaller) size classes during an aggregate expansion (contraction).
5. The higher cyclical sensitivity of large employers 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

as job flows by initial or end-of-period employer size; employer panels with full longitudinal information. The data often break down the separate contributions of continuing, entering, and closing establishments and firms, of firms of different age, and all entries by industry and location. Particularly useful proved the new Census Bureau's Business Dynamic Statistics (BDS), which covers 1976-2005, as well as matched employer-employee datasets from Denmark, France and Brazil. The different formats of these datasets allow us to address, and relieve concerns about, the effects of two potential sources of bias. First, the *Regression Bias* is a well-known fallacy that creates the illusion of a negative size/growth relationship. We are not interested in the sign of that bias but in whether it changes over the business cycle, and we show that it does not. Second, 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 to 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. 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 empirical findings depict the following view of how business cycles propagate. After a recession, when unemployed workers abound, firms hire mostly and cheaply from unemployment. As the reservoir of unemployment dries out, more productive, larger 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 firm size/growth relationship 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, that is, 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. Interacting size with the detrended unemployment rate in a growth/size regression is likely to invalidate this conclusion. Verifying this conjecture is a task on our current agenda.

The new business cycle Fact 1 that we uncover is reminiscent of Okun's (1973) idea of *Cyclical Upgrading of Labor*, a cross-industry pattern whereby employment reallocates from low- to high-paying industries in booms, and vice versa in recessions (see Bils and McLaughlin (2001) for a recent new interpretation). Instead, the phenomena that we emphasize in this section hold within broad industries and not across. This is surely worth noticing, although it does not apply equally well to many industries.

Our facts appear to plainly contradict a well-established set of facts regarding the sensitivity of small firms to cyclical conditions and to monetary shocks. In a very influential paper, Gertler and Gilchrist (1994) [GG] present evidence that small firms, which they argue are more credit constrained, are more sensitive to monetary policy shocks as measured by Romer and Romer (1989) [RR]. GG use the Quarterly Financial Report for Manufacturing Corporations (QFR), 1958-1992. This dataset defines firm size in terms of nominal sales, raising the issue of industry-specific price indexes within manufacturing, and is a set of repeated cross-section, which lack longitudinal links, leading GG to make an ingenious yet ad hoc correction to avoid the Reclassification Bias. Even then, GG's conclusion that small firms are more cyclically sensitive holds, if any, at RR dates, which are notoriously contro-

versial, and not during NBER-dated recessions. Of the six recessions in GG’s 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 RR and NBER dates in the QFR time series, and extend GG’s analysis 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 RR shocks.

Our measure of performance and size is employment, not sales and capital, because this is what we are primarily interested in and because we have employment growth data that are immune from Reclassification Bias, most notably the BDS for the U.S. and longitudinal business micro data from other countries. Another U.S. dataset which still suffers from this bias, the Business Employment Dynamics (BED), is updated more promptly and allows us to give a first look at the current recession that began in 2007. Consistently with common wisdom, we see the telltale mark of a credit crunch, in that small firms suffered more in the crucial last quarter of 2008. When considering the entire 1992:Q3-2008:Q4 BED sample period, large employers are still more cyclically sensitive. Going further back in time, King (1923) found that in the first quarter of 1920, the onset of the deep 1920-1921 recession in the U.S., firms employing less than 21 workers started with 1/3 of total employment, but were responsible for just 1/20 of the (sharp) reduction in aggregate employment.² All in all, the different focus (RR dates vs. NBER dates and unemployment) explains most of the difference between our and GG’s findings. The conventional wisdom that “small businesses are the engine of job creation” finds some empirical support in our data for periods of high

¹Sharpe (1994) replicates their findings at RR dates for employment growth, by initial size defined in terms of net capital. He uses the NBER Manufacturing Panel from Compustat. As we showed in MPV08, and discuss again in this paper, in the full Compustat panel comprising all industries, over a longer time period (1975-2005), our pattern of differential growth rate by initial size emerges quite clearly at NBER-dated business cycles.

²See Tables VI-VIII and page 31: “[...] these records give unequivocal evidence that it is primarily the large concern which is affected by a business depression [...]”. We thank Mark Bills for this reference.

unemployment, recessions and their aftermaths, which is presumably when jobs are more needed. This statement clearly fails in tight labor markets, when job creation is taken over by large employers.

Finally, our findings may suggest a solution to a long-standing empirical puzzle, the discrepancy between two monthly measures of aggregate U.S. employment.³ The BLS interviews about 50,000 households with 140,000 individuals in the Current Population Survey (CPS) to generate the unemployment rate, and surveys about 160,000 firms with 400,000 establishments in the Current Employer Survey (CES), to generate official employment figures. The CES/CPS ratio of employment measures has been increasing and strongly procyclical in the last two decades (Bowler and Morisi, 2006). The increase in this ratio suspiciously mirrors the one in the average size of U.S. firms, and the cycle the relative performance of large and small employers, that we find in our data. Given well-know problems that any payroll survey encounters when tracking small employers, the CES sample may be skewed towards large employers and may miss small, young firms, which are instead captured by the CPS.

In Section 2 we illustrate 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. We then present auxiliary evidence broken down by employer size, in Section 4 from repeated cross-sections of employment distribution, and in Section 5 from worker flows, specifically job-to-job quits and layoffs. Section 6 presents evidence from several countries other than the U.S.. Section 7 illustrates a simple model of firm size dynamics. Concluding remarks follow.

³We thank Bob Hall for suggesting a possible connection.

2 Methodology

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.

2.1 Definitions

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 accurate measures of productivity are hard to obtain, and those that do exist 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.

Our measure of relative employer performance is 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 fewer 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. Our main measure of the economy’s business cycle conditions is the detrended civilian unemployment rate. Again, this is motivated by our theoretical work.

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 :

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

Here $\alpha : \mathbb{N}^2 \rightarrow [0, 1]$ is the weight on end-of-period size, and $1 - \alpha(L_{it}, L_{it-1})$ on initial size. This is a weighting function that can depend on both numbers. Let the weighted growth rate

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

Finally, let β denote another weighting function, and $\bar{L} > \underline{L} > 0$ two integers that define “large” employers ($L_i^{(\beta)} \geq \bar{L}$) and “small” employers ($L_i^{(\beta)} \leq \underline{L}$) according to β -weighted size. Let the growth rate between $t - 1$ and t of employment at all employers that are classified at time $s \leq t$ as large:

$$g_{s,t,LARGE}^{(\alpha,\beta)} = \frac{\sum_{i:L_{is}^{(\beta)} \geq \bar{L}} (L_{it} - L_{it-1})}{\sum_{i:L_{is}^{(\beta)} \geq \bar{L}} L_{it-1}^{(\alpha)}}$$

Notice that, using the weighting β , we can choose the size class over which to compute net job creation (the numerator) by either initial, average, or end-of-period size observed at any given date $s < t$. This initial date s could be either fixed once and for all with longitudinal data, or reset every period, either at the beginning ($s = t - 1$) or the end ($s = t$) of the period. Also, we can assign to size classes the numerator (net job creation) independently of the denominator (baseline employment), using different weighting functions α and β . Similarly for the “small” size class, $g_{s,t,SMALL}^{(\alpha,\beta)}$ for $L_{is}^{(\beta)} \leq \underline{L}$.

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

$$\Delta g_{s,t}^{(\alpha,\beta)} = g_{s,t,LARGE}^{(\alpha,\beta)} - g_{s,t,SMALL}^{(\alpha,\beta)}$$

and in particular in how its deviations from trend correlate with the business cycle conditions at $t - 1$. Our preferred cyclical indicator is the civilian unemployment rate, detrended over the post-war period. Our main statistics of interest is $\text{corr} \left(\Delta \hat{g}_{s,t}^{(0,0)}, \hat{u}_{t-1} \right)$, where hats denote absolute deviations from trend. As we showed in MPV08, using the employment/population ratio instead makes little difference, as participation has no clear cyclical pattern. To detrend series, we use a Hodrick Prescott filter. For the unemployment rate, following Shimer (2005) we use a high smoothing parameter (8.1E6 at monthly frequency). For relative growth rates and shares’ growth rates, a high smoothing parameter is also necessary so that no obvious cyclical pattern is left visible in the fitted trend. Fitting a linear trend makes little difference in this case.⁴

⁴The Census Bureau publishes County Business Patterns aggregate statistics that include the number of

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. A common one 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. 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. We now analyze this issue in more detail.

2.2 The Regression Bias

If employer size L_{it} is mean-reverting, then employers that are initially small 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, Galton's fallacy, 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 the previous one, $\alpha = 1$, is probably much more problematic than this time convention, which is $\alpha = 0$, 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 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:Q3-2008:Q1. 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 typically less 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.

the growth rate between times $t - 1$ and t is the average of employment over $t - 1$ and t . This reduces the mean reversion fallacy, and is well defined both for entrants, who have $L_{it-1} = 0 < L_{it}$, and closing employers, who have $L_{it-1} > 0 = L_{it}$.

Whether mean reversion affects also the *cyclical* of the *relative* growth rate $\Delta g_{s,t}^{(\alpha,\beta)}$ is questionable, and depends on the specific statistical model of firm growth implied by the underlying structural model. To err on the side of safety, in some datasets we use both $\Delta g_{t-1,t}^{(0,0)}$, which is differential growth between initial size classes, and $\Delta g_{t-1,t}^{(1/2,0)}$, differential growth between employers classified by their initial size but where the growth rates are computed dividing net job creation by average employment over the period. We find 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 $\Delta g_{t_0,t}^{(0,0)}$, 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. Our empirical results prove robust.

2.3 The Reclassification Bias

Consider $\Delta g_{t-1,t}^{(\alpha,1)} = \Delta g_{t,t}^{(\alpha,0)}$, which is relative growth between $t-1$ and t of employers classified by their end-of-period size at t . If the economy grows, and all employers with it, while the size cutoffs $\bar{L} > \underline{L}$ 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 employment falls. This is a bias that produces the illusion of procyclical relative growth by size, precisely the fact that we aim to document.

This bias appears in different forms in the various types of datasets that we employ. Its most straightforward manifestation arises when using repeated cross-sections of employment shares of size classes. That is, if we lack longitudinal links and only observe a time series of

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

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 employment growth for size class j . When small firms grow faster than large ones, their share of total employment rises. However, as small employers gain size, they are reclassified into larger size classes, so repeated cross-sections are subject to Reclassification Bias. Specifically, define the growth rate of aggregate employment as $g_t = \sum_i L_{it} / \sum_i L_{it-1} - 1$. Then simple algebraic manipulations yield the following approximation near zero:

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

The growth rate of the employment share of a given size class j on the left hand side decomposes into the sum of three terms: $g_{jt}^{(0,0)}$, 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 the net growth rate of employment imputed to size class j . This third term reflects 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.

The Reclassification Bias also appears in a more complex form in the Business Employment Dynamics dataset, maintained by the Bureau of Labor Statistics. Recall that Job Creation (JC) is the addition of employment positions at all units that expand, and vice versa for Job Destruction (JD), so JC–JD is net job creation. The “Dynamic allocation” method used by the BLS to impute firm size in BED modifies 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 β is $\beta' = (L_{it} - L_{j(i),t}) / (L_{it} - L_{it-1})$

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 $\beta = 0$. 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 $\Delta g_{t-1,t}^{(1/2,\beta')}$. This method is vulnerable to the Reclassification Bias, both in the numerator (dynamic allocation) and the denominator (average employment).

The potential quantitative importance of the Reclassification Bias is illustrated in the following 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, “small”, “medium” and “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 class, and 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, and again -57.2% for the large class, coming to a difference of -90.5% . If we apply 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 $2(1566 - 1000 + 999 - 1690) / (1566 + 1000 + 1690 + 999) = -4.75\%$, difference 6.9% . Finally, the employment shares of the small and large class both decrease from (resp.) 3% and 79% at $t - 1$ to 2.8% and 66.5% at t , but the share of large firms falls much more. Qualitatively, all methods except dynamic allocation suggest that small firms are doing relatively better. Quantitatively, the results vary widely, but the example is intentionally extreme.

In order to quantify the Reclassification Bias in the data, we use longitudinal business micro data, where we can fix an initial size class for each employer either every period before

computing growth, or 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 years later, in a way that depends on cyclical conditions.

Nevertheless, we will also present evidence from datasets where a Reclassification Bias may arise either because the allocation of employment stocks and flows is not made by initial size or because we proxy growth rates of employment by size with those of employment shares from repeated cross-sections. By comparing these results to those from longitudinal datasets, we will show that the Regression Bias is quantitatively modest. As data on employment shares by size classes are relatively easy to come by in many countries and over long time periods, this suggests a promising avenue to expand further the scope of our investigation to more countries.

2.4 Entry and Exit

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. Similarly, exit may be systematically related to the size of the firm or establishment in its last period of existence. A natural question is whether the higher cyclical sensitivity of large employers is due entirely or in part to differential entry and exit by large and small firms in booms and recessions.

We found no compelling reason to exclude exiting employers if we know their initial size class. Nonetheless, we provide 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. The datasets at our disposal, however, do not allow a distinction between firm exit and establishment closing by surviving firms. Thus, we either include or exclude establishment deaths, which contain both types of exit.

For entrants, the main issue is the attribution of their first growth to their initial size class. Even longitudinal datasets, which allow us to classify JC and JD by initial size, do not contemplate an initial size class “0 employees”, hence they attribute the net JC of entrants to their post-entry size class (when first observed). Using information on firm age and on establishment birth status, we can correctly attribute the employment growth due to entrants, and we present our evidence with entrants both included in and excluded from the “small” size class. We can and will distinguish between entry of a new firm and opening of a new establishment by an existing firm. 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.

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

3.1 Business Dynamic Statistics (BDS)

3.1.1 The Dataset

The primary source of information on the identity, location, employment, sales, payroll, and industry of all U.S. businesses is the Census Bureau’s Business Register, formerly known as Standard Statistical Establishment List. For establishments in the Business Register, information is updated annually with the Company Organization Survey. The list of establishments in the Business Register is updated at every quinquennial Economic Census.⁵ Starting in December 2008, the Center for Economic studies at 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. private employment, and contains information on establishment-level employment stocks and job flows, for continuing, entering, and exiting establishments, at annual frequency for the 1976-2005 period, broken down by location and industry of the establishment, and by age and size of

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

the parent *firm*. A firm is an aggregate of all establishments (the original unit of observation) that share a certain identifier, mainly the tax Employer Identification Number. Two notions of firm size are available in BDS: average size between last year and this year, partly after the job flows have taken place, and “initial firm size” *before* the flows are measured. 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 running from March of year $t - 1$ to March of year t , and a measure of initial ($t - 1$) employment. Everything is classified by a measure of initial firm size, as of March of year $t - 1$. In BDS, 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 $t + 1$. One exception are new firms (initial age 0), who are attributed to the initial size class according to their end-of-period size. Hence, we first reset their initial size to zero and remove their contribution to net JC from the size class they reach ex post, to impute it to the initial size class “0 employees”, which is part of the “small firm” group. For pre-existing firms, the employment at time t by initial (year $t - 1$) size class reported in BDS is the stock in year t . We recover the employment of those firms, in each initial size class, in year $t - 1$ by subtracting from employment at t their net JC between $t - 1$ and t . We then study the differential growth rate of employment between initially large and small employers including all firms and establishments, which is our main focus. To highlight the role of entry and exit, we then exclude one by one the contributions to both initial employment and subsequent net JC of: new firms; new establishments, even if added by pre-existing firms; closing establishments; new and closing establishments (hence focusing only on continuing

⁶In 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.

establishments).

Our choice of size cutoffs that define a “large” and “small” employer is necessarily arbitrary, but guided by the following consideration. Let \mathfrak{s}_j denote the average employment count in an establishment owned by a firm in size class $[L_j, L_{j+1} - 1]$, for example $[1,4]$ employees, $[500,999]$ etc. Clearly, $\mathfrak{s}_j < L_{j+1}$, 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 $\mathfrak{s}_j \geq L_j$, 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 $\mathfrak{s}_j/L_j < 1$. 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).

Table 1 shows that in the BDS the ratio \mathfrak{s}_j/L_j 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 multi-establishment firms. Indeed, the average establishment size stabilizes around 60 employees for all firms of total size above 500. Therefore, to define small and large firms, we choose employment size classes ≤ 49 and ≥ 1000 .

Firm size category L_j to $L_{j+1} - 1$	Mean establishment size \bar{s}_j	Size ratio \bar{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. Source: BDS and authors' calculations.

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, including the contribution of entry and exit, starting in 1979. Consistently with our definitions, our time convention in this and all following graphs is that observations in period t denote growth between $t - 1$ and t . In the case of the BDS, observations occur every year in mid-March, so employment growth plotted in year t takes place mostly during calendar year $t - 1$. It is important to keep this in mind to correctly interpret the behavior of the differential growth rate during NBER-dated recessions, the shaded vertical bars.

Employment at initially large firms suffers much more as a consequence of recessions. The growth rate of large firms is higher in the few years preceding each of the three NBER peaks and lower in the few years after. There is a slight delay in the 1990-1991 episode, but that was a very short recession, and one third of the March 1990 - March 1991 employment growth reported as the 1991 observation took place between March and July 1990, when the economy was still expanding.

This picture corroborates only in part the common wisdom that small businesses are

the engine of (net) job creation. This wisdom has been criticized as subject to Galton's regression fallacy, which is known to generate a negative size-growth relationship (see for example Davis et al., 1996). 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 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 account for 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 destruction 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. 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 relative to small ones late in expansions to account for their better performance in those phases, including again in the mid 2000s.

To better visualize the cyclical effect, we take the difference in growth rates between (initially) large and small firms. If this difference is procyclical, as we will find, then large firms are more cyclically sensitive: they shed proportionally more of their employment when aggregate employment is falling. Because our underlying theoretical hypothesis emphasizes the role of unemployment and labor market tightness as a cyclical indicator, and because NBER monthly datesmark cyclical slumps and recoveries too narrowly for our purposes, due to the annual frequency of the data, Figure 3 reports the differential growth rate against the civilian unemployment rate, specifically the average of monthly deviations from trend of the unemployment rate from March of the previous year to March of the current year. Here is

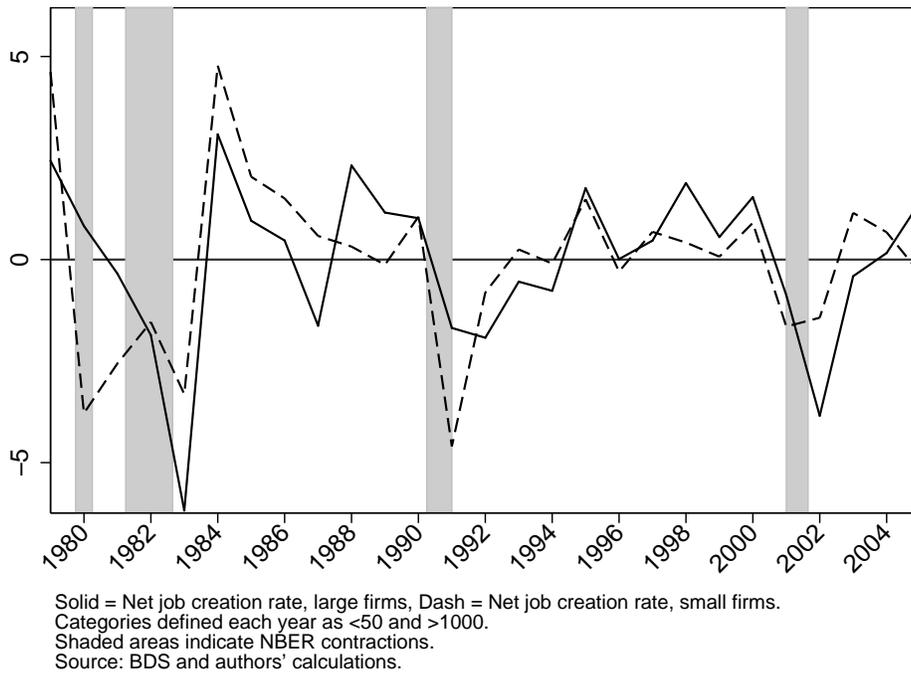


Fig. 1: Net job creation rates

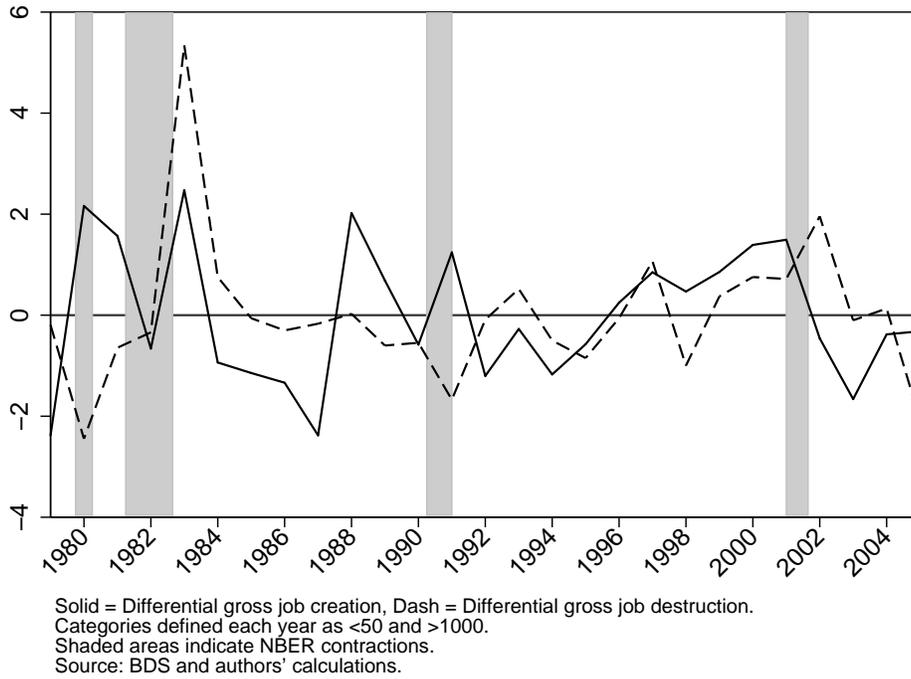


Fig. 2: Differential gross job flows

the central finding of this paper, visually clear and confirmed by the -0.54 correlation:

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. Therefore, 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 year immediately following the NBER trough (November 1982, March 1991 and November 2001), so that most of the growth reported in year t took place in year $t - 1$, and plot it in a worm graph for several years before and after, without any prior detrending. 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 Fact 1, we point out that BDS, which begins in 1977 (that is, with job flows between 1976 and 1977), shows some troubling discrepancies with comparable, well-proven datasets in the first two years, 1977 and 1978. In particular, total net JC in BDS in 1976-1977 and 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 the behavior of our differential growth rate, which is quite unusual in 1976-1978, suggesting serious data issues of unknown nature. We omit the plot for brevity (available upon request). Because CBP is a true and tested dataset while BDS is a new resource, but is immune from reclassification, we begin our analysis in 1979, since when CBP and BDS coincide. Fortunately, 1976-1978 were not recession years, and we do not lose much in terms of cyclical variability in the economy.

⁷Although 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.

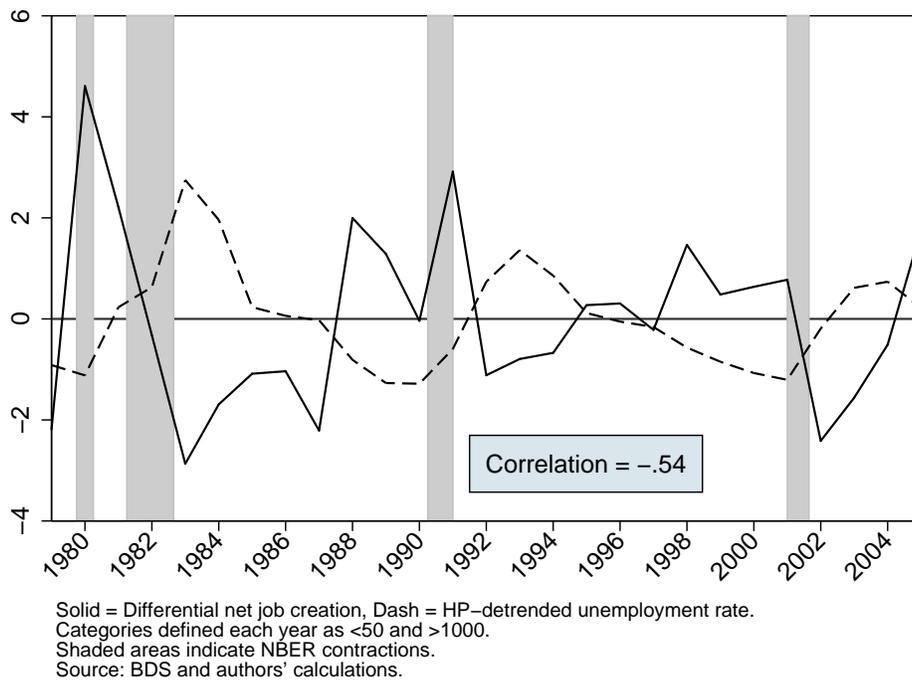


Fig. 3: Differential growth

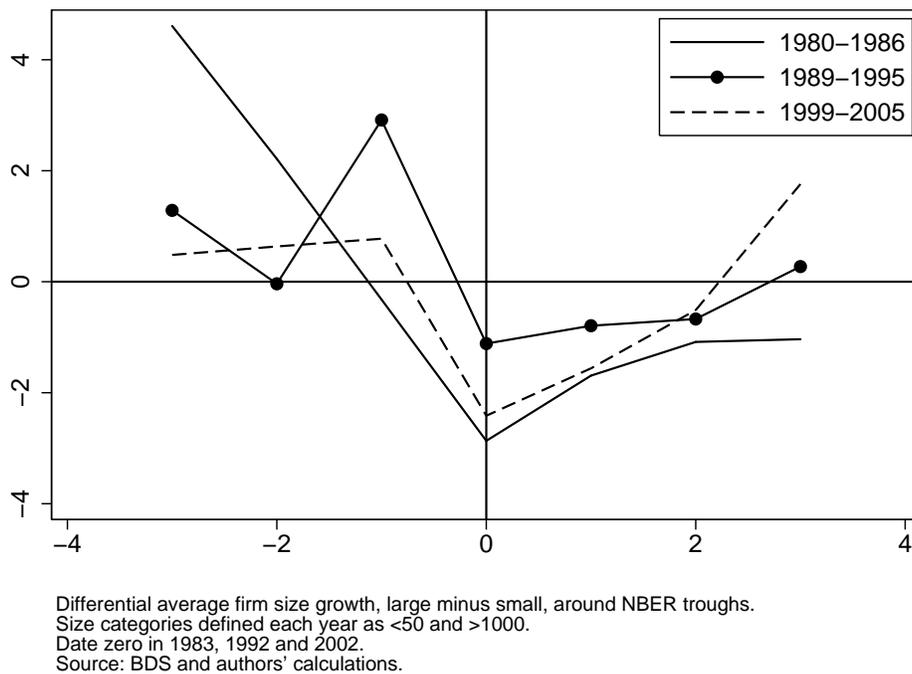


Fig. 4: Differential growth, around three consecutive NBER troughs

3.1.3 Entry and Exit

As mentioned, we can isolate the contributions of entry and exit by eliminating from our computations new firms, new establishments, and/or closing establishments. In the interest of brevity we do not report plots but only correlations between the relative growth rate of employment between large and small employers and the unemployment rate over the period 1979-2005. When we exclude new firms, whose initial employment is zero (so they belong in the “small” group, less than 50 employees) but which create much new employment, this correlation is -0.59 . When we also exclude new establishments, even if added by pre-existing firms, the correlation is -0.57 . To exclude closing establishments, we need to subtract their employment from the initial stock, as well as their JD from the total employment growth, for each size class. The correlation then equals -0.42 . Finally, when we exclude both new and closing establishments, the correlation is -0.47 . Qualitatively, our main Fact 1 is immune to all possible treatments of entry and exit. Quantitatively, the only significant difference is made by establishment closings by large firms. When excluding such closings, the correlation of interest weakens, still remaining quite negative. This suggests that our Fact 1 is due (in a statistical sense) in part to large firms shutting down entire establishments more than small firms in and after recessions. This should not be surprising, as almost all of our small firms are mono-establishment, hence closing an establishment for them means closing shop altogether, while large firms can choose between an intensive and an extensive margin of employment reduction.

Recent work on business micro data (e.g. Foster et al., 2008) as well as established theories of firm dynamics (Jovanovic, 1982) point to the age of a firm as a major predictor of its future behavior. Older firms are much less volatile and growing, conditional on survival, than younger firms. Furthermore, our underlying theoretical framework focuses on firm size as a proxy for productivity. This approximation is likely to be more accurate for older firms, which are closer to their steady state size, while very young firms are still learning their true quality, so we would expect Fact 1 to emerge even stronger among established firms. We

repeat our computation by restricting attention to firms that are initially at least 3 years old, where the age of a firm is the one of the oldest establishment it owns, and is typically not reset to zero by mergers and acquisitions. This automatically excludes entrants, and focuses on the behavior of more established companies. We find a -0.61 correlation between the differential growth rate of employment between initially large and small older firms and the unemployment rate. As we expected, when restricting attention to older firms, Fact 1 is strengthened. If we exclude new establishments, the correlation becomes -0.54 . If we exclude closing establishments, the correlation is -0.48 . Finally, if we exclude new and closing establishments, and focus only on continuing ones, the correlation becomes -0.38 . If we focus on firms that are at least 4 years old, the results are virtually unchanged.

Fact 2. *The higher cyclical sensitivity of large employers is not due (only) to their different entry and exit patterns, but holds also for continuing firms and establishments, as well as for older, established firms.*

3.1.4 Industry Patterns

We now dig deeper and check whether the basic Fact 1 that we uncover, that initially larger firms have more cyclically sensitive employment, holds within or across geographical locations and industries. One important proviso is that in BDS the location and industry refer to the establishment, not to the parent company like initial size and net job creation. It is of course impossible to attribute a unique location and industry to most large firms that have establishments in many U.S. states and industries. 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 by the Census 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.

One natural question is whether the more pronounced cyclical sensitivity of large em-

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 would not diminish the substance of our finding, although it may appear less surprising in light of what we already know about the average 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	
FIRE	-.634	Wholesale Trade	-.34
Retail Trade	-.499	Mining	-.236
TCPU	-.446	Services	-.199
Construction	-.421	Agr. Svc., Forestry, Fishing	-.026
Manufacturing	-.35	All	-.54

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

We maintain our classification in nine broad sectors, which is reasonably consistent through the changeover from SIC to NAICS in 1998. In Table 2 we report the correlations between the differential growth rate of (initially) large minus small firms in each broad sector and the civilian unemployment rate, both detrended.⁸ The small and large firm cutoffs are set here to < 50 and > 500 employees for all industries. Similar calculations using as cutoffs the first and third quantiles of the average (over the entire period) firm size distribution within each sector — 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. The familiar pattern emerges also within most sectors, including the larger ones.

Finally, we perform 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

⁸Caution should be used in interpreting the figure for Agricultural Services, Forestry and Fishing as the BDS leaves agricultural production workers out (among a few other categories of workers). Overall in our BDS sample, that particular industry accounts for less than 0.7% of total employment.

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 almost disappears. The relative growth differential has a correlation with the unemployment rate of just -0.099 .

3.1.5 Geographical Patterns

We associate with each U.S. state the corresponding local civilian unemployment rate from the BLS, and correlate the differential (large minus small firms) employment growth rate and the local unemployment rate, both detrended. We should note, however, that in the BDS file presenting job flows and stocks by initial size, state, and age of the firm that we use to compile Table 3 about 20% of all observations are suppressed, which is likely to attenuate any correlation we find. Table 3 reports the results. The phenomenon that we identify takes place within most states, and is not driven by employment moving from locations with small firms to locations with large firms in a boom, and vice versa in recessions.

Wisconsin	-.527	California	-.278	Nebraska	-.052
North Carolina	-.507	North Dakota	-.266	Washington	-.049
Kansas	-.471	New Jersey	-.265	Ohio	-.048
Florida	-.431	Tennessee	-.26	Montana	-.035
Georgia	-.415	New York	-.253	Maine	-.03
Illinois	-.408	Missouri	-.25	New Mexico	-.019
Virginia	-.364	Connecticut	-.233	Alaska	-.017
Idaho	-.358	Utah	-.224	Oklahoma	.024
Iowa	-.348	South Carolina	-.174	Michigan	.033
Alabama	-.338	Oregon	-.171	Massachusetts	.037
Pennsylvania	-.333	South Dakota	-.136	Delaware	.048
Kentucky	-.33	Nevada	-.13	Louisiana	.075
Arkansas	-.308	Vermont	-.107	New Hampshire	.094
Minnesota	-.308	Maryland	-.106	Indiana	.099
Mississippi	-.303	Texas	-.074	Rhode Island	.109
D. C.	-.29	Hawaii	-.072	West Virginia	.2
Colorado	-.278	Arizona	-.056	Wyoming	.22

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

We summarize the findings of this and the previous Subsections in:

Fact 3. *The higher cyclical sensitivity of large employers holds principally within, not across, sectors and States.*

3.2 Business Employment Dynamics (BED)

A different but still exhaustive source of information on the distribution of employment and on job flows by employer size in the U.S. is available from the Business Employment Dynamics program at the BLS. This program collects information accruing from the States' unemployment insurance programs. 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 magnitude of that bias found in BDS (and documented below in Section 4) suggests that valuable information is contained also in BED flows as broken down by dynamically classified size. The BED only begins in 1992, but has the dual advantage over the BDS of a quarterly frequency and also of currently including a new cyclical peak in 2007:Q4 and the critical 2008:Q4, arguably the nadir of the current recession.

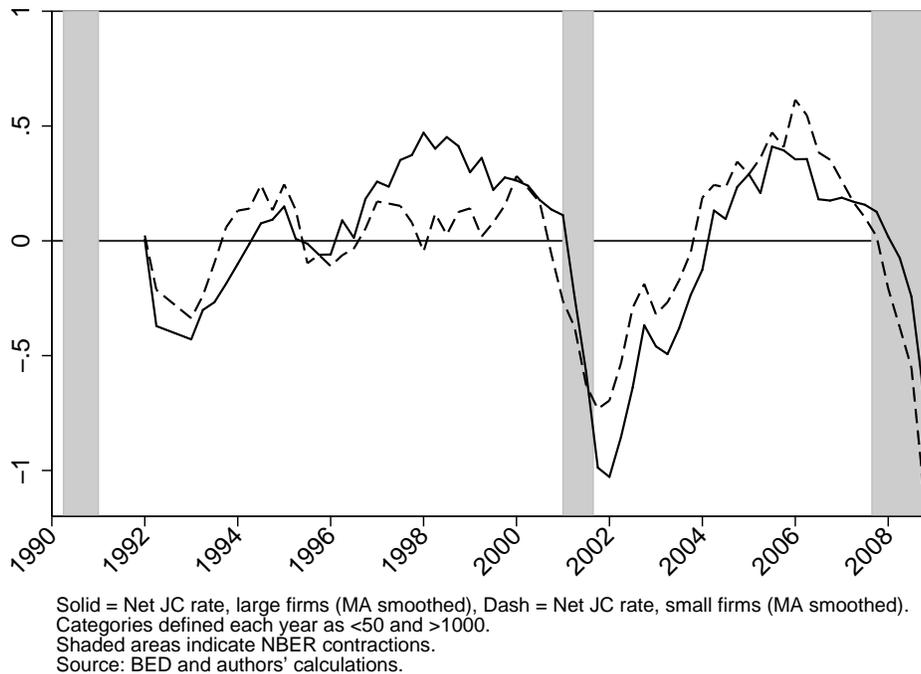


Fig. 5: Net job creation rates

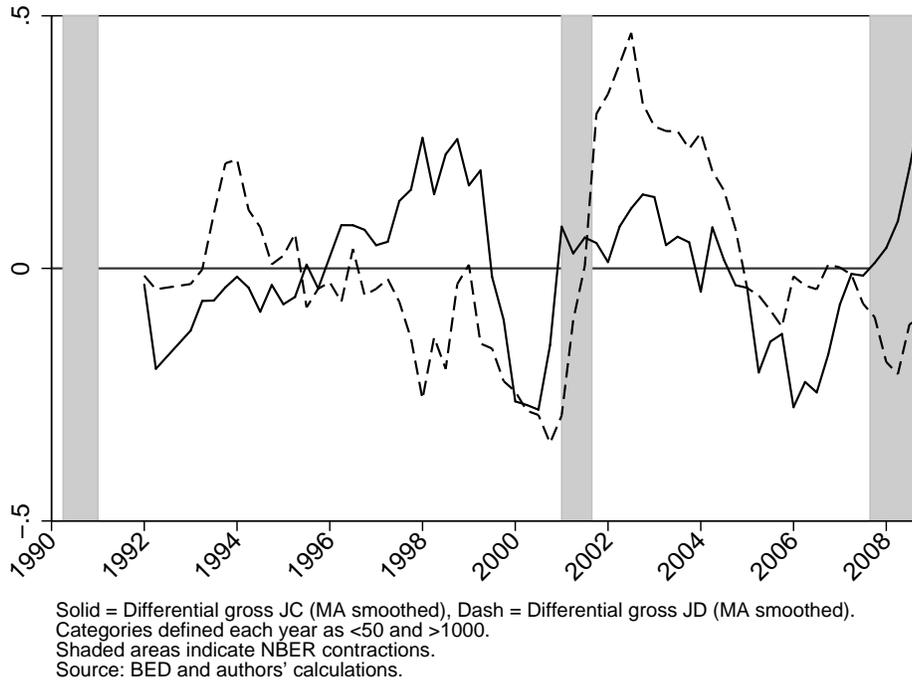


Fig. 6: Differential gross job flows

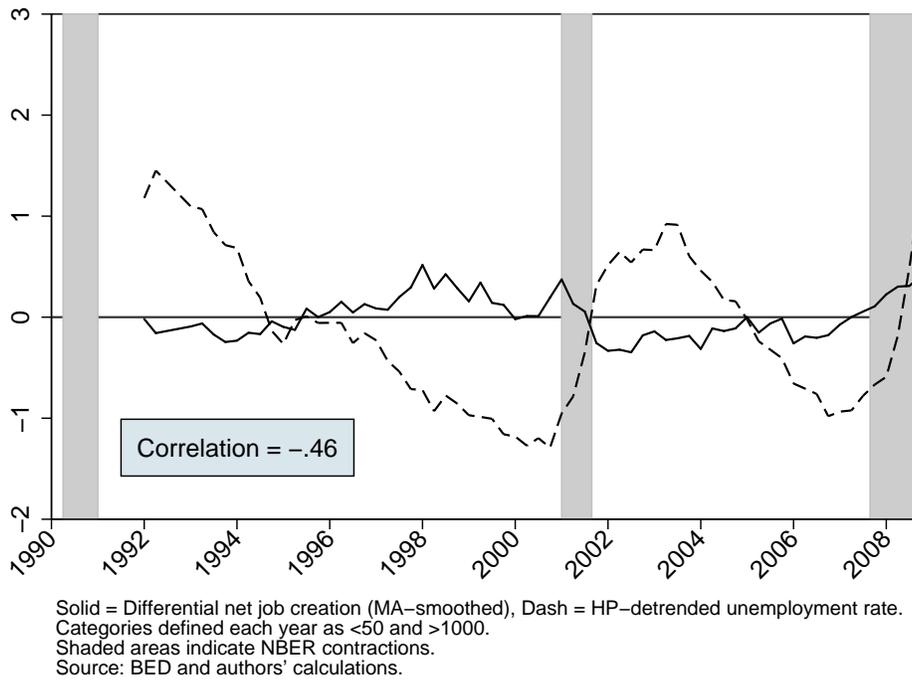


Fig. 7: Differential growth

Figure 5 plots the net JC rate of large and small firms, according to the usual size cutoffs, and NBER recessions, including the (censored) current one.⁹ Figure 6 decomposes these net flows into differential gross JC and differential JD, large minus small firms. In Figure 7 we take the difference in net growth rates and plot it against the detrended unemployment rate. The two series mirror each other, except in the current recession, which shows the telltale mark of a credit crunch, hitting small firms harder. Despite this episode and its obvious importance in such a short sample, the correlation between the two series over the last 17 years is -0.46 . The message from these BLS data is the same as from 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 given date and to follow its growth over a long period of time including several business cycles, as a function of its initial size. This procedure eliminates reclassification altogether and lets any regression to the mean play out. But it requires selecting the set of companies that survive through the entire sample period, introducing a strong survivorship bias. If we tracked all firms in existence at the initial date, 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.

The only set of fully longitudinal micro business U.S. data 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. In 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 we calculate the growth rate of employment over the past year at (initially) large firms and subtract the growth rate of the other, initially small firms in the sample. In Figure 7, MPV08 we plot

⁹Recall 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 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. The correlation with the detrended unemployment rate is -0.61 .

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

As discussed in Section 2, 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 now report evidence on changes in the employment distribution from repeated cross-sections, both to gauge (by comparison with our previous results) the magnitude of the Reclassification Bias, as formally measured by the third term in Equation (1), 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 BDS

The BDS also allows construction of employment shares of size classes each year. We can thus repeat the exercise of subsection 3.1.2 only using the growth rates of employment shares across size classes, whose correlation with the growth rates of employment by initial firm size is 0.91. Figure 8 parallels the aggregate picture of Figure 3, and plots the difference in growth rates of employment shares against unemployment. The message is essentially unchanged. The correlation with unemployment is slightly more negative in Figure 8 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, and vice versa in recessions. Yet, the bias appears to be quantitatively modest.

Fact 4. *The higher cyclical sensitivity of large employers is not due to reclassification of employers into larger (smaller) size classes during an aggregate expansion (contraction).*

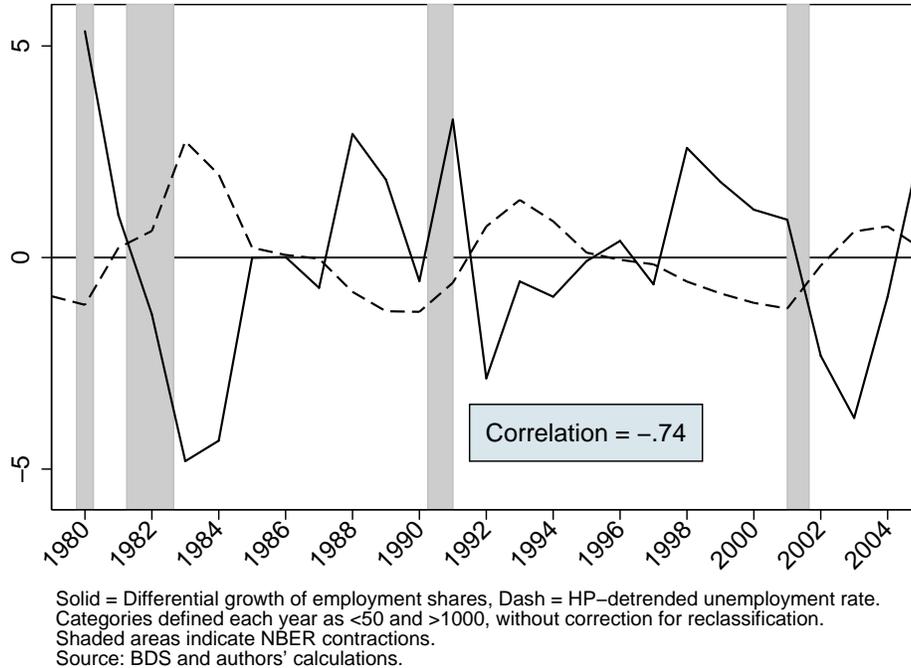


Fig. 8: Differential growth of employment shares

4.2 BED

The BED dataset reports 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. Dividing flows by rates we can back up average employment stocks $L_{it-1}^{(1/2)}$. Again, we detrend the growth rates of employment shares and correlate the deviations from trend with those of the unemployment rate, at the same quarterly frequency. The results corroborate our previous findings from BDS: the negative correlation is again strong, at -0.55 . 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.

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.

5.1 Hires from Other Employers

In MPV08, MPV09 we argue that the reason for the faster growth of small firms in a 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. In MPV08 we corroborate this hypothesis with evidence from the Survey of Income and Program Participation (SIPP) over 1997-2004. 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 11 in MPV08 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. In other words, it plots a measure of the importance of poaching in the recruitment activity of establishments, by size of the hiring establishment. While this admittedly constitutes very limited evidence, we still draw the following two conclusions. 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: 1-25, 26-99, and 100 employees or more. Second, larger establishments almost always poach more than smaller ones. This difference in “poaching intensity”, however, is pronounced in 1997-1999, when the labor market turns tight, and almost vanishes in 2002-2004, when it is slack.

5.2 Layoffs

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.¹⁰ 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.

By comparing mass layoffs (MLS), which can only belong to medium-to-large firms, with

¹⁰A 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.

layoffs at all employers (JOLTS), we can check whether in the 2001 recession and again in the current slump mass layoffs rise faster than layoffs in general. Figure 9 leaves no doubt that mass layoffs, as % deviations from their trend, are manyfold more volatile and concentrated in recessions than total layoffs.¹¹

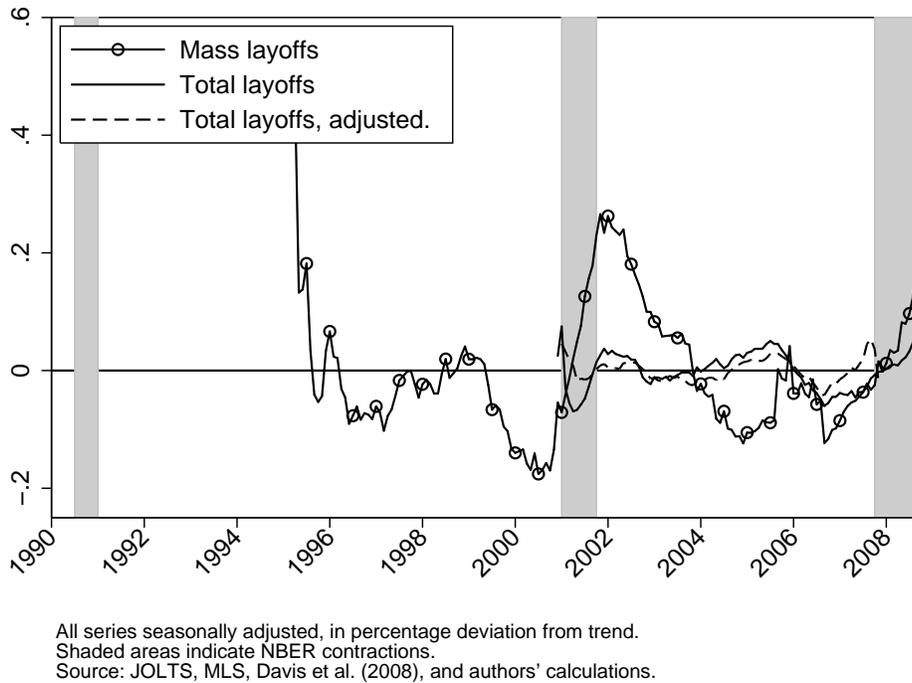


Fig. 9: Mass layoffs vs. all layoffs

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

¹¹JOLTS 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 9 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. Almost all of this literature focuses on cross-sectional patterns, either at one point in time or on average over some time period. We could find two articles that fit a power law to cross-sectional data year by year, in large samples of firms from the Netherlands (Marsili, 2006) and Italy (Delli Gatti et al., 2006). They both find a countercyclical Pareto slope, implying a flatter power law and a more unequal size distribution in booms, and a steeper law in recessions. Although this very indirect evidence, our Fact 1 does imply that the size distribution of employers tends to be more heavily weighted on large ones, hence to be more unequal, when unemployment is low, and vice versa.

Much closer to us is the recent analysis of West-German data by Bachmann and David (2009). These authors use two different datasets, the IAB Employment Sample (IABS) and the LIAB, both derived from a German administrative source (the Employment Statistics Register) to investigate cyclical patterns of job and worker flows in Germany. By design, those two samples make it possible to correct for the Reclassification Bias, as discussed earlier in this paper. The longer of the two samples effectively covers 1980-2003, a period which spans three recessions and two complete business cycles. Bachmann and David reach the following main conclusions. First, the cyclical behavior of labor market flows varies considerably with both worker and establishment characteristics, with establishment size playing an especially important part. They establish that smaller establishments increase their share of total hirings in recessions, while the hiring activity of larger employers declines during recessions and mainly takes place in mature phases of expansions. Second, concerning the nature of observed worker flows, they observe a clear shift from EU to EE transitions in the

mature phase of the expansion. While large establishments tend overall to hire proportionately more from employment (30% v. 42% on average), Bachmann and David conclude that “larger establishments hire an increased number of employed job searchers during the later phase of the expansion, while smaller establishments mainly hire individuals out of unemployment, and start doing so earlier than large establishments”. Not only are Bachmann and David’s results exactly in line with the cyclical pattern of relative employment growth across size classes that we document for the U.S. in this paper, but they also lend strong support to the hypothesis formulated in MPV08 and MPV09 (and restated earlier in this paper in Subsection 5.1) about the nature of labor market adjustments over the cycle.

We now look directly at new data from several countries.

6.1 The Employers’ Growth/Size Relationship over the Business Cycle

For three countries (viz. Denmark, France and Brazil), we have been able to access full longitudinal business microdata on employment from censuses at annual frequency for an extended period of time, allowing us to replicate the exercises that we performed for the U.S. with BDS and Compustat. In all three cases we compute the differential growth rate between initially large and small firms, where “initially” refers either to last year or to a given year fixed once and for all at the beginning of the sample. Also in all three 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

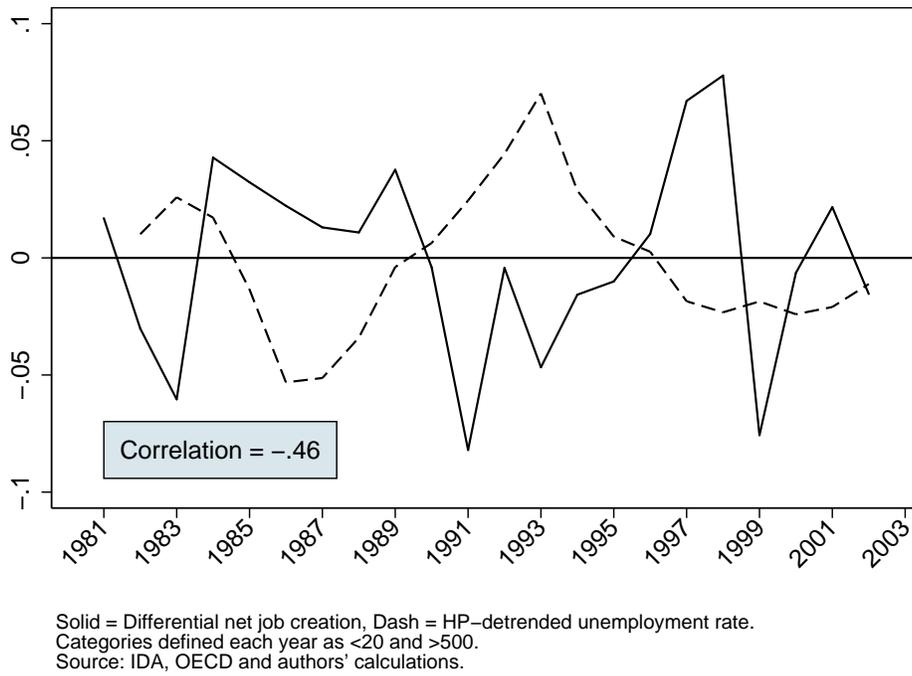


Fig. 10: Differential growth, Denmark

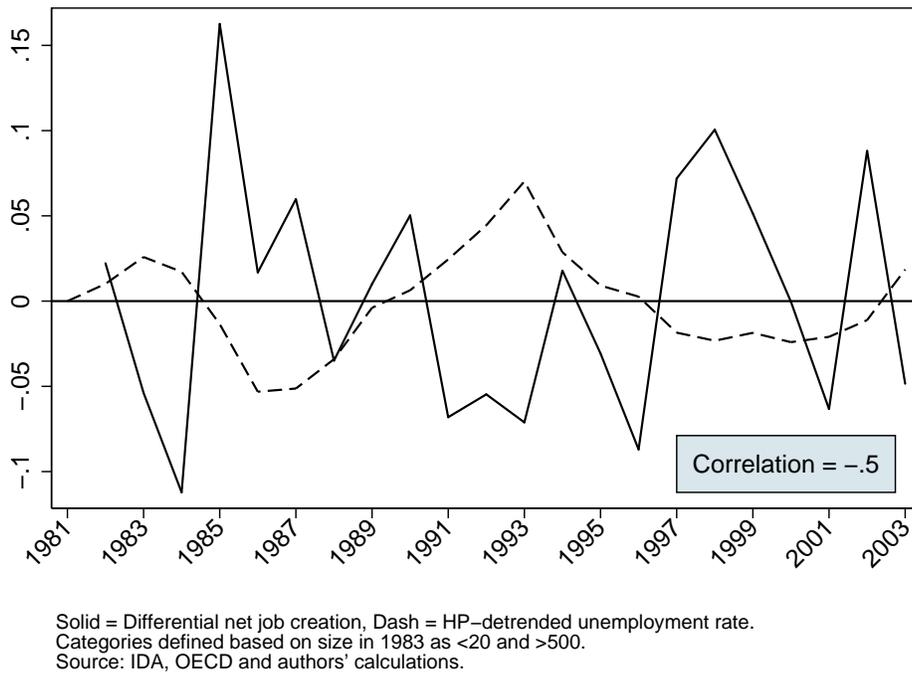


Fig. 11: Differential growth, Danish private companies classified by size in 1981

the entire Danish population aged 16 to 69.¹² 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 10, which plots for 1981-2003 the (detrended) Danish 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 -0.46 .

As an alternative, the Danish employer panel also enables us to fix the composition of size classes once and for all. In Figure 11 we allocate firms to the two size classes in 1981 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 through 1994, and then appears more blurred. This is not too surprising, and we do not take it as contradictory evidence, as firm size class membership is not perennial and significant reshuffling occurs over decades.

Our underlying theoretical framework considers firm size mostly as an indirect measure of productivity, based on the amply documented positive relationships between employment size and either revenue-based measures of productivity or wages. For IDA, 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 underlying productivity. It is plausible that wages reflect productivity better than size, especially for young firms

¹²See Bagger et al. (2009) for a detailed description of the IDA data set.

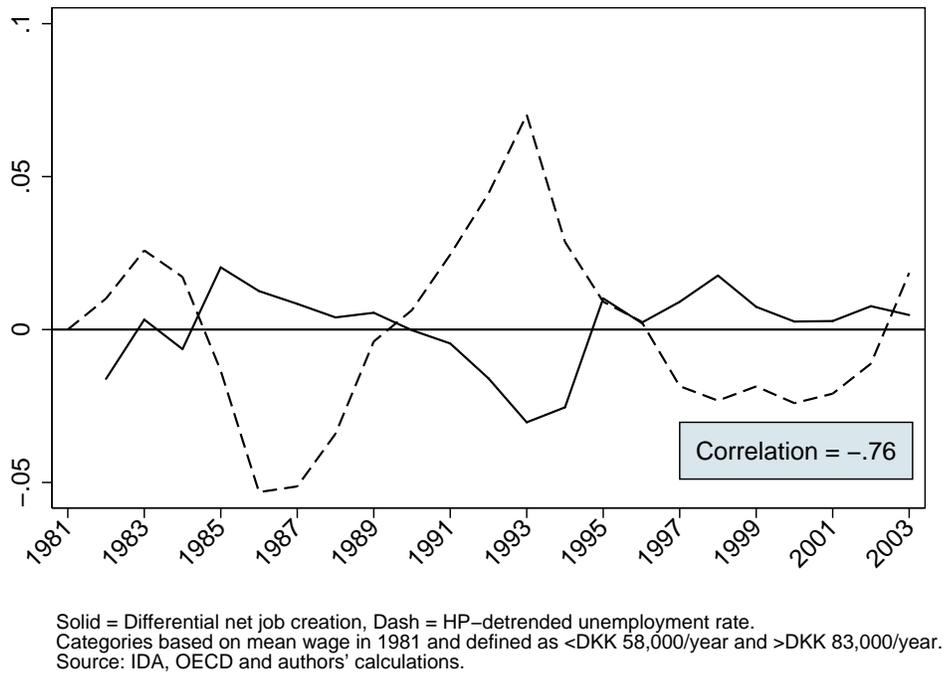


Fig. 12: Differential growth, Danish private companies classified by mean wage in 1981

that are still in their initial growth phase. In Figure 12 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).

6.1.2 France

The French Institute for National Statistics (INSEE) collects and maintains a panel of firm-level data, extracted from the *Bénéfices Réels Normaux* (BRN) register, and which conveys standard annual accounting information on all private companies (*not* establishments) with an annual sales turnover in excess of (roughly) €550,000 and liable to corporate taxes. From the BRN data set, we were able to access an exhaustive panel of private firms covering the

years 1985-2005 with information about end-of-year employment size (as of December 31st each year) and value added.¹³

By its nature and design, the French BRN panel is very similar to the Danish IDA firm file. We thus use it in exactly the same way: Figures 13, 14 and 15 are the French counterparts of the Danish Figures 10, 11 and 12, the only difference being that firm are classified using value added per worker (used here as an alternate measure of productivity) in Figure 15, as opposed to mean wages in the Danish case (Figure 12).¹⁴ As usual in all figures we also show a plot of the HP-filtered unemployment rate (taken from the OECD labor force data) to materialize the business cycle.

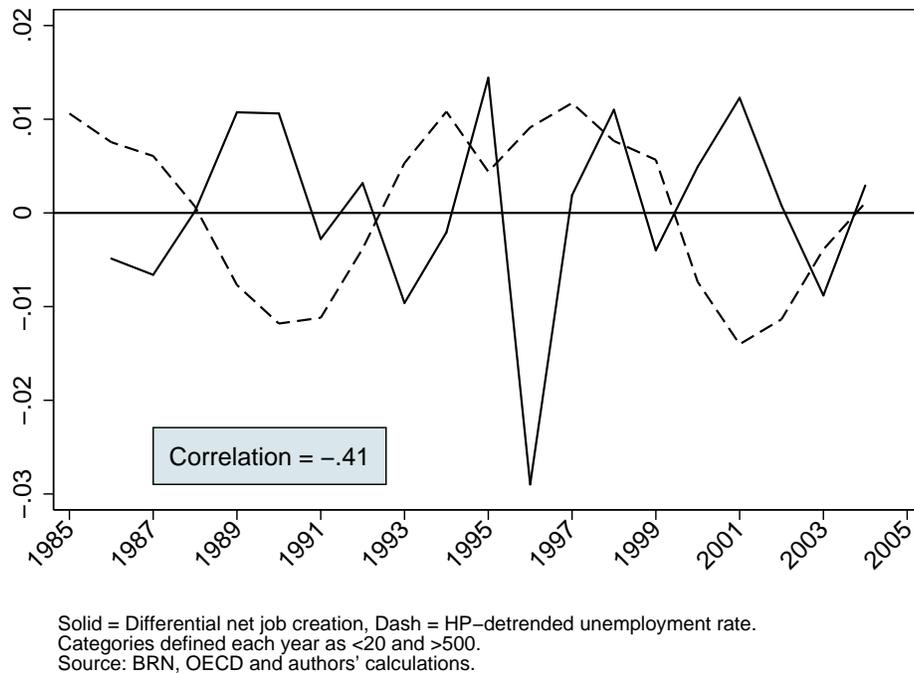


Fig. 13: Differential growth, France

The familiar cyclical pattern of relative growth emerges again from all three figures: initially large employers (in size or in value added per worker) grow faster when unemployment

¹³Access to BRN data is restricted to authorized researchers. We are grateful to Linas Tarasonis of University Paris I and CREST-INSEE for performing the data extraction and analysis for us.

¹⁴Specifically, firms in Figure 15 are classified once and for all into high- and low-productivity classes according to their mean value added per worker over the initial period 1985-88. This time-averaging is warranted by the notorious volatility of value added per worker.

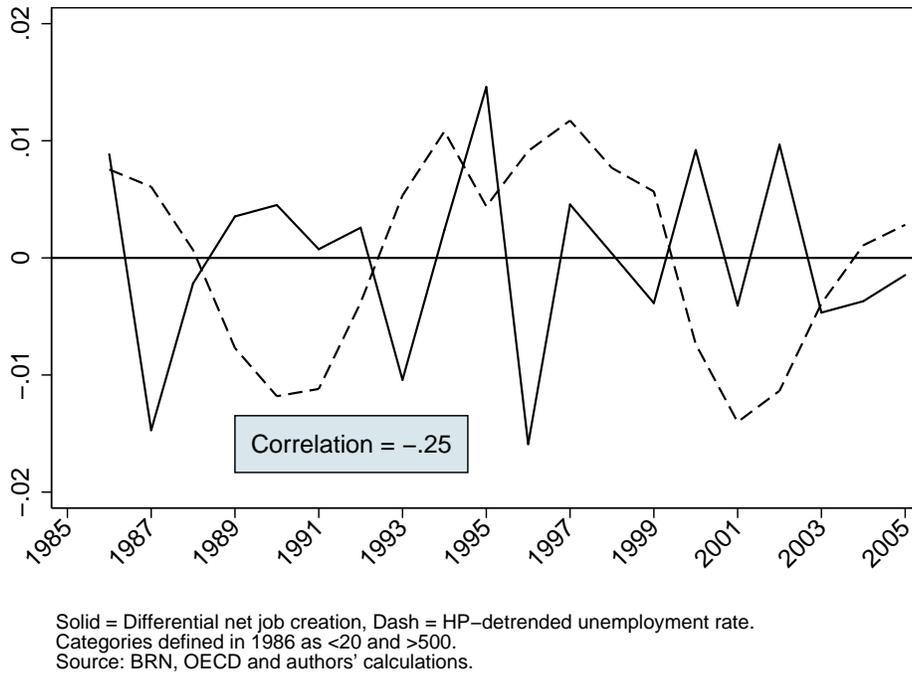


Fig. 14: Differential growth, French private companies classified by size in 1986

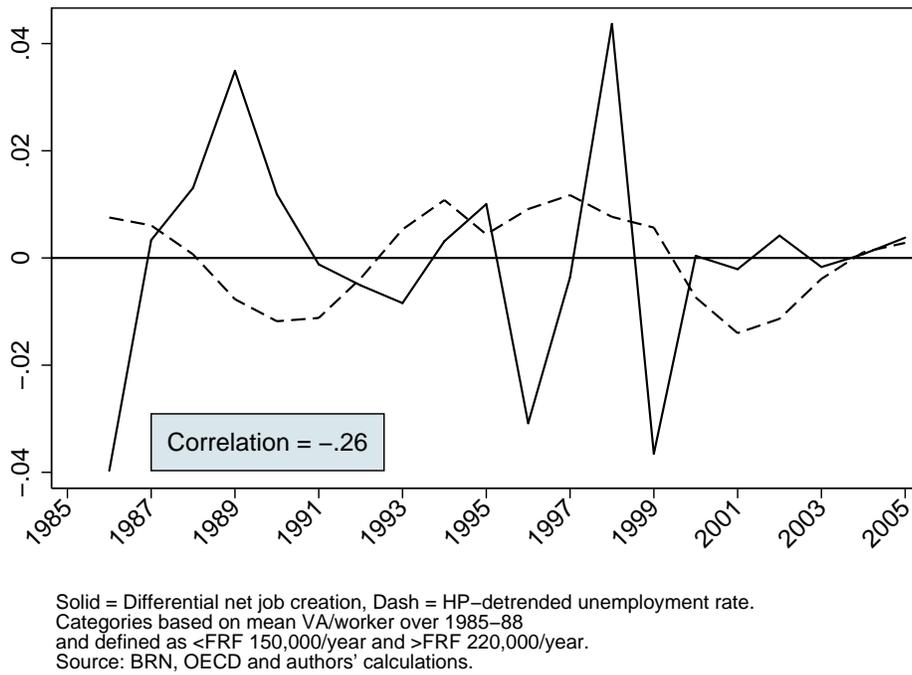


Fig. 15: Differential growth, French private companies classified by mean value added per worker in 1986

is unusually low, and vice versa. The only qualification to that statement is the presence of two spikes in all three relative growth rate series, in 1994-5 and 1997-8. Looking at average growth among small and large employers separately reveals that those spikes are mainly driven by abnormally strong growth of large firms in those years. Exactly what causes those spikes is difficult to determine. The problem of “spurious” job flows resulting from companies changing identifier for reasons that are completely orthogonal to their hiring and firing decisions is now known to be particularly severe in the French data (Picart, 2008).¹⁵ Yet while it undoubtedly adds noise to our relative growth series, we have found no compelling reason to believe that it was especially acute in those particular spike years. Whatever their cause, those two spikes are not enough to falsify the statement of Fact 5: the correlation between detrended unemployment and relative growth by size class remains decidedly negative on all three graphs.

6.1.3 Brazil

RAIS (*Relação Anual de Informações Sociais*) is an administrative, longitudinal dataset collected annually by the Brazilian labor ministry, which includes all firms in the Brazilian formal sector and provides information for all their workers such as age, education and sex, some information about establishments, such as sector and location, and 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.¹⁶

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 16 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 blow suffered by the Brazilian economy in 1998

¹⁵We thank Pierre Cahuc for bringing this recent study to our attention.

¹⁶Access 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.

as a consequence of the Asian crisis 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 phase, which works against Fact 5, reflects in part structural, rather than cyclical, factors, a downward trend in the relative size of the informal sector, which halts after the crisis.

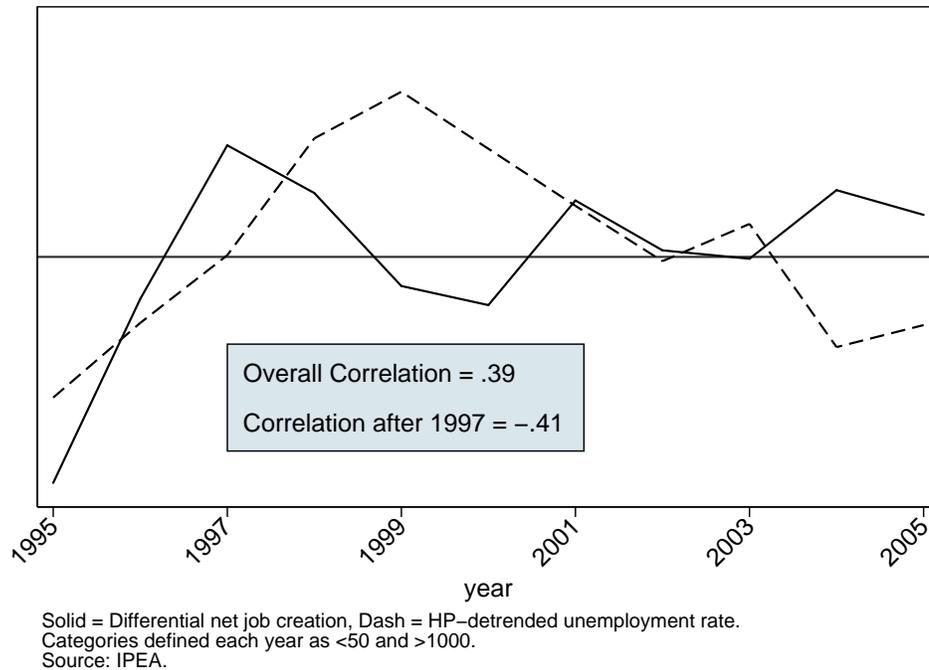


Fig. 16: Differential growth by firm size fixed in 1995, RAIS census of Brazilian employers

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 underlying pattern of growth by initial size, because the Reclassification Bias appears to be quantitatively very modest at annual or higher frequencies. We present evidence from two

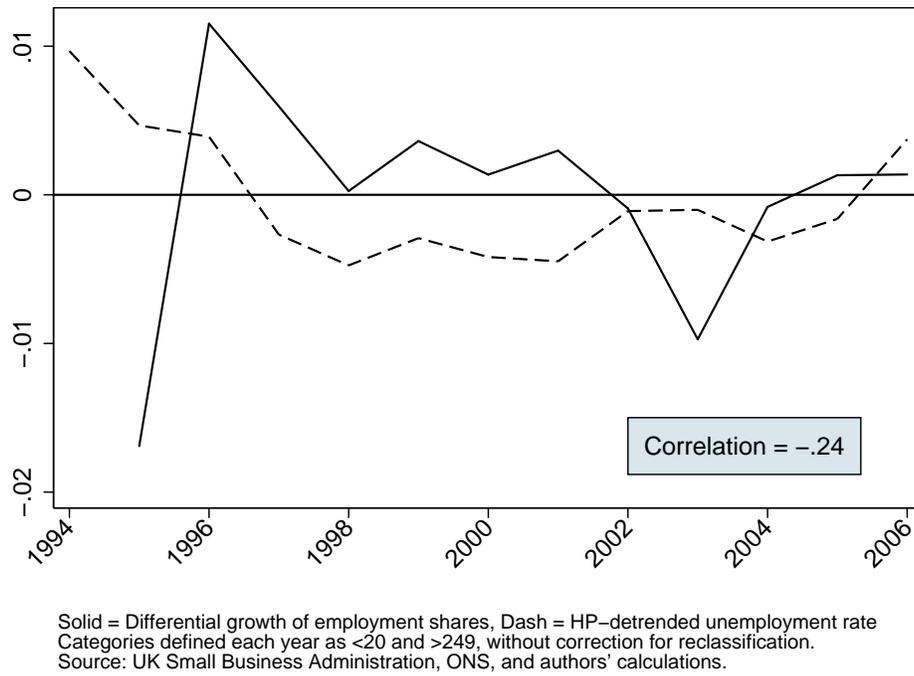


Fig. 17: Differential growth of employment shares, United Kingdom

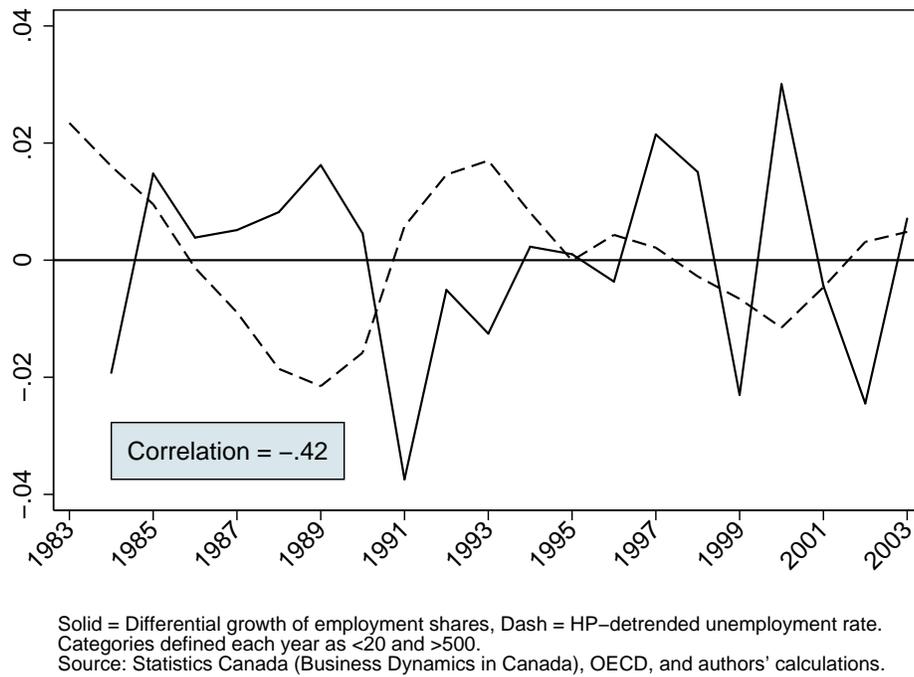


Fig. 18: Differential growth of employment shares, Canada

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 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 negative, as visually clear in Figure 17.

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 1983-2003.¹⁷ The largest size category available is 500. From these data, we can compute employment growth rates, with reclassification and including entrants. In Figure 18 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 -0.42 .

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.

¹⁷Visit www.statcan.gc.ca/bsolc/olc-cel/olc-cel?catno=61-534-XWE&lang=eng for details.

Let u_t denote the unemployment rate at (continuous) time t . Let ω denote an aggregate state of the economy, which changes over time according to a Poisson process. We model transitions in and out of unemployment as a standard “bathtub” model. Assume constant labor force participation. Then the evolution of the unemployment rate is described by:

$$\dot{u}_t = \delta^\omega (1 - u_t) - \lambda_0^\omega u_t.$$

The inflow into unemployment, equal to the separation rate δ^ω times the employment rate $1 - u_t$, and the outflow, equal to the job-finding rate λ_0^ω times the unemployment rate, determine the change in the unemployment rate in state ω .

Next, we model job-to-job transitions as a “firm ladder”. Workers can search randomly both off and on the job and receive offers at rates λ_0^ω and λ_1^ω respectively. When unemployed, a worker accepts any offer. 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 firm A over firm B, so do all other workers, and any A employee will reject offers from B and any B employee will accept offers from A. Additionally, some workers lose jobs from all rungs of the firm ladder and become unemployed, as in the bathtub model, and must restart climbing the ladder from the bottom rung.

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

$$\dot{L}_t(\theta) = -\{\delta^\omega + \lambda_1^\omega [1 - Q(\theta)]\} L_t(\theta) + q(\theta) \left[\lambda_0^\omega u_t + \lambda_1^\omega \int_0^\theta L_t(x) dx \right] \quad (2)$$

The outflow occurs at rate δ^ω to unemployment and at rate λ_1^ω , the contact rate on the job, times $1 - Q(\theta)$, the probability of sampling a higher-ranked firm, to other employers. The inflow occurs at rate λ_0^ω from unemployment and at rate λ_1^ω 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” $q(\theta)$.

Equation (2) 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 $u_0 = 1 - \int_0^1 L_0(\theta) d\theta$. At every countable time when the state ω changes value stochastically to some ω' , the new turnover rates $\delta^{\omega'}$ and $\lambda_i^{\omega'}$ apply until the next state switch. This PDE has the solution

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

If ω remains constant, this converges to an ergodic stationary distribution

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

The normalized steady state cdf of employment across firms can then be written

$$\frac{\int_0^\theta L_\infty(x) dx}{\int_0^1 L_\infty(x) dx} = \frac{\delta^\omega + \lambda_1^\omega [1 - Q(1)]}{\delta^\omega + \lambda_1^\omega [1 - Q(\theta)]}$$

As Q is an increasing function, it is straightforward to show that this cdf is increasing in δ^ω and decreasing in λ_1^ω for every θ in $[0,1]$.

Now consider two realizations of the aggregate state, ω and ω' , such that the job offer rates λ_i^ω are low and the separation rate δ^ω is high in state ω , and vice versa in state ω' . So ω corresponds to a depressed labor market, as in a typical recession, and ω' to a tight labor market, as in a mature expansion. Then, the normalized cdf of employment across firms in the steady state with constant state ω' first-order stochastically dominates (FSD) that in a steady state ω . It is then natural to expect that in the stochastic economy when the aggregate state fluctuates randomly between ω and ω' , a recession and a boom, the distribution of employment FSD-shifts up and down: in a recession, while total employment declines, the low-ranked (low θ) firms shrink less fast, while high-ranked firms grow faster in a boom, when employment climbs the job ladder and moves faster towards them.

The natural application to our setting is as follows. Let θ denote a firm-specific productivity parameter, fixed over time. Then, the assumed worker preference ordering over employers implements efficient turnover: given the opportunity, workers always move from less to more productive employers. Next, assume that $L_0(\theta)$ is increasing in θ . That is, more productive firms start out bigger. Under efficient turnover, this initial size ranking is never reversed. In MPV08 we calibrate the parameters of this PDE ($\lambda_0^\omega, \lambda_1^\omega, \delta^\omega, L_0$ and Q) to match stylized facts about worker turnover (to/from unemployment and job-to-job) and the average distribution of employment by size (L_0), but not the evolution of firm sizes. When we compute the resulting path of $L_t(\theta)$, just 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, and plunges from peak to trough.

The intuition behind these robust distributional dynamics emerges from carefully inspecting Equation (2). All firms have access to an inflow from unemployment that is independent of firm size, although it may depend on firm productivity through sampling weights Q . As the economy expands, this inflow declines with unemployment. 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. So the first stage of an expansion “loads” small firms with employment; later, when hiring from unemployment becomes difficult, larger firms tap this reservoir.

In MPV09 we show that this efficient turnover is implemented by the unique equilibrium of a dynamic 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 produces the efficient turnover dynamic described above. In our quantitative exercise, the 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. 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.

8 Concluding Remarks

We present abundant empirical 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 U.S. datasets, BDS and BED, cannot be extended back in time given the availability of underlying micro data. Repeated cross-sections like County Business Patterns are available in archival form for the post-war period, but with many gaps before the late 1960s. Alternative sources that extend over longer periods cover only specific sectors, mostly manufacturing. 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, France, 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 will 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] Bachmann, R. and P. David (2009), “The Importance of Two-Sided Heterogeneity for the Cyclicalities of Labour Market Dynamics”, RWI Essen *mimeo*.
- [2] Bagger, J., H. Bunzel, B. J. Christensen and N. Pedersen (2009), “Danish Matched Employer-Employee Data: Patterns and Puzzles”, Aarhus University *mimeo*.
- [3] Bils, M. and K. McLaughlin (2001), “Inter-Industry Mobility and the Cyclical Upgrading of Labor”, *Journal of Labor Economics*, 19(1), 94-135.
- [4] Bowler, M. and T. Morisi (2006), “Understanding the employment measures from the CPS and CES survey”, *Monthly Labor Review*, February, 23-38.
- [5] Chari, V.V., L. Christiano and P. Kehoe (2007), “The Gertler-Gilchrist Evidence on Small and Large Firm Sales”, *mimeo*.

- [6] 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.
- [7] 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.
- [8] Davis, S. J., J. C. Haltiwanger and S. Schuh (1996), *Job Creation and Destruction*, Cambridge, MA: The MIT Press.
- [9] 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.
- [10] Foster, L., J. Haltiwanger and C. Syverson (2008), “Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability?”, *American Economic Review*.
- [11] 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. [GG]
- [12] 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.
- [13] Jovanovic, B. (1982), “Selection and the Evolution of Industry”, *Econometrica* 50(3), 649-670.
- [14] Kanagarajah, S. (2003), “Business Dynamics in Canada”, Statistics Canada: Catalogue No. 61-534-XIE.
- [15] King, Wilford Isbell (1923), *Employment Hours and Earnings in Prosperity and Depression, United States, 1920-1922* NBER, New York.

- [16] Lentz, R. and D. T. Mortensen (2008), “An Empirical Model of Growth through Product Innovation”, *Econometrica* 76(6), 1317-73.
- [17] 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.
- [18] Moscarini, G. and F. Postel-Vinay (2008), “The Timing of Labor Market Expansions: New Facts and a New Hypothesis”, in D. Acemoglu, K. Rogoff and M. Woodford., eds, *NBER Macroeconomics Annual*, 23, 1-51. [MPV08]
- [19] Moscarini, G. and F. Postel-Vinay (2009), “Nonstationary Search Equilibrium”, Yale University and University of Bristol *mimeo*. [MPV09]
- [20] Okun, A. (1973), “Upward Mobility in a High Pressure Economy”, *Brooking Papers on Economic Activity*, 207-52.
- [21] Picart, C. (2008), “Flux d’emploi et de main-d’œuvre en France: un réexamen”, *Economie et Statistique*, 412, 27-56.
- [22] 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. [RR]
- [23] Sharpe, S. (1994), “Financial Market Imperfections, Firm Leverage, and the Cyclicity of Employment”, *American Economic Review*, 84(4), 1060-1074.
- [24] Shimer, R. (2005), “Reassessing the Ins and Outs of Unemployment”, University of Chicago *mimeo*.
- [25] Sutton, J. (1997), “Gibrat’s Legacy,” *Journal of Economic Literature*, 35(1), 40-59.