Volume Title: Measuring Entrepreneurial Businesses: Current Knowledge and Challenges

Volume Author/Editor: John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-45407-8 (cloth); 978-0-226-45410-8 (e-ISBN)

Volume URL: http://www.nber.org/books/halt14-1

Conference Date: December 16-17, 2014

Publication Date: September 2017

Chapter Title: High Growth Young Firms: Contribution to Job, Output, and Productivity Growth

Chapter Author(s): John Haltiwanger, Ron S. Jarmin, Robert Kulick, Javier Miranda

Chapter URL: http://www.nber.org/chapters/c13492

Chapter pages in book: (p. 11 – 62)

1

High-Growth Young Firms
Contribution to Job, Output, and Productivity Growth

John Haltiwanger, Ron S. Jarmin, Robert Kulick, and Javier Miranda

1.1 Introduction

Business start-ups and high-growth young firms disproportionately contribute to job creation in the United States. In a typical year, start-ups account for about 10 percent of firms and more than 20 percent of firm-level gross job creation. Less well known is that most US business start-ups exit within the first ten years, and the median surviving young business does not create jobs but remains small. A small fraction of young firms create jobs rapidly and contribute substantially to job creation. These high-growth young firms are the reason that start-ups make a long-lasting contribution to net job creation.[1]

Most of the limited evidence on high-growth firms has been about their contribution to job creation. Less is known about the nature of their contribution to output and productivity growth due primarily to data limitations.

John Haltiwanger is a Distinguished University Professor of Economics at the University of Maryland and a research associate of the National Bureau of Economic Research. Ron S Jarmin is the Assistant Director for Research and Methodology at the US Census Bureau. Robert Kulick is a PhD candidate in economics at the University of Maryland. Javier Miranda is principal economist in the Center for Economic Studies, US Census Bureau.

We thank David Brown, Tim Dunne, participants of the NBER/CRIW Conference on Measuring Entrepreneurial Businesses, and participants of the 2015 ASSA meetings for helpful comments. We thank the Kauffman Foundation for financial support. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the US Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see http://www.nber.org/chapters/c13492.ack.

1. This discussion is based on Haltiwanger, Jarmin, and Miranda (2013) and Decker et al. (2014). Note that the statistic that start-ups account for more than 20 percent of firm-level job creation is based on gross job creation by firms, not establishments. Start-ups account for slightly less than 20 percent of establishment-level job creation.

For the United States, substantial progress has been made in developing longitudinal business databases that permit tracking growth and survival of businesses in terms of jobs. Studies of the role of business dynamics in output and productivity growth are largely limited to the manufacturing sector with some limited analysis of the retail trade sector where the data are most suitable.

In this chapter, we describe our efforts to extend the data infrastructure on business dynamics to permit tracking real output and labor productivity growth at the firm level for the entire US private sector on an annual basis. To our knowledge, this is the first database at the firm level that tracks both output and employment outcomes for all types of firms in the private sector on an annual basis.[2] This enables us to study the contribution of young high-growth firms to real output and productivity growth (i.e., real output per worker).

High-growth firms are part of the ongoing dynamics of real output and input reallocation that characterize economic growth in the United States and other market economies. Since at least the work of Dunne, Roberts, and Samuelson (1989) and Davis and Haltiwanger (1990, 1992), we have known that underlying net growth in the United States is a high pace of job reallocation. Early work focused on decomposing net employment growth into gross job creation and destruction. More recent work has shown that there is a high pace of real output and capital reallocation that accompanies the employment reallocation (see, e.g., Foster, Haltiwanger, and Krizan 2001; Becker et al. 2006), at least for selected sectors. One of the earliest findings in this literature is that young businesses exhibit a high pace of reallocation relative to more mature businesses. A second key finding in the early literature is that most of the job reallocation reflects reallocation within industry. While early work focused on US manufacturing, recent work has extended the analysis to the entire US private sector (e.g., Haltiwanger, Jarmin, and Miranda 2013; Decker et al. 2014).[3]

The high pace of within-industry reallocation has been interpreted through the predictions of the canonical firm dynamics models of Jovanovic (1982), Hopenhayn (1992), and Ericson and Pakes (1995), among others. In these models and in the subsequent literature, firms in the same industry differ in their productivity and the reallocation dynamics reflect moving resources away from less productive to more productive businesses. Such productivity differences can be endogenous given the role of endogenous innovation and research and development (R&D) activities. Entrants and

2. For publicly traded firms, COMPUSTAT provides a rich source of output, asset, and other data. The quinquennial economic censuses can be used to provide output data for most sectors every five years. Annual surveys of specific sectors can be used to generate samples of firms for most sectors but they are less well suited for longitudinal analysis at the firm level.

3. Hereafter we often refer to these as HJM (2013) and DHJM (2014).

young businesses play a critical role in these dynamics. They put competitive pressure on incumbents, and in some models they are critical for innovation (see, e.g., Acemoglu et al. 2013).

The high pace of real output and input reallocation of young businesses is interpreted as part of the learning and selection dynamics as well as the endogenous innovation dynamics that are present in this class of models. Jovanovic (1982) argues that entering firms initially do not know their type, but learn about it over time. In that model, high-growth young firms are those that learn that they are high productivity or high demand. In contrast, high-decline young firms are those that learn that they are low productivity or demand. Ericson and Pakes (1995) extended these learning ideas to environments where all firms engaging in some new form of activity have to learn whether they are profitable in that activity. Moreover, with endogenous innovation such as in Acemoglu et al. (2013), productivity evolves based on the amount and success of innovative activity. In these models with more active learning and endogenous innovation, high-growth young firms are those that innovate and learn successfully.

While some theoretical models highlight the potentially critical role of high-growth young firms to growth, it is increasingly understood that the contribution of high-growth young firms is likely to be much more important in some sectors than others. For example, the recent work of Hurst and Pugsley (2012) highlights the heterogeneity in the motivation for starting a business and hence their potential growth. They point to sectors dominated by small businesses that reflect occupational and lifestyle choices of business owners (such as wanting to be their own boss) rather than an entrepreneurial desire to innovate and grow. In such sectors it may be the case that high-growth firms do not play a significant role in contributing to job creation and productivity growth.

Most previous efforts to analyze the role of high-growth firms focused only on one single dimension of growth—employment. We create a revenue-enhanced version of the Census Bureau's Longitudinal Business Database that has been the workhorse of much research on firm dynamics. These data permit us to examine high-growth firms along both the employment and output dimensions, as well as to examine their role in productivity growth as in the models discussed above.

We find that the patterns for high-growth-output firms largely mimic those for high-employment-growth firms. High-growth-output firms are disproportionately young and these firms make outsized contributions to output and productivity growth. The share of activity accounted for by high-growth-output- and employment firms varies substantially across industries—in the post-2000 period the share of activity accounted for by high-growth firms is significantly higher in the high-tech- and energy-related (for the latter, the share of output) industries. A firm in a small business-intensive industry is less likely to be a high-growth-output firm, but small business-intensive industries do not have significantly smaller shares of activity accounted for by high-growth firms for either output or employment.

The chapter proceeds as follows. Section 1.2 presents a description of the data developed and used in this chapter. Section 1.3 presents our main empirical findings. Our findings are mostly descriptive findings about the joint distribution of employment, real output, and productivity growth. Given our interest in entrepreneurship, in section 1.4 we focus considerable attention on the role of young firms in these dynamics. Concluding remarks that summarize our main findings and discuss next steps are in section 1.5.

1.2 Business Dynamics Data

We use two core-related databases in this chapter. Both are based on the Census Business Register (BR). We use the Census Bureau's Longitudinal Business Database (LBD) to construct measures of firm employment growth and firm age. We then append to these core business dynamics data firm-level revenue data contained in the BR and sourced from administrative records. First, we discuss the basic LBD data and then describe our work to enhance the LBD with revenue information.

1.2.1 Business Dynamics Measurement with the LBD

Like the BR, the LBD covers the universe of establishments and firms in the US nonfarm business sector with at least one paid employee. The LBD includes annual observations beginning in 1976 and currently runs through 2013. It provides information on detailed industry, location, employment, and parent firm affiliation for every establishment. Employment observations in the LBD are for the payroll period covering the 12th day of March in each calendar year. The LBD's high-quality longitudinal establishment and firm-ownership information make possible the construction of our measures of firm growth and firm age. In what follows, we first discuss the key features of the LBD and then return to discussing the data we use from the BR to measure real output.

A unique advantage of the LBD is its comprehensive coverage of both firms and establishments. Firm activity is captured in the LBD up to the level of operational control instead of being based on an arbitrary taxpayer ID.[4] The ability to link establishment and firm information allows firm characteristics such as firm size and firm age to be tracked for each establishment.

4. A closely related database at the BLS tracks quarterly job creation and destruction statistics (Business Employment Dynamics). The BED has advantages in terms of both frequency and timeliness of the data. However, the BED only can capture firm dynamics up to the level of establishments that operate under a common taxpayer ID (EIN). There are many large firms that have multiple EINs—it is not unusual for large firms operating in multiple states to have at least one EIN per state.

Firm-size measures are constructed by aggregating the establishment information to the firm level using the appropriate firm identifiers. The construction of firm age follows the approach adopted for the BDS and based on our prior work (see, e.g., Becker et al. 2006; Davis et al. 2007; Haltiwanger, Jarmin, and Miranda 2013). Namely, when a new firm ID arises for whatever reason, we assign the firm an age based on the age of the oldest establishment that the firm owns in the first year in which the new firm ID is observed. The firm is then allowed to age naturally (by one year for each additional year it is observed in the data), regardless of any acquisitions and divestitures as long as the firm continues operations as a legal entity. This permits defining start-ups as new firms with all new establishments and shutdowns as firms that cease operations and all establishments shut down.

We utilize the LBD to construct annual establishment-level and firm-level employment growth rates. The measures we construct abstract from net growth at the firm level due to mergers and acquisitions (M&A) activity. We use Davis, Haltiwanger, and Schuh (1996) net growth-rate measures that accommodate entry and exit.[5] We refer to this as the DHS growth rate.

Computing establishment-level growth rates is straightforward, but computing firm-level growth rates is more complex given changes in ownership due to mergers, divestitures, or acquisitions. In these instances, net growth rates computed from firm-level data alone will reflect changes in firm employment due to adding and/or shedding continuing establishments. This occurs even if the added and/or shed establishments experience no employment changes themselves. To avoid firm growth rates capturing changes due to M&A and organization change, we compute the period $t - 1$ to period t net growth rate for a firm as the sum of the appropriately weighted DHS net growth rate of all establishments owned by the firm in period t, including acquisitions, plus the net growth attributed to establishments owned by the firm in period $t - 1$ that it has closed before period t. For any continuing establishment that changes ownership, this method attributes any net employment growth to the acquiring firm. Note, however, if the acquired establishment exhibits no change in employment, there will be no accompanying change in firm-level employment induced by this ownership change. The general point is that this method for computing firm-level growth captures only "organic" growth at the establishment level and abstracts from changes in firm-level employment due to M&A activity (see supplementary data appendix to Haltiwanger, Jardin, and Miranda [2013] for an example).

The LBD permits us to characterize the comprehensive distribution of firm employment growth rates including the contribution from firm entry, firm exit, and continuing firms.[6] We begin our analysis with the LBD to

5. This growth-rate measure has become standard in analysis of establishment and firm dynamics because it shares some useful properties of log differences, but also accommodates entry and exit (See Davis, Haltiwanger, and Schuh 1996; Törnqvist, Vartia, and Vartia 1985).

6. By continuing firms we mean firms that continue between $t - 1$ and t.

characterize the distribution of firm net employment growth rates for both continuing and exiting firms. Much of our analysis focuses on firms that are age 1 and older so that we do not focus on start-ups in their first year. Our recent work (see Haltiwanger, Jarmin, and Miranda 2013) highlights the contribution of start-ups to job creation in their first year. As we noted in the introduction, start-ups account for slightly more than 20 percent of firm-level gross job creation (and slightly less than 20 percent of establishment-level job creation). The focus of the current chapter is postentry dynamics.

1.2.2 Enhancing the LBD with Firm-Level Measures of Revenue

A key innovation of this chapter is that we introduce real output and productivity-growth measures to the analysis of high-growth firms. Our measure of output is a gross-output measure derived from revenue data from the Census Bureau's Business Register (BR), which also provides the source data for the LBD. The BR's revenue measure is based on administrative data from annual business income tax returns. Unlike payroll and employment, which are measured at the establishment level going back to 1976, the nominal output data are available at the tax reporting or employer identification number (EIN) level only and then only starting in the mid-1990s. The tax reporting unit is equivalent to a particular physical location (an establishment) only in the case of single-unit firms. In the case of multiunit firms, the administrative data does not apportion output to particular establishments. Thus, in the BR, revenue is only measured at the establishment level for single-location firms. Constructing a comprehensive revenue measure is further complicated by the fact that the content of the receipts fields on the BR vary substantially by type of activity and the legal structure of the firm according to different tax treatments.

For sole proprietorships, business income taxes are filed on the business owner's individual income taxes. Administrative data enable linking these individual income tax returns to the payroll EINs for sole proprietors, but these links are imperfect (see Davis et al. 2009). Corporations and partnerships file their business income taxes with an EIN but a challenge is that firms may have multiple EINs. Information from the Economic Censuses, Company Organization Survey, and administrative records are used to develop high-quality links between all the payroll EINs of a firm and the parent firm ID. This implies that for most corporations and partnerships, we link the business income tax EIN to one of the payroll EINs. Given the links of the payroll EINs to the parent firm identifier, this enables us to construct a consistent measure of employment and output at the firm level. However, multiple EIN firms are not required to report income using the same EIN they use to report quarterly payroll. As a result, income EINs can become "detached" from their payroll EINs. We discuss these issues in more detail in appendix A, but overall we successfully added nominal revenue measures to over 80 percent of the firm records in the LBD in our sample period. We denote this as the revenue enhanced subset of the LBD.

We find that the pattern of missingness of revenue is only weakly related to observable indicators in the full LBD like firm age, firm size, broad industry, the employment growth rate, or multiunit status. Consistent with this finding, the relationship between the distribution of firm employment growth rates and firm age for the revenue enhanced subset of the LBD and the full LBD are very similar. However, to mitigate possible selection issues we weight our subset data with inverse propensity score weights (IPW). These weights are based on estimation of propensity score models separately for continuers, deaths, and births from the full LBD. The propensity score models use logistic regressions with the dependent variable equal to one if the firm has revenue and zero otherwise. Observable firm characteristics from the full LBD used in the models include firm size, firm age, employment growth rate, industry, and a multiunit status indicator. The propensity score-weighted data yields patterns of employment growth rates, employment-weighted entry, and employment-weighted exit that are quite similar to those obtained from the full population of continuers, entrants, and exiters. Additional details are provided in the data appendix.[7]

We deflate the nominal revenue measures with a general price deflator (the GDP Implicit Price Deflator). As such, our measures of real gross output will reflect both real output changes and changes in relative prices across industries. Revenue fields in the BR can be noisy so we adopt filters to clean out unreasonable values. These filters are discussed further in the data appendix and include minimum and maximum productivity value cutoffs, maximum revenue cutoffs, and maximum revenue-growth values. Subsequent references to output in what follows should be interpreted as real revenue or equivalently real gross output.

A limitation of our real gross output measure is that it does not capture the contribution of intermediate inputs. In many of our exercises, we control for interacted industry and year effects. Doing so effectively controls for industry-specific deflators. Moreover, this also acts as a control for industry-specific variation in intermediate input shares.[8] Controls for industry and year effects is especially important when we examine labor productivity, since cross-industry variation in gross output per worker are difficult to interpret. We also note that for output growth we use DHS measures of growth. Another limitation of our output growth measures is that since we do not have the underlying establishment-level output growth we cannot abstract from the contribution of M&A activity to output growth.

7. We note that we exclude 2001 and 2002 from our statistics since the 2001 data are problematic (which impacts the growth-rate distributions in both 2001 and 2002).

8. Most of our analysis focuses on the distribution of growth rates of gross output. Growth rates abstract from any industry-level differences in gross output from differential intermediate input shares.

The filters we design partly take care of this as M&A activity can lead to spurious large flows of output. We have checked and found that the broad patterns we find for employment growth largely hold when we do not adjust for M&A growth—but still we regard this as a limitation that should be acknowledged (and also as an area for future research).

1.3 The Role of High-Growth Firms for Job Creation, Real Output Growth, and Productivity Growth

1.3.1 The Up or Out Dynamics of Young Firms in the United States

Employment Dynamics

We begin by comparing results we obtain with the output-enhanced subset of the LBD with prior findings from HJM and DHJM that make use of the full LBD. Those papers emphasized two features of the employment-growth dynamics of young firms in the United States: (a) the up or out dynamic of young firms, and (b) differential patterns of dispersion and skewness of firm-growth distribution by firm age.

As highlighted in HJM, decomposing overall net growth into the net growth from continuers and the contribution from exit reveals the up or out pattern of young firms. Figures 1.1A and 1.1B show the net employment growth rate for surviving firms as well as the job destruction rate from firm exit by firm age. Figure 1.1A shows results from the full LBD and figure 1.1B from the output-enhanced subset adjusted using inverse propensity score weights. We exclude years not covered by the output-enhanced subset.[9] Firm exit is defined as discussed above. All statistics are employment weighted. Figures 1.1A and 1.1B focus on the postentry dynamics of firms; in our nomenclature, age one is the year after entry. We exclude entrants in these figures since age zero businesses only create jobs in their year of entry.[10] The weighted sum of net job creation yields overall net employment growth for a given age group.[11] Conditional on survival, young firms have much higher growth rates than more mature firms. Young firms also have a substantially

9. In particular, the statistics are based on tabulations of pooled data from 1996 to 2013 from the Longitudinal Business Database (LBD), excluding the 2001 and 2002 years. We exclude those years here since the output data for 2001 has been partially lost. As we discuss below, the focus on the 1996–2013 period implies that our statistics are influenced by the Great Recession.

10. See HJM (2013) and DHJM (2014) for an extensive analysis of the contribution of start-ups to job creation. We have noted their average contribution. Those papers highlight that there has been a declining pace of entry in the United States. They also note that entry rates vary substantially across sectors and geographic regions. But interestingly, the papers note that even with variation in the entry rates the postentry dynamics are similar across sectors in terms of up or out dynamics.

11. Overall net growth is the sum of the weighted net-growth rate for continuers plus job destruction from exit. The weight is the share of employment for continuing firms. See HJM (2013) for details.

Fig. 1.1A Up or out dynamics of firms, 1996–2013, LBD

Source: Statistics computed from the Longitudinal Business Database and the revenue-enhanced subsets 1996–2000 and 2003–2013.

Notes: Figures 1.1A and 1.1B show patterns of net employment growth for continuing firms and job destruction from firm exit for firms age one and older.

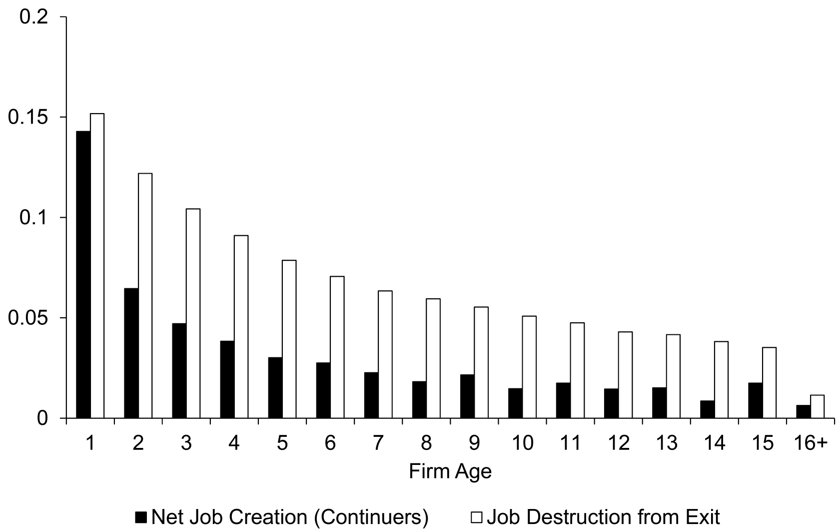

Fig. 1.1B Up or out dynamics of firms, 1996–2013, revenue-enhanced subset of LBD

Source: Statistics computed from the Longitudinal Business Database and the revenue-enhanced subsets 1996–2000 and 2003–2013.

Notes: Figures 1.1A and 1.1B show patterns of net employment growth for continuing firms and job destruction from firm exit for firms age one and older.

higher (employment-weighted) exit rate than more mature firms. Slightly over 50 percent of an entering cohort of firms in figure 1.1A will have exited by age five (on an employment-weighted basis). The very high failure rate of young firms is partially offset by the contribution of the surviving firms. For the sample period in figure 1.1A, five years after the entry of an average cohort, the employment is about 70 percent of the original contribution of the cohort. This is in spite of losing over 50 percent of employment to business exits.[12] Figure 1.1B shows very similar patterns for our propensity score-weighted, revenue-enhanced subset of the LBD.

Figures 1.2A and 1.2B examine job creation from firm births by size class for both the LBD population and the revenue-enhanced subset. The job-creation rate from births is particularly high among the smallest firms and decreases monotonically with the firm size. Patterns are again very similar across figures 1.2A and 1.2B.

One implication of figures 1.1A and 1.1B is that the overall net employment growth rate is negative for all firm age groups for age greater than firm age equal to zero. This pattern is evident from the job destruction rate from exit exceeding the net employment growth rate for continuing firms for all firm age groups. This pattern partly reflects our sample period, which includes the sharp contraction and slow recovery of 2007–11. But it also reflects the more general pattern that even in a typical year of overall positive net growth, continuing firms tend to be mildly contracting on average with overall (economy-wide) net employment growth being positive because of the contribution of firm start-ups (depicted in figures 1.2A and 1.2B). HJM show that this pattern holds for the sample period 1992–2005.[13] A related implication of figures 1.1A and 1.1B is that overall net employment growth rates are increasing with firm age.[14] Again, this partly reflects our sample period since young firms were hit especially hard in the Great Recession (see Fort et al. 2013), but is also a common pattern more generally (see figure 4 of HJM).

The second finding, highlighted in DHJM, highlights the dispersion and skewness of the employment growth rate distribution of continuing young firms. Figures 1.3A and 1.3B show the 90th, 50th (median), and 10th percentiles of the net job-growth distribution of surviving firms by firm age. As before, figures 1.3A and 1.3B show the LBD population and the revenue-enhanced subset, respectively. Percentiles are from the

12. These calculations of the five-year contribution of each cohort are low relative to those reported in HJM (2013) or in DHJM (2014). These differences reflect differences in sample periods and, in particular, whether the years of the Great Recession are included. HJM (2013) use the period 1992–2005. They find that for five years after the entry of an average cohort, the employment is about 84 percent of the original cohort. DHJM (2014) use the period 1992–2011 and find the same calculation yields 80 percent.

13. The BDS shows that in the years of most robust net growth, both very young and very old firms tend to have positive overall net growth inclusive of the contribution of exit.

14. This can be inferred by computing the overall net growth implied by figure 1.1.

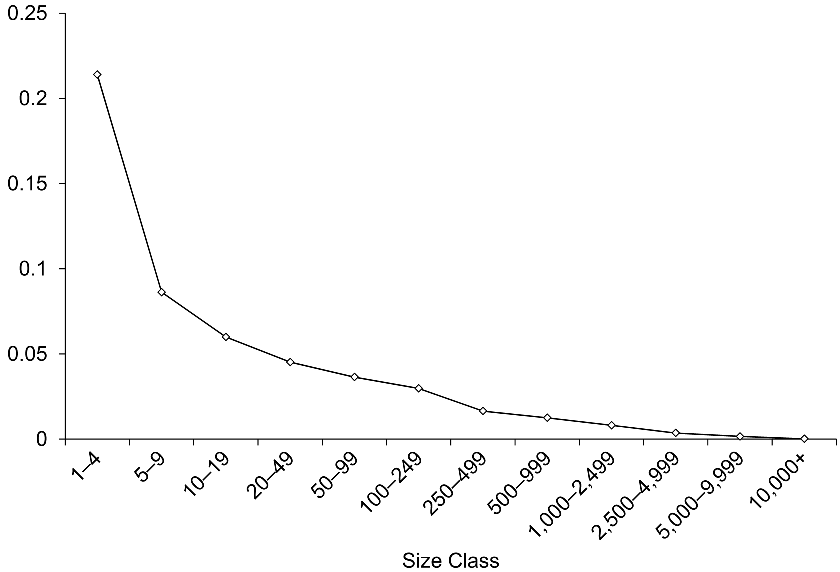

Fig. 1.2A Job creation from births, 1996–2013, LBD

Source: Statistics computed from the Longitudinal Business Database and the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: Figure 1.2A shows the pattern of job creation from firm births by size class.

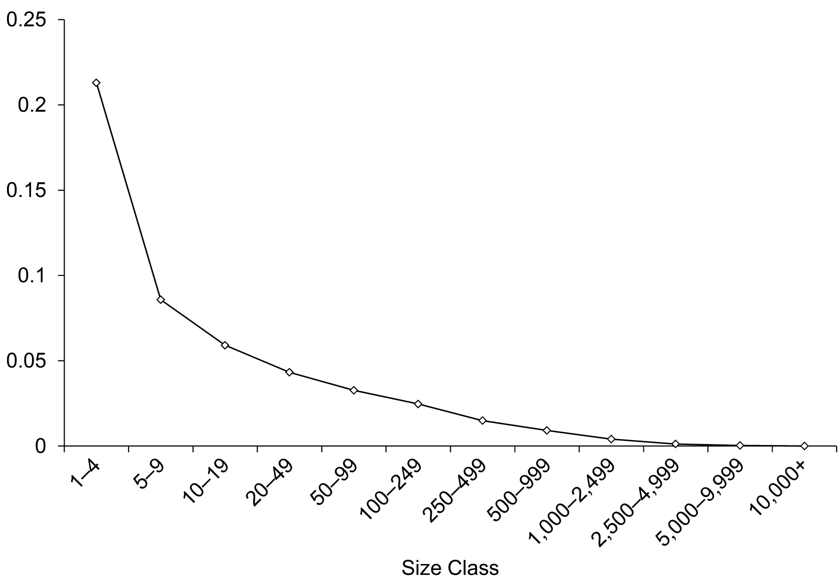

Fig. 1.2B Job creation from births, 1996–2013, revenue-enhanced subset of LBD

Source: Statistics computed from the Longitudinal Business Database and the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: Figure 1.2B shows the pattern of job creation from firm births by size class.

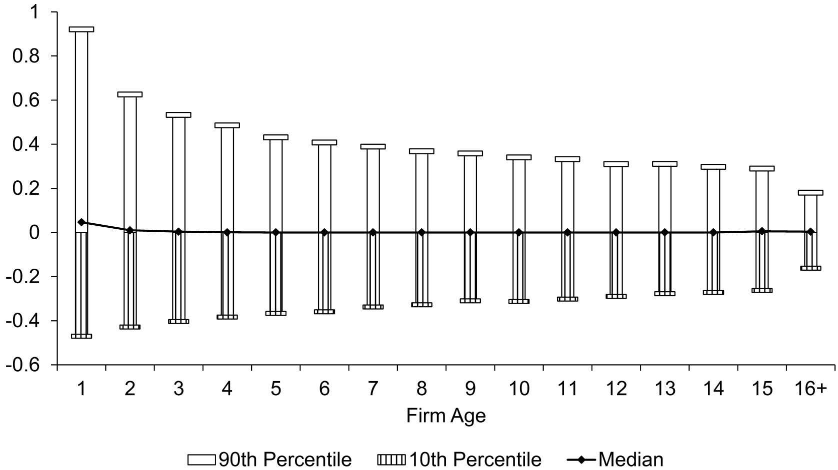

Fig. 1.3A Net employment growth, 1996–2013, LBD

Source: Statistics computed from the Longitudinal Business Database and revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: The 90th, 10th, and median are all based on the employment-weighted firm-level employment growth rate distribution for each firm.

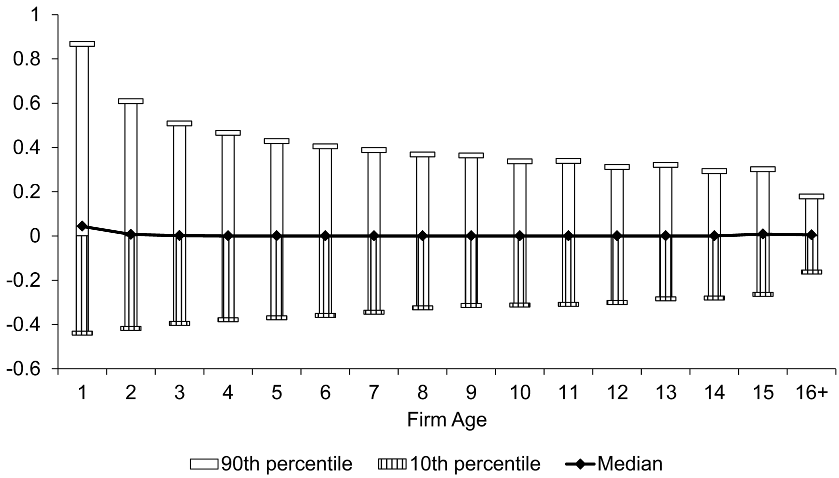

Fig. 1.3B Net employment growth, 1996–2013, revenue-enhanced subset of LBD

Source: Statistics computed from the Longitudinal Business Database and revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: The 90th, 10th, and median are all based on the employment-weighted firm-level employment growth rate distribution for each firm.

employment-weighted distribution, which mitigates the impact very small firms have on these statistics. We discuss dispersion by examining the patterns of the 90–10 differential and skewness by comparing the difference between the 90–50 and the 50–10 differentials.

Results from the full LBD and the propensity-weighted, revenue-enhanced sample are again very similar. Young continuing firms have very high dispersion of employment growth, and also very high positive skewness. The median employment growth rate for young firms is close to zero (and for that matter the median is close to zero for all firms) so the positive skewness is seen in the relative magnitudes of the 90th and 10th percentiles where the employment growth rates of younger firms are much more skewed to the right (positive) compared to more mature firms. This accounts for the high mean net employment growth rate of young firms relative to older firms from figures 1.1A and 1.1B. Taking figures 1.1A and 1.1B and 1.3A and 1.3B together, the typical young continuing firm (as captured by the median) exhibits little or no employment growth, even conditional on survival; however, among all the young firms, a small fraction exhibit very high rates of growth.

Our results thus far show that the full LBD and the revenue-enhanced subset yield very similar patterns for continuing firms, for entrants, as well as exiters. Comparison of figures 1.1A through 1.3B, and more extensive analysis contained in appendix A, indicate that by using propensity-score matching we are able to capture the basic patterns of firm behaviour from the LBD, giving us the confidence to proceed with our revenue-enhanced subset of LBD firms for the remainder of the analysis.

Output Dynamics

Keeping the pattern in figures 1.1A, 1.1B, 1.2A, and 1.2B in mind, we now characterize the distribution of output growth rates. We again use inverse propensity-score weights in calculations with the revenue-enhanced subset that permits measuring real gross output.

Figures 1.4A and 1.4B examine the output dynamics from continuers and from births, respectively. We first note that the patterns depicted in figures 1.4A and 1.4B are very similar to those in figures 1.1A and B and 1.2A and B. Young continuing firms experience on average high output growth rates relative to more mature firms. Young firms also experience higher rates of output destruction from exit. However, there are also some notable differences. We find output growth by continuers exceeds output destruction from exit for all age classes. Indeed, for most age classes, output growth for continuers exceeds destruction from exit. Comparing figures 1.1A and 1.4A, we find that young business exits generate larger percentage job losses than output losses. This is consistent with young business exits having relatively low productivity—a result emphasized in Foster, Haltiwanger, and Krizan (2001, 2006) for selected sectors. Turning to figure 1.4B, we can examine

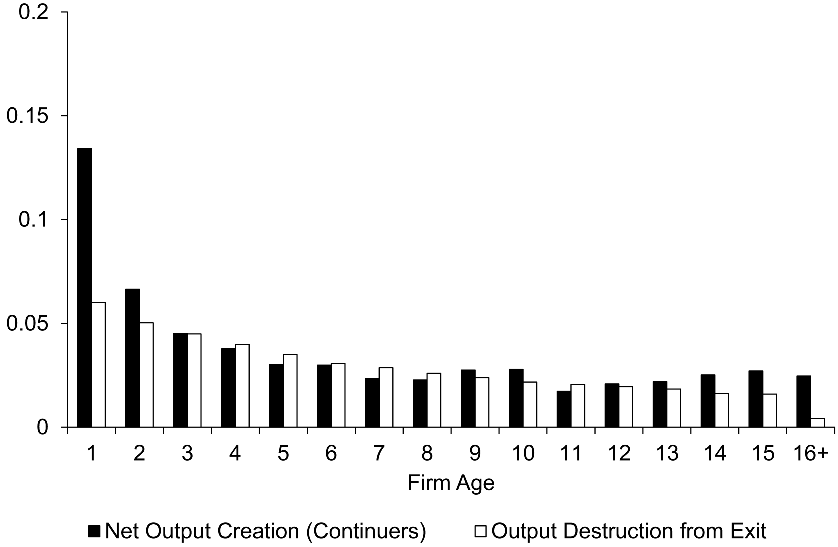

Fig. 1.4A Up or out output dynamics for firms, 1996–2013, revenue-enhanced LBD

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: Figure 1.4A shows patterns of net output growth for continuing firms and output destruction from firm exit for firms age one and older.

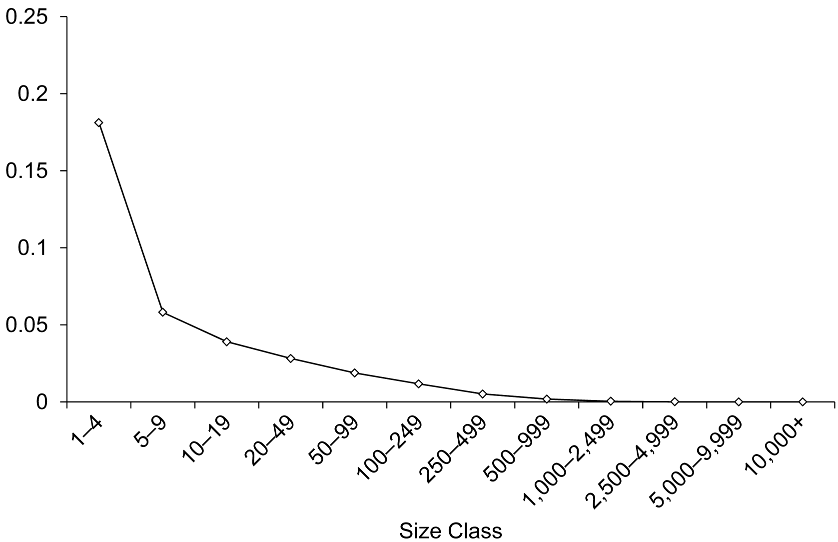

Fig. 1.4B Output creation from births, 1996–2013, revenue-enhanced subset of LBD

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: Figure 1.3B shows the pattern of output creation from firm births by size class.

the contribution of start-ups to output in their size classes. We see that the smaller start-up firms account for 18 percent of overall output in their size class. This is smaller when compared to their job contribution in figures 1.2A and 1.2B, but still a considerable amount.

1.3.2 Real Output versus Net Employment Growth Rate Distributions by Firm Age

Figure 1.5 characterizes the distribution of firm output growth rates by firm age for continuing firms. Depicted are the 90th, 50th, and 10th percentiles of the output-weighted distribution. As before, activity weighting mitigates the impact that very small firms have on these statistics since they account for only a small fraction of output. Comparing figures 1.3 and 1.5 yields many similarities, but also some notable differences. Output growth rates exhibit high dispersion and positive skewness for young firms in a similar manner to employment growth rates. However, while net employment growth for the median surviving firm is close to zero (except for age one firms), we find that real output growth for the median continuing firm is in excess of 4 percent per year in each of the first four years and in excess of 3 percent in all years.

The skewness of firm growth for young firms is less pronounced for output growth than for employment growth. However, we find that this is driven in part by cyclical dynamics.[15] Our revenue-enhanced subset of the LBD is only available from 1996–2013 so that the Great Recession plays a potentially important role. Figures 1.6A and 1.6B depict the 90–50 and 50–10 differentials for output growth (1.6A) and net employment growth (1.6B) for the subperiods 1996–06 (prior to the recession), 2007–10 (the recession), and 2011–13 (postrecession). The cycle clearly influences the skewness patterns, especially for output growth. In figure 1.6A, we find that the 90–50 exceeds the 50–10 for output growth for all firm ages at or below 5 and that the 90–50 is about the same as the 50–10 for firm ages greater than 5 for the period 1996–06.[16] However, in the recession period the 50–10 differential increases substantially for all ages so that rather than positive skewness, the output growth distribution exhibits negative skewness for most ages, and especially for older firms. In the postrecession period, we again see a pattern resembling that for the years 1996–06, although some of this may be the cyclical recovery. Figure 1.6B shows similar but more muted patterns when

15. Decker et al. (2015) emphasize that skewness of employment dynamics exhibits a negative trend. We have not investigated that pattern in the current chapter. Our sample is less well suited for examining changing trends since it starts in 1996 compared to 1981 for Decker et al. (2015). Still, the latter analysis emphasizes that the post-2000 period is a period of rapid decline in skewness in employment growth rates with differential patterns across sectors. Investigating these patterns using the output data would be of great interest.

16. The exclusion of 2001 and 2002 from our 1998–06 sample period may be playing a role here as well. However, we note that the full LBD shows substantial positive skewness in the employment growth-rate distribution using all years from 1998–06.

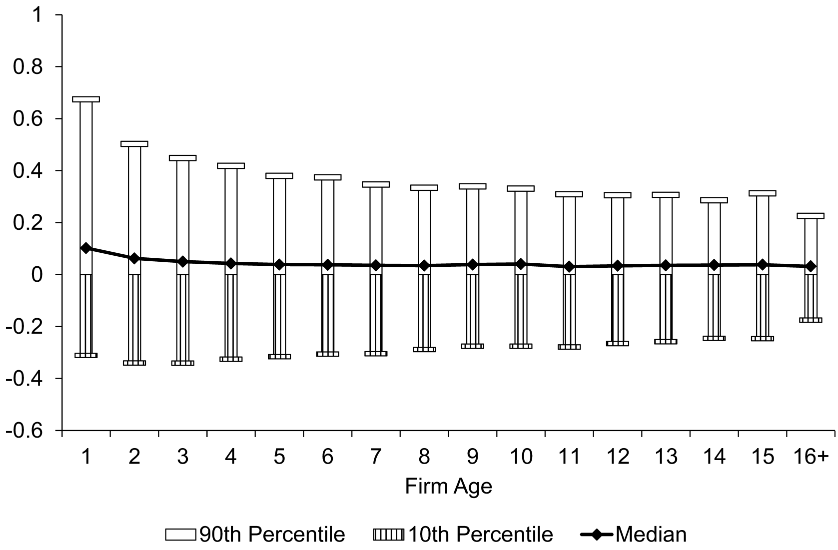

Fig. 1.5 Output growth distribution for continuing firms, 1996–2013

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: The 90th, 10th, and median are all based on the output-weighted, firm-level output growth rate distribution for each firm.

employment growth is considered as opposed to output growth. In short, we find that the positive skewness for young firms exhibited in terms of both output and net employment growth is procyclical. In what follows, for the sake of brevity, we will mostly present results for our entire sample period but we will note when patterns are especially sensitive to the business cycle.

Turning to figure 1.6C, we find that the mean output and net employment growth rates for surviving firms exhibit very similar patterns that decline sharply with firm age. Based on figures 1.3A, 1.3B, 1.5, 1.6A, and 1.6B, we know that underlying these quite similar mean patterns are differences in the shapes of the underlying distributions. For net employment growth, the high mean for young firms is driven by the positive skewness for young firms. Or put more simply, the high average is driven by high growth firms. For output growth, the high mean for young firms reflects both the high median for young firms and the greater positive skewness for young firms.

In either case, figures 1.3A, 1.3B, and 1.5 highlight the very high net employment and output growth of the 90th percentile firms, particularly for young firms. We quantify their importance in Table 1.1A where we decompose output and employment growth. We find that 12 percent of continuing firms have output growth in excess of 25 percent accounting for about

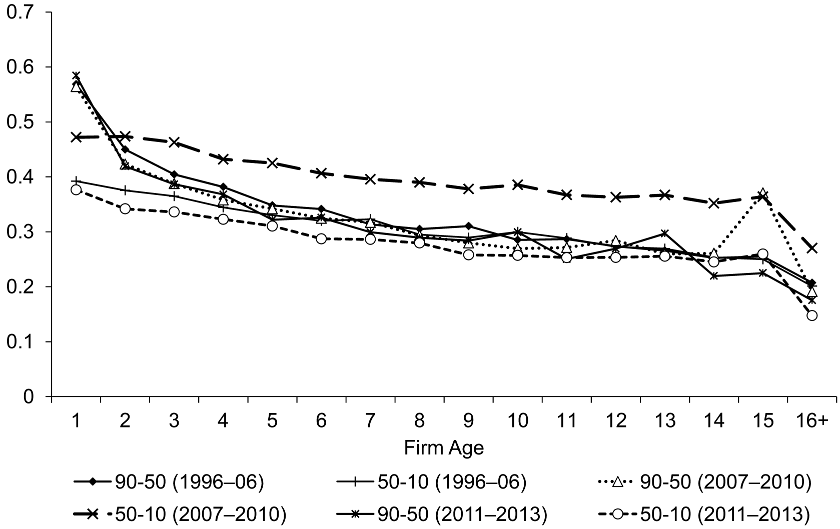

Fig. 1.6A The 90–50 versus 50–10 differentials for output growth for continuing firms by subperiods

Source: Statistics computed from revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: The 90th, 10th, and median are all based on the output-weighted firm-level output growth rate distribution for each firm.

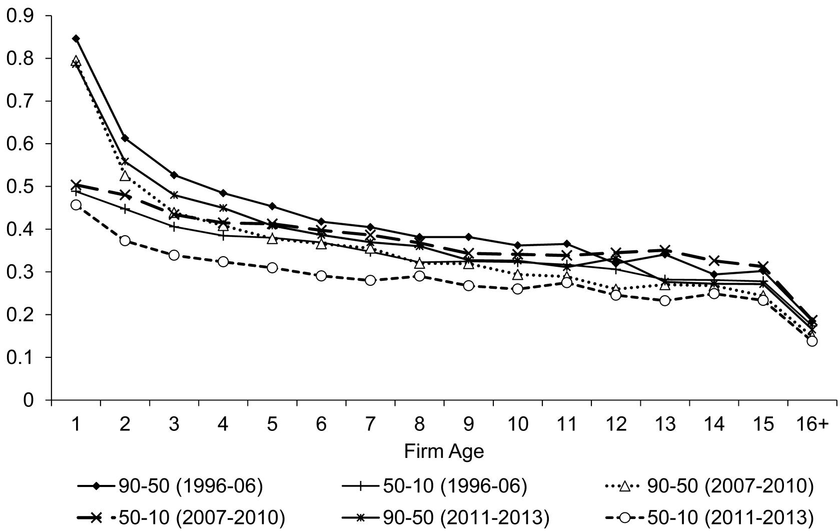

Fig. 1.6B The 90–50 versus 50–10 differentials for employment growth for continuing firms by subperiods

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Notes: The 90th, 10th, and median are all based on the employment-weighted firm-level employment growth rate distribution for each firm.

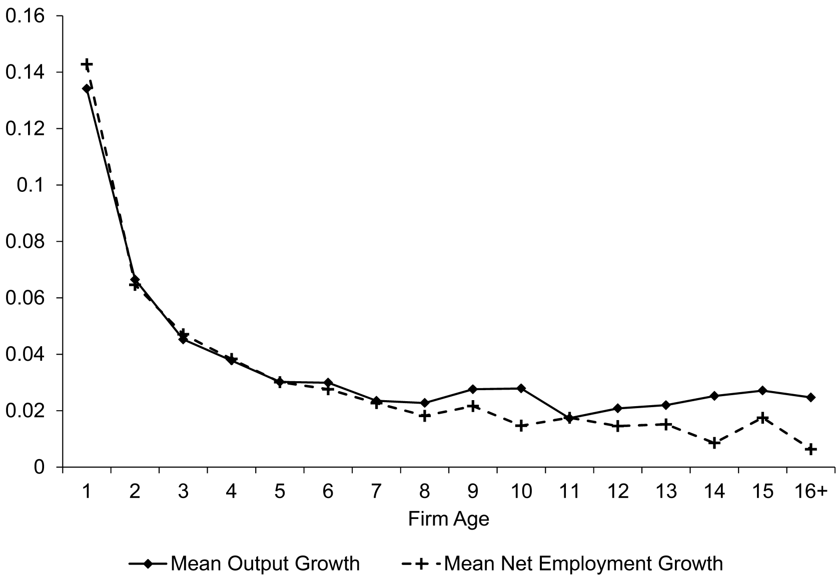

Fig. 1.6C Mean output and employment growth for continuing firms, 1996–2013

Source: Statistics computed from the revenue-enhanced LBD subsets 1998–2000 and 2003–2011.

Notes: The mean net employment growth is the employment-weighted average firm-level employment growth rate for each firm age. The mean output growth is the output-weighted average firm-level output growth rate for each firm age.

50 percent of the gross output creation for continuing firms.[17] Analogously, about 17 percent of continuing firms have net employment growth in excess of 25 percent accounting for close to 60 percent of gross job creation for continuing firms. Start-ups and exiting firms also contribute to employment and output growth. Table 1.1B looks at the contribution to output and employment growth from the entry and exit margins. Start-ups contribute disproportionately to employment and output growth. The contribution to employment growth is particularly large, accounting for an additional 25 percent of gross job creation versus 15 percent for output creation. Exiting firms account for a disproportionate share of employment, but this is less true for output.

In what follows, we explore the characteristics of high-growth firms on a number of margins. In particular, we consider not only firm age, but firm size, industry, and geographic location. We turn to that analysis below. Before doing so, we provide evidence on the connection between output and net employment growth rates.

17. By this we mean the output creation from growing firms.

Table 1.1A **The share of output and job creation accounted for by high-growth firms**

	High-growth firms	
	Output (%)	Employment (%)
Share of gross creation	49.8	58.5
Share of firms	12.3	16.7

Source: Tabulations from revenue-enhanced LBD subsets 1996–2000, and 2003–2013.

Table 1.1B **The share of output and job creation accounted for by births and deaths**

	Firms (%)	Output (%)	Employment (%)
Births	10.2	14.8	24.8
Deaths	8.8	10.7	26.1

Source: Tabulations from revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

1.3.3 The Joint Distribution of Real Output and Net Employment Growth Rates

Theoretical models of firm adjustment in response to shocks suggest a positive correlation between output and employment growth. This correlation may depend on the nature of the adjustment costs and frictions. We find that output and net employment growth rates for surviving firms are positively correlated but the contemporaneous correlation is not high (about 0.218). Further analysis shows that this reflects in part the pattern that output growth rates tend to lead employment growth rates. Table 1.2A shows the estimates from a simple reduced form one-lag vector autoregression (VAR) model relating firm-level net employment growth and output growth for continuing firms.[18] Net employment growth estimates reported in the first column show there is negative serial correlation reflecting the well-known regression to the mean in employment growth rates. Interestingly, however, lagged output growth is associated with higher net employment growth in the current period. The same is not true (to the same extent) for the relationship between lagged net employment growth and current period output growth shown in column (2), suggesting first that the output shock leads the employment adjustment, and second that output growth is only weakly correlated with prior growth shocks.

The patterns in table 1.2A are consistent with standard adjustment cost models for employment dynamics (see, e.g., Cooper, Haltiwanger, and

18. We weight the regressions with the LHS employment growth with employment weights and the regressions with the RHS output growth output weights. We have tried common weights and obtain similar results.

Table 1.2A VAR relating net employment and output growth

	Dependent variable	
Explanatory variables:	Net employment growth	Output growth
Lagged net employment growth	–0.103	0.023
	(0.000)	(0.000)
Lagged output growth	0.125	–0.003
	(0.000)	(0.000)

Source: Estimated specifications using revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Standard errors in parentheses. Specifications are weighted by employment and output, respectively.

Table 1.2B VAR relating probability of being high-growth-output and high-growth-employment firms

	Dependent variable	
Explanatory variables:	High-growth-employment firms	High-growth-output firms
Lagged high-growth-employment firms	0.066	0.071
	(0.000)	(0.000)
Lagged high-growth-output firms	0.131	0.123
	(0.000)	(0.000)

Source: Estimated specifications revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Standard errors in parentheses. Specifications are weighted by employment and output, respectively.

Willis 2007).[19] In such models, firms facing a positive profit (e.g., demand or productivity) shock exhibit immediate increases in output but a delayed adjustment for factors such as capital and labor.

We now explore whether the patterns at the mean of the growth rate distributions carry over to the upper tails of the joint growth rate distribution. Table 1.2B shows results for similarly estimated simple one-lag VAR models for indicators for high-growth episodes firms. For this purpose, a firm experiences a high-output (employment) growth episode in a particular year if the firm's output (net employment) growth rate is greater than 25 percent.

Table 1.2B shows that having a high-growth output episode in the previous year is positively associated with having both high-output and employment growth episodes in the current year. Interestingly, in spite of the overall negative serial correlation for employment growth in table 1.2A, there is

19. This likely also reflects the timing of the data. Employment growth from $t-1$ to t represents a March-to-March change, while output growth represents annual output changes during the calendar year from $t-1$ to t. Our primary focus is not on dynamics so we do not explore this issue further.

Table 1.3 **Distribution of high-growth and high-decline events for five-year-old firms**

Number of events	Output		Employment	
	High decline	High growth	High decline	High growth
0	61.30	39.42	50.69	32.69
1	25.99	28.01	33.47	32.73
2	9.57	18.67	12.94	21.78
3	2.66	9.41	2.62	9.40
4	0.41	3.42	0.27	2.86
5	0.06	1.07	0.01	0.54

Source: Tabulations from revenue-enhanced subset of LBD. Shares are percentages of columns.

some positive persistence in high employment growth episodes. These patterns are consistent with high-growth-output (employment) events extending beyond a single year, with high-growth-output events tending to precede high-growth-employment events.

Table 1.2B implies high-growth events exhibit positive persistence. But this simple VAR does not tell how often high-growth events are reversed. Table 1.3 provides insights for this latter question. For each five-year-old firm, we count the number of high-growth and high-decline events that the firm has experienced. A five-year-old firm can have between 0 and 5 high-growth and high-decline events. Table 1.3 shows the distribution of high-growth and high-decline events for both employment and output. The skewness highlighted earlier for young firms is self-evident in the much higher share of five-year-old firms having N high-growth compared to high-decline events.

Conditional probabilities are also easily computed from the joint distribution of high-growth and high-decline events.[20] The probability that a five-year-old firm with one, two, three, and four high-output growth events has zero high-decline events is 54 percent, 50 percent, 59 percent, and 74 percent, respectively.[21] Thus, most five-year-old firms with one or more high-output growth events have no high-decline events. Similar remarks apply to conditional probabilities for high-employment growth events compared to high-employment decline events.

1.3.4 The Characteristics of High-Growth Firms: By Firm Age, Firm Size, and Industry

Our objective in this section is to provide descriptive statistics about the characteristics of firms in the top of the growth-rate distribution. To this

20. The joint distributions are depicted in figures 1B.1 and 1B.2 of appendix B.
21. The (output weighted) probability that a firm with zero high-growth events has zero high-decline events is about 70 percent. This is not surprising since most output is at firms with zero high-decline events.

end, we estimate linear probability regressions pooling across firm years. We consider discrete dependent variables that take on a value of one if the firm is a high-growth output (employment) firm. As before, we define high-growth firms as those with annual growth in excess of 25 percent.[22] For the specifications with high-growth-output indicators, we weight by output (averaged in period $t - 1$ and t) and for the specifications with high-growth-employment indicators we weight by employment (averaged in $t - 1$ and t).

We first focus on firm age and firm size characteristics. For firm age, we consider firm age classes between one and sixteen and older. For firm size, we use within-industry deciles of the size distribution. In the case of the output-growth specifications, these are output-weighted deciles of output size. For the employment-growth specifications, we use employment-weighted deciles of employment size. For calculating these deciles, we use two alternative measures of size for output and employment. We use base year size (e.g., output or employment in period $t - 1$) and current average size (i.e., the average of output or employment in period $t - 1$ and t). We consider both, since as discussed in HJM, using base year size yields regression to the mean effects (i.e., given transitory shocks a firm classified as small in the prior period is more likely to grow). The use of current average size is a compromise between using base year and current year size (where the latter suffers from the opposite problem from base year size). We present our estimated firm size and firm age coefficients via a series of graphs. We do not report standard errors, but note given the very large sample size (in excess of 30 million) all of the standard errors for the reported size and age effects are less than 0.001. The same remarks apply to the state and industry effects that we report below.

Figures 1.7 and 1.8 report the estimated firm age effects for high-growth employment and output firms, respectively, with and without size controls. The likelihood of being a high-growth employment and output firm is decreasing with firm age even with firm size controls. The latter have relatively little influence on the patterns. It is apparent that our earlier findings in figures 1.3 and 1.5 are robust to controlling for firm size effects. We note that in unreported results we also find that these patterns are robust to controlling further for industry and year effects.

Figures 1.9 and 1.10 report the analogous estimated firm size effects for high-growth employment and output firms with and without age controls. For the firm size effects, we report results using both base year and current average size categories. If we do not control for firm age, there is an inverse relationship between firm size and the likelihood of being a high-growth firm using both the base year and current average size approaches. But once we control for firm age, these patterns are substantially mitigated. For high-growth-employment firms, the relationship between the likelihood of being

22. There is nothing inherently special about the 25 percent cutoff. We have found our results are robust to using alternative cutoffs.

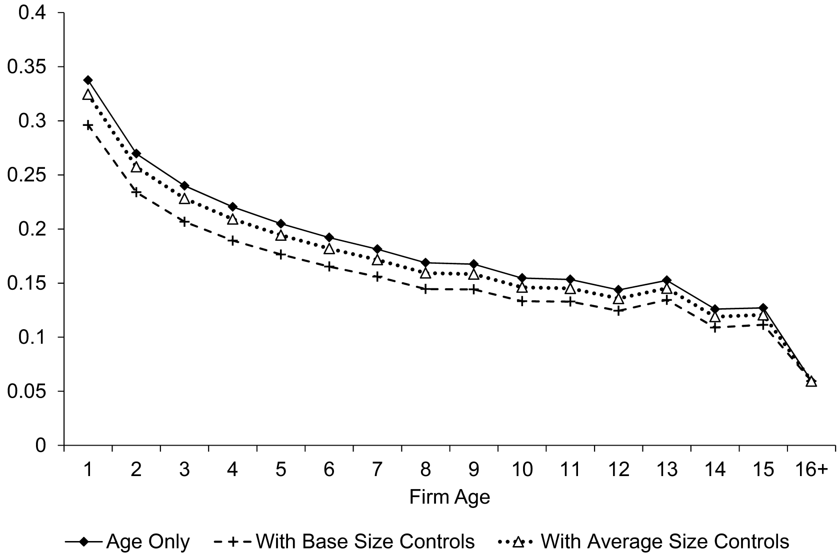

Fig. 1.7 Employment high-growth firms by firm age, 1996–2013

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on controls as listed. All coefficients are reported relative to unconditional mean for 16+.

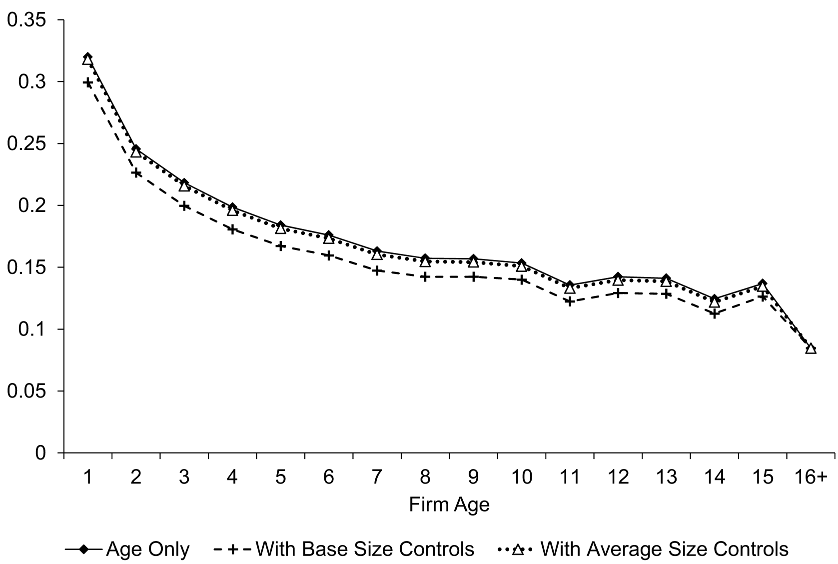

Fig. 1.8 Output high-growth firms by firm age, 1996–2013

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on controls as listed. All coefficients are reported relative to unconditional mean for 16+.

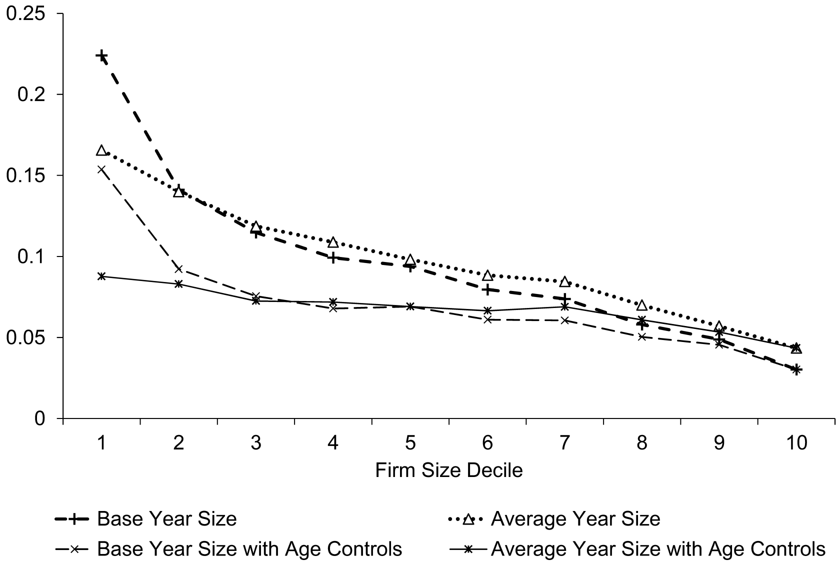

Fig. 1.9 Employment high-growth firms by firm size, 1996–2013

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on controls as listed. All coefficients are relative to unconditional means for top size decile.

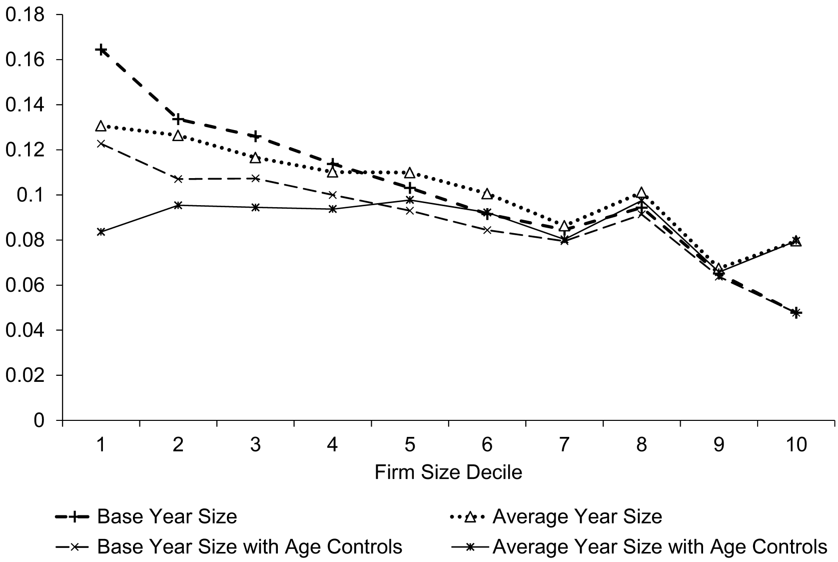

Fig. 1.10 Output high-growth firms by firm size, 1996–2013

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on controls as listed. All coefficients are relative to unconditional means for top size decile.

a high-growth firm and size is relatively flat using current year average size and age controls. For high-growth-output firms, the relationship becomes partly positive.

The inference we draw from figures 1.7–1.10 is that firm age is a robust and key determinant of the likelihood of being a high-growth firm. In contrast, once we control for firm age, firm size has relatively little influence. The role of firm age as opposed to firm size is reminiscent of the findings in HJM that found that young firms grow faster than more mature firms, but that small firms do not grow faster than large firms once firm age is taken into account. We note, however, that while firm age is a key determinant that the adjusted R-squared (see table 1.4) from age effects alone is 2 percent for the output growth distribution and 5 percent for the employment growth rate distribution. With size, industry, state, and year effects the adjusted R-squared rises to about 8 percent for output growth (using either base year or current average year size) and between 8 and 9 percent for employment growth. Industry effects alone yield 5 percent and 4 percent, respectively, in terms of adjusted R-squared. These patterns imply that the factors that determine which firm is a high-growth firm largely are factors within firm age, firm size, industry, and year cells that we do not observe in our data. Still there is systematic variation by industry to which we turn to now.

Figure 1.11A shows the top fifty industries and figure 1.11B shows the bottom fifty industries for high-growth-output firms. Figures 1.12A and 1.12B show the analogous patterns for high-growth-employment firms. Reported are the regression estimates with industry effects alone. Regressions are either employment or output weighted. We begin by noting that all four-digit North American Industry Classification System (NAICS) sectors have some high-growth firms. The top-ranked industries have high-growth firms that account for as much as 39 percent of industry output and 29 per-

Table 1.4 **Adjusted R-squared for effects accounting for high-growth firms**

	High-growth-output firms	High-growth-employment firms
Industry	0.050	0.041
Age	0.021	0.053
Base year size	0.011	0.033
Average size	0.004	0.015
Year	0.010	0.004
State	0.007	0.002
All effects (base year size)	0.082	0.092
All effects (average size)	0.078	0.081

Source: Estimated adjusted R-squared with dependent variable using the revenue-enhanced LBD subsets 1996–2000 and 2003–2013. Specifications are weighted by output (column [1]) and employment (column [2]), respectively.

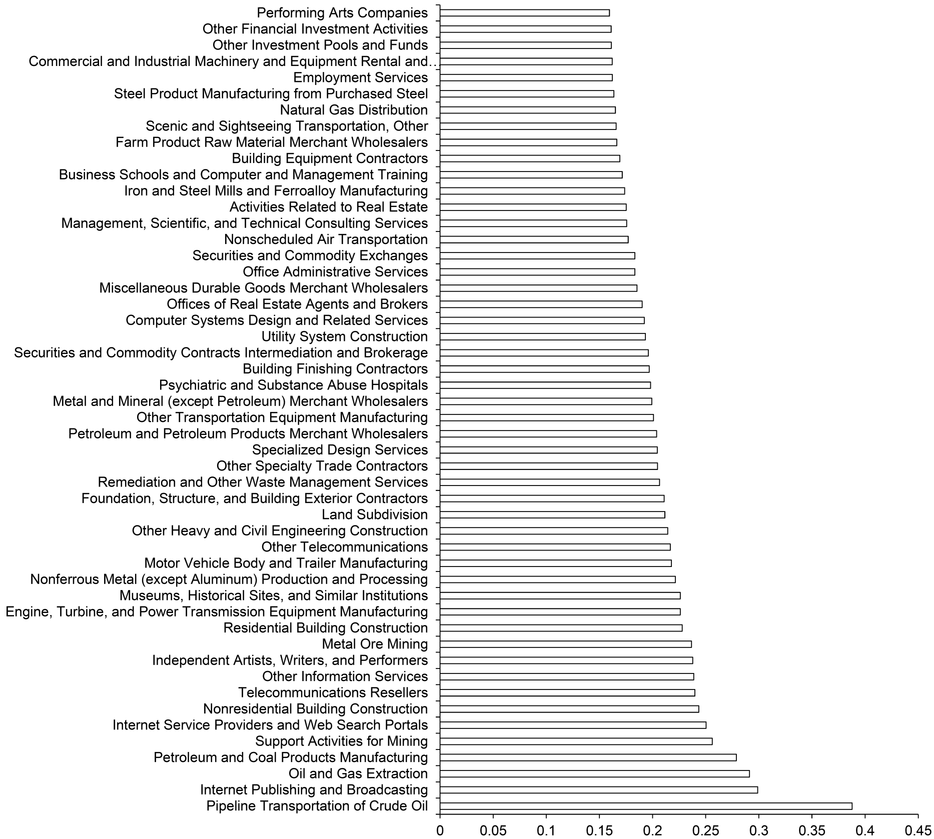

Fig. 1.11A Top fifty high-growth firms industry effects: Output

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on industry effects.

cent of industry employment. In contrast, the bottom-ranked industries have high-growth firms that account for less than 1 percent of industry output and employment.

Table 1.5A reports analysis of whether there are industry clusters that are more or less likely to have high-growth firm activity. The industry clusters we consider are sectors that can be classified as tradable, construction, high tech, bio tech, and energy related. We also include a small business-intensive sector dummy. This dummy is equal to 1 for the forty industries with the largest share of activity accounted by small firms where small is defined as having twenty employees or less.[23] Hurst and Pugsley (2012) suggest these

23. This follows Hurst and Pugsley (2012).

Fig. 1.11B Bottom fifty high-growth firms industry effects: Output

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on industry effects.

industries are disproportionally dominated by entrepreneurs with little interest or motivation for growth.

The dependent variable in table 1.5A is the share of either output or employment accounted by high-growth firms. We find that the energy-related sectors have greater high-growth firm activity in terms of output, but not employment. High-tech sectors have greater high-growth firm activity in terms of both output and employment. Tradeable sectors have lower high-growth activity, especially in terms of employment. The latter is consistent with the view that employment gains from high-growth firms in the tradable sectors have largely been offshored during our sample period. The construc-

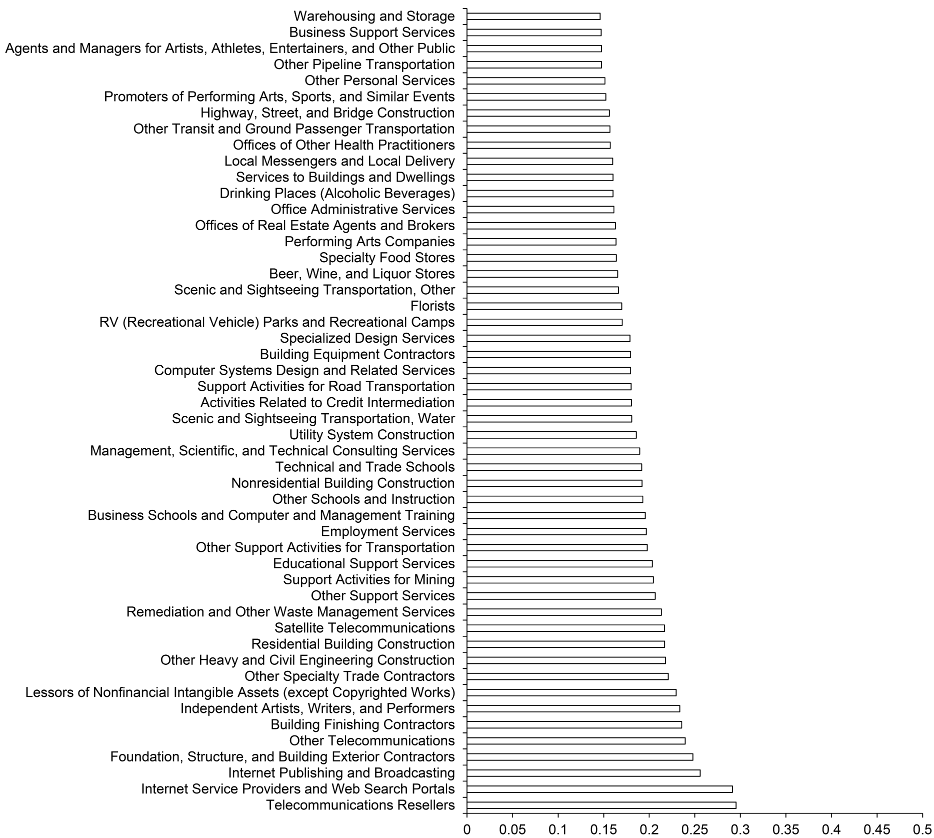

Fig. 1.12A Top fifty high-growth firms industry effects: Employment

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on industry effects.

tion sector also has especially high output-growth activity. This likely reflects the housing boom in the first decade of the twenty-first century. We find that the biotech sectors do not have significantly higher or lower high-growth activity. This contrast with the high-tech sectors is interesting and deserves further investigation. It may be that this is due to the way innovation takes place in this sector. One view is that successful biotech firms are much more likely to be bought up by large, mature firms rather than grow internally, since the process of bringing new pharmaceuticals from testing to the market (with all of the required approvals) favors the large firms.

We also find that small business-intensive sectors do not have significantly higher or lower high-growth activity in terms of output or employment. This finding might seem to be at odds with the hypothesis of Hurst and Pugsley

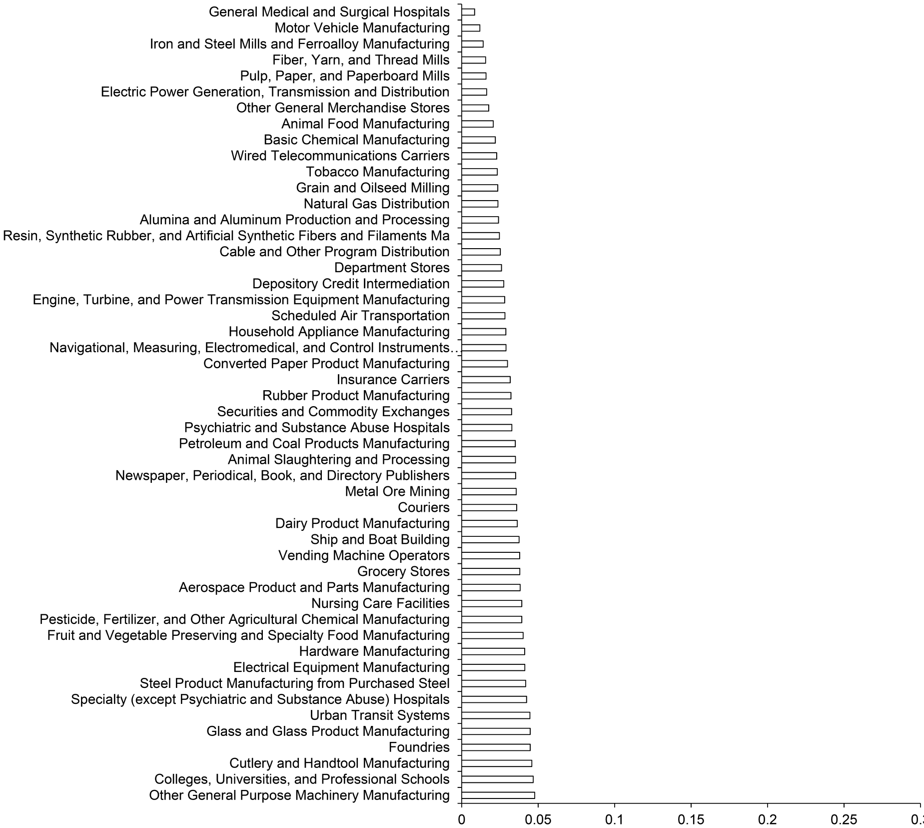

Fig. 1.12B Bottom fifty high-growth firms industry effects: Employment

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on industry effects.

(2012) that these are sectors where the typical firm is not growth oriented. That is, based on this hypothesis, we might have anticipated statistically significant negative effects. As we show in table 1.5B, there is an important and subtle difference between investigating what the typical firm is doing in a sector versus what share of activity is accounted for by high-growth firms. The key issue here is the role of activity weighting. The weighted rows in table 1.5B use the same dependent variables as in table 1.5A (namely, the share of industry activity accounted for by high-growth firms). The unweighted rows in table 1.5B use the share of firms that are high-growth firms (on either an output or employment basis) as the dependent variable.

When we examine what the average firm is doing on an unweighted basis we find some evidence in support of Hurst and Pugsley's hypothesis. Specifi-

Table 1.5A **The role of industry groupings in accounting for high-growth firms**

	Dependent variable: High-growth firm industry effects	
Explanatory variables:	Output	Employment
Tradable	–0.017	–0.053
	(0.008)	(0.007)
Construction	0.036	0.020
	(0.013)	(0.011)
High tech	0.033	0.037
	(0.016)	(0.013)
Bio tech	–0.042	–0.016
	(0.034)	(0.028)
Energy	0.044	–0.010
	(0.014)	(0.012)
Small business intensive	–0.008	0.015
	(0.010)	(0.009)

Note: Dependent variables are the estimated industry effects on high-growth firms.

Table 1.5B **Small business-intensive dummy variable coefficients by dependent variable**

Dependent variable	Small business-intensive dummy coefficient
Employment HGF weighted	0.036
	(0.009)
Employment HGF unweighted	0.004
	(0.005)
Output HGF weighted	0.006
	(0.010)
Output HGF unweighted	–0.022
	(0.008)

Note: Dependent variables are the estimated industry effects on high-growth firms.

cally, the probability that a firm is a high-output growth firm is lower in small business-intensive sectors; however, that same relationship does not hold for high-employment growth firms. But these findings do not imply there is less overall activity accounted for by high-growth firms in these industries. Weighted results show small business-intensive sectors do not have lower shares of activity accounted for by high-growth firms.

These results help reconcile the alternative perspectives in Hurst and Pugsley (2012) and our work. The average firm in these small business-intensive sectors is unlikely to grow (at least in terms of output). But even in these sectors, there are on an activity-weighted basis sufficient high-growth firms that these sectors have no less activity than other sectors in high-growth activity. Looking back at figures 1.3 and 1.5, recall that overall the median firm exhibits little growth, so it is not surprising that in small business-intensive sectors the typical or average firm exhibits lower than average propensity to grow rapidly. But even in these sectors there are enough high-growth firms in the tail of the activity-weighted growth rate distribution that small business-intensive sectors have broadly similar levels of high-growth firm activity.

1.3.5 Sensitivity Analysis

We now briefly discuss the results of a number of sensitivity analyses we conducted and that are described more fully in appendix A. The rankings of industries in terms of high-growth output and high-growth employment events exhibit substantial positive correlation (see figure 1B.3). We also find that the industry rankings for high-growth firms are reasonably stable over time (see figure 1B.4). We also investigated the relationship between the first three moments of the growth-rate distribution by industry. We find that industries exhibiting greater high-growth firm activity also tend to have high overall mean growth on both an output and employment basis (see table 1B.1). But industries with a large fraction of high-growth activity also have a large fraction of high-decline activity and greater volatility as captured by the 90–10 differential. High-growth activity is also associated with greater positive skewness in the respective growth-rate distribution (as measured by the difference between the 90–50 and 50–10 differentials in net output and employment growth rates). These patterns are also exhibited in scatter plots (see figures 1B.5 and 1B.6).

Appendix B shows related types of exercises for state variation in the share of high-growth firm activity (see figures 1B.11, 1B.12, and table 1B.2). We find that states with a larger share of output in high-growth-output firms are also states with a larger share of employment in high-growth-employment firms. States at the top on the basis of output growth are energy-intensive states such as Oklahoma, South Dakota, Texas, and North Dakota. For the ranking by high-growth firms by employment, Oklahoma is toward the top but Texas and North Dakota are not. There is a very strong correlation in high-growth output and employment states compared to industries. There is somewhat less stability of state rankings relative to industry rankings. We find that high-growth firm effects for a state are positively related to overall growth, dispersion, and skewness.

1.4 Firm Age and Productivity Dynamics: The Role of High-Growth Young Firms?

We now turn to the relationship between high-growth young firms and productivity dynamics. The revenue-enhanced LBD is the first economy-wide database to include measures of output and productivity on an annual basis. We are especially interested in the contribution that high-growth young firms have on the reallocation components of productivity growth. As such, we focus on continuing firms in this section. We use the output and employment measures to construct a labor productivity measure for each firm. Since we use gross output and not value added, these statistics are not comparable across industries so we focus on within-industry patterns. This controls for industry-specific differences in intermediate input shares. In addition, since we do not have industry-specific output deflators for all industries, we always control for industry by year effects in this section. This is equivalent to controlling for industry-specific deflators.

Figures 1.13A and 1.13B show the mean and standard deviations of the within-industry (log) labor productivity measure by firm age. We construct these figures as follows. First, we compute the within-industry means and standard deviations within each detailed six-digit industry for each firm age group. In figure 1.13A we generate these means on an unweighted basis, and in figure 1.13B we use employment weights to weight up to the industry level. Then, in both figures, we take an average across all industries where we use gross output weights following the procedures used in Foster, Haltiwanger, and Krizan (2001) and Baily, Hulten, and Campbell (1992). For the mean calculation we index the average productivity of sixteen-year-olds and older at 1 so that the reported effects reflect differences from that oldest group.

Figure 1.13A shows that, relative to other firms in the same industry, mean (log) labor productivity rises with firm age whether we use the unweighted or weighted approach within industries. However, the differences by firm age are much larger in magnitude using the weighted approach. When we weight by employment, the patterns reflect both the unweighted mean within the industry, firm age cell, and the covariance between size and productivity within the cell as per the Olley and Pakes (1996) decomposition.[24] The weighted mean patterns show a more dramatic increase with firm age. By construction, this pattern reflects a sharp rise in the covariance between size and productivity within an industry by firm age cell. The latter pattern is not surprising, since for young firms the relationship between size and productivity is likely weak as firms have not sorted themselves out in terms of the relationship between relative size and productivity. Another possible factor is that measurement error is greater for young firms, but this should be less problematic in this setting given the use of administrative data.

Figures 1.13A and B also show that the within-industry dispersion of productivity declines monotonically with firm age. For both the unweighted

24. The Olley-Pakes (1996) decomposition of the level of productivity is given by:

$$P_{it} = \Sigma_{e\epsilon i}\, \omega_{et} P_{et} = \widetilde{P_{et}} + \Sigma_{e\epsilon i}(\omega_{et} - \widetilde{\omega_{et}})(P_{et} - \widetilde{P_{et}}),$$

where a tilde represents the simple average across all plants in the same industry. When we compute the weighted average productivity for each age group and compare it to the unweighted average, the difference is the Olley-Pakes covariance term for the age group.

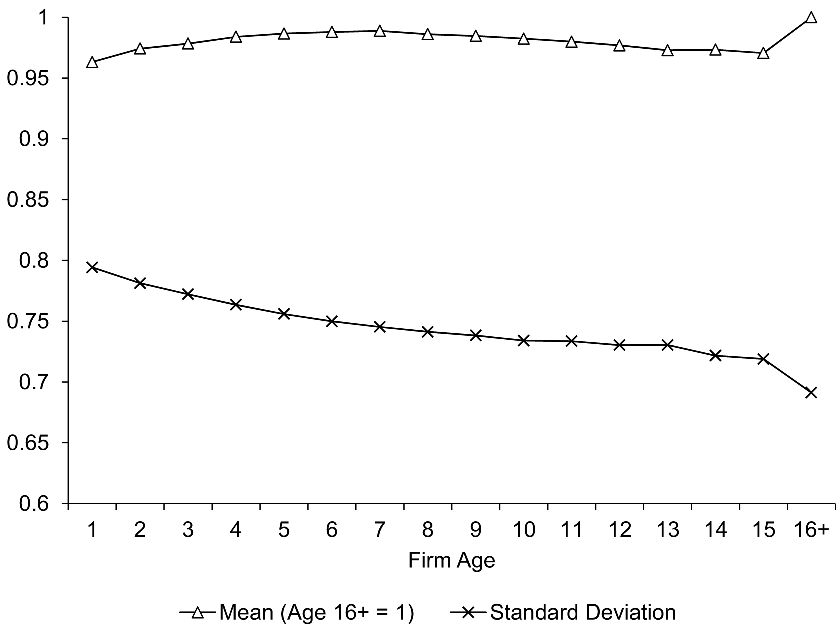

Fig. 1.13A Mean and standard deviation of within-industry labor productivity by firm age (unweighted within industries)

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

and weighted results, we find similar patterns. The patterns are consistent with our findings of much greater dispersion in both output and net employment growth for young firms.

To explore the contribution of high-growth firms, we turn to examining within-industry decompositions of industry-level productivity growth for continuing firms.[25] We control for industry and time effects via the decomposition itself. We define an index of industry-level productivity as given by:

$$P_{it} = \sum_{e \epsilon i} \omega_{et} P_{et},$$

where e indexes firms, i indexes industry, P is log labor productivity, and ω is the share of employment. Note that for the purposes of a labor productivity index the appropriate weight is employment, since then the index is the geometric mean of firm-level labor productivity. Then the change in this index at the industry level (which is log based so that it can be interpreted

25. From the existing literature (see, e.g., Foster, Haltiwanger, and Krizan 2001, 2006), we know that net entry contributes disproportionately to within-industry productivity growth as exiting businesses are much lower in productivity than entering businesses. We are focusing on continuing firms given the limitations of our output-restricted database.

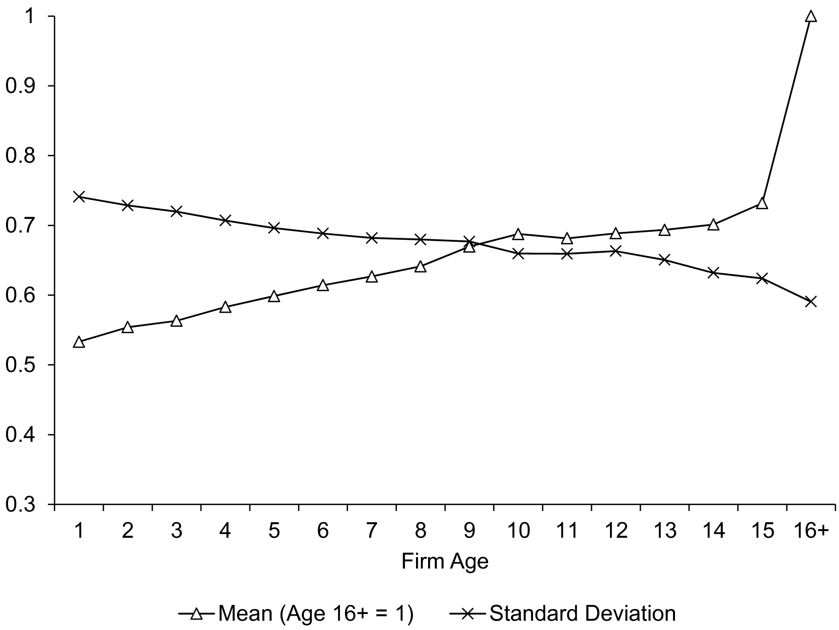

Fig. 1.13B Mean and standard deviation of within-industry labor productivity by firm age (weighted within industries)

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

as an index of industry-level productivity growth) can be decomposed into within and between effects as given by:

$$\Delta P_{it} = \sum_{e \epsilon i} \bar{\omega}_e \Delta P_{et} + \sum_{i} (\bar{P}_e - \bar{P}_i) \Delta \omega_{et},$$

where a bar over a variable represents the average over $t - 1$ to t. The first term on the right-hand side (RHS) is the within term and the second term is the between term. The within term captures the weighted average of within-firm productivity growth, while the between term captures the contribution of changes in employment shares. A firm contributes positively to the between term if it has labor productivity higher than the industry average. In this decomposition, we focus on within-industry patterns by using an industry-specific decomposition.

We calculate this decomposition for every industry/year pair in our data. To compute an aggregate average we use average gross output weights following the approach of Foster, Haltiwanger, and Krizan (2001) and Baily, Hulten, and Campbell (1992). Figure 1.14 shows the results of this decomposition for the average (output weighted) industry for all years and for the subperiods 1996–06, 2007–10, and 2011–13. We find that the within term

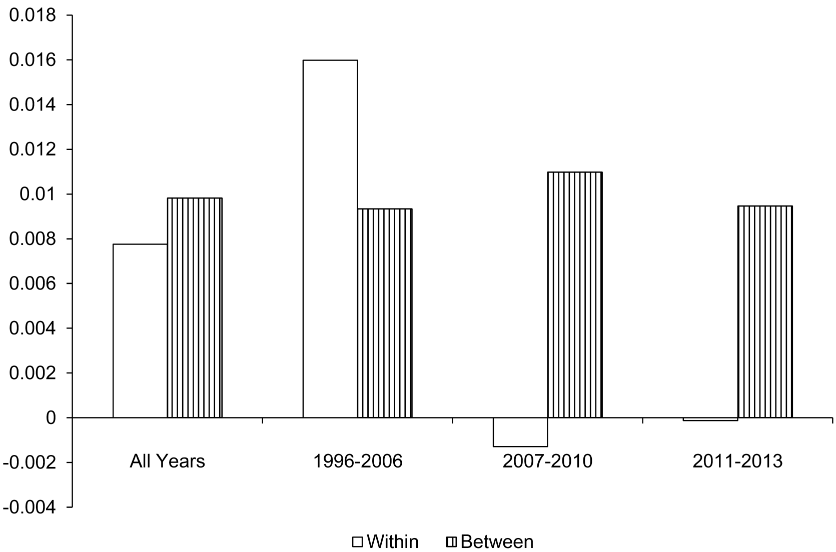

Fig. 1.14 Within versus between components of within-industry labor productivity growth

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

is highly procyclical. It is positive for the overall period, much higher in the 1996–06 period, negative in the 2007–10 period, and almost zero for the 2011–2013 period. This is consistent with within-firm productivity being procyclical—likely for reasons associated with varying capacity utilization and the adjustment costs discussed earlier. In contrast, the between term is more stable over time and is always positive. It is not surprising then that the between term accounts for much of the overall increase for the full period and about half the increase for the period 1996–06.

To explore the role of high-growth firms we focus on the between term, since it is both more stable but also captures the reallocation dynamics where high-growth firms play such a critical role. Figure 1.15A shows the contribution of each of the output growth rate classes, and figure 1.15B shows employment growth rate classes by time period. Interestingly, we find that it is especially high-growth and high-decline businesses that account for the between term and the patterns are roughly similar for both output and employment-growth rate classes. For the growth-rates classes with relatively modest increases or decreases we find little contribution of the between component. Since the between term is only positive for a group of growing (shrinking) firms if they have productivity that is on average higher (lower) than the overall average, these findings imply rapidly growing firms have above average productivity, while rapidly shrinking firms have

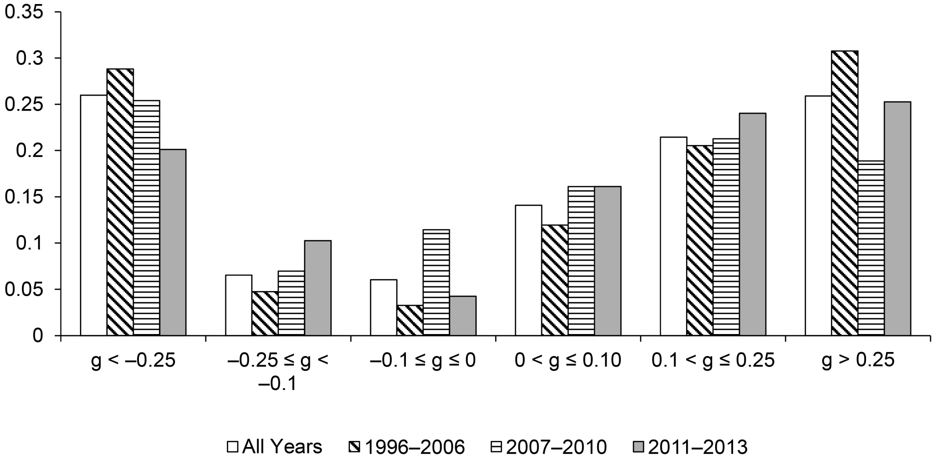

Fig. 1.15A Share of between accounted for by output growth rate classes
Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

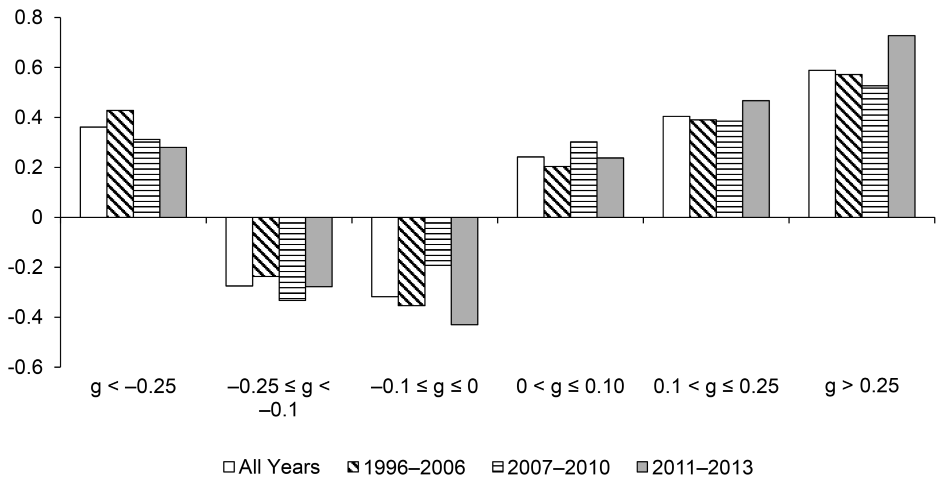

Fig. 1.15B Share of between accounted for by employment growth rate classes
Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

below average productivity. In this respect, Figures 1.15A and 15B remind us that the high-growth firms are part of the overall dynamic contributing to productivity-enhancing reallocation with an equally important role for rapidly shrinking low-productivity firms.

Where do young high-growth firms fit into this picture? First, we note that young firms that are less than ten years old only account for about 13 percent of output and 19 percent of employment. But we find that young firms contribute about 50 percent to the between term—much higher than their share of activity. Also, high-growth young firms contribute about 40 percent to the high-growth component of the between term. Thus, we find that young firms disproportionately contribute to the between term overall, and that high-growth young firms contribute disproportionately to the between-term contribution of high-growth firms.

1.5 Concluding Remarks

We find that high-growth young firms contribute disproportionately to job creation, output, and productivity growth. Young firms are very heterogeneous. Many fail in their first few years, and even among those that survive there is considerable dispersion in the growth patterns they experience. Conditional on survival, young firms have higher average net employment growth and output growth than their more mature counterparts. For employment growth, this is especially striking since median net employment growth for young firms is about zero. As such, the higher mean reflects the substantial positive skewness with a small fraction of very fast-growing firms driving the higher mean net employment growth. For output growth, young firms have higher median growth than their more mature counterparts. Still, young firms exhibit more positive skewness in growth rates than their mature counterparts on both an employment and output growth basis—although the positive skewness of output growth for young firms is highly procyclical.

Given these findings, we explored the characteristics of the high-growth firms further. Consistent with the above, we find that high-growth firms are more likely to be young than mature, even controlling for firm size. We also found that there is considerable variation across industries and states in the fraction of activity accounted for by young firms. The range across industries is substantial. Industries at the top of the ranking have as much as 40 percent of activity in high-growth firms, while industries at the bottom of the ranking have close to zero. In the post-2000 period, the share of activity accounted for by high-growth firms is significantly higher in the high tech and energy-producing (for the latter the share of output) industries. A firm in a small business-intensive industry is less likely to be a high-growth-output firm, but small business-intensive industries do not have significantly smaller shares of activity accounted for by high-growth firms for either output or employment. These findings are not in conflict with each other, but rather emphasize the importance of distinguishing between the typical (median or even average firm) and the activity accounted for by high-growth firms. Small business-intensive sectors often still have a small but highly influential contribution of high-growth firms.

We find that the ongoing reallocation dynamics of which high-growth young firms play a critical part contributes substantially to within-industry labor productivity growth. Our findings suggest that at least half of within-industry labor productivity growth for continuing firms is attributable to employment being reallocated from less productive to more productive firms within the industry. Young firms contribute disproportionately to this contribution from reallocation. But in this respect both high-growth and high-decline firms contribute substantially to the productivity-enhancing reallocation.

Industries and states with a greater fraction of high-growth firms exhibit high overall net growth, higher volatility, and also higher positive skewness. In this respect, a propensity for high growth is an indicator related to first, second, and third moments of the growth-rate distribution.

We interpret our findings as being consistent with models of innovation and growth that impact the first, second, and third moments. A rough storyline that we think fits the patterns we observed is as follows: firms with positive productivity realizations (exogenously or through endogenous innovative activity) leads to growth that contributes to both dispersion and positive skewness in the firm growth-rate distribution. The latter reflects the rareness of being a successful innovator (i.e., being in the right tail of the productivity distribution), and those that do succeed exhibit rapid growth. Those rare rapidly growing firms contribute substantially to net job creation, output growth, and labor-productivity growth. But often accompanying growth are those that do not succeed, so that volatility accompanies growth.

This storyline is obviously incomplete on many dimensions. It may be that (and we have presented some evidence of this) shocks and innovation in some sectors do not involve this complex dynamics of entry, exit, volatility, and skewness. Another set of issues relate to industry life cycle; that is, what do the dynamics of industries and locations in decline look like. They may also involve volatility. Similarly, there may be shocks that induce reallocation without much productivity growth or even adverse consequences for growth. For example, uncertainty shocks of the type emphasized by Bloom (2009) may have this character.

Our analysis has been intentionally descriptive. We think the data infrastructure we have developed and the basic facts we have presented provide a framework for more direct analysis of the process of innovation and growth. Our findings suggest that exploring patterns by firm age and examining first, second, and third moment effects will be important for detecting and understanding periods of growth and innovation. Moreover, we think our data infrastructure and approach should be helpful to explore factors that distort innovation and growth. The recent findings of Hsieh and Klenow (2014) that show that young firms grow rapidly in the United States relative to their counterparts in India and Mexico is highly relevant in this context. Our findings show that the rapid growth of young firms in the United States involves substantial skewness and dispersion. As such, distortions that may be adversely impacting the growth of young firms in India and Mexico (among other countries) may be impacting many of the different margins that underlie the patterns we have detected.

Appendix A

Data Appendix

In this appendix, we describe the construction of the firm-level revenue variable that serves as the basis for our analysis. We then describe how this variable is used to construct a revenue-enhanced subset of the LBD that includes continuers, births, and deaths, and discuss our methodology for cleaning the data. Finally, we describe our implementation of propensity-score matching to control for potential selection effects. In presenting our propensity-score models, we compare propensity-score-adjusted job creation and job destruction statistics from the revenue-enhanced subset to the results for the full LBD to indicate the effectiveness of our strategy.

Construction of the Revenue Variable

The US Census Bureau Business Register files contain revenue data sourced from business administrative income and payroll filings. These data are used for statistical purposes, including the Economic Census program and the Nonemployer Statistics program. There are a number of different tax forms and different revenue items within those forms that are relevant for calculating firm-level revenue depending on the sector that a firm operates in (or more specifically, the particular reporting tax unit, the EIN, within a firm), as well as the legal form of organization of the firm (nonprofits, partnerships, corporations, or sole proprietors). In an effort to build revenue measures reasonably comparable across firms, starting in 2002 the census developed an algorithm that takes these differences in tax forms and revenue concepts into account.[26] Within the census, this "best receipts" variable has previously been applied to single-unit firms only. Thus, we extended the

26. Algorithms are available to Census Bureau employees and RDC researchers that have an approved project and a need to know. Depending on the form and industry these algorithms may include total revenue, net receipts, gross revenue, receipts from interest, receipts from gross rents, total income, cost of goods sold, and direct as well as rent expenses.

original methodology in two ways. First, we apply the Census Bureau methodology to multiunit firms. Multiestablishment firms can report different parts of their operations under different and independent EIN filings. There are many possible reasons firms organize across multiple EINs including geographic, tax status, or business considerations. Given these within-firm sources of variation, we apply the algorithm at the EIN level first, using the EIN's self-reported NAICS classification to assign an industry to the EIN. The taxable revenue items that are included in the EINs total revenue are determined by this industry designation. We then compute a firm-level revenue measure by summing up all of the EINs associated with a particular firm.

Second, we developed an analog of the algorithm for years prior to 2002. The Business Register went through a complete redesign in 2002, which made it possible to keep additional fields that had been combined in prior years. We modify the pre-2002 algorithm to adjust for the different revenue items available before 2002. For any given year of revenue, we use prior-year revenue variables from the following year's BR. Previous research from the census has indicated that due to extended filing schedules, late filing, and other factors, these prior-year revenue variables provide significantly improved revenue information. Thus, in applying our algorithms we always use revenue for a given year from the BR file for the following year. Figure 1A.1 shows the results of applying these algorithms on the BR revenue measures and after filtering. Revenue is deflated using the GDP Implicit Price Deflator. Real revenue is in 2009 dollars.

The Revenue-Enhanced LBD Subset

Based on the revenue variable describe above, each observation in the LBD falls into one of four revenue categories: revenue continuers with revenue data in both year $t-1$ and year t, revenue deaths with revenue in year $t-1$ but no revenue in year t, revenue births with no revenue in year $t-1$ and revenue in year t, and observations with no revenue data at either time. Observations in the fourth category are dropped in their entirety from the sample, while the subsets represented by the first three categories are cleaned to ensure that the observations are suitable for analysis.

Inspection of the revenue data reveals a number of outliers. These can come about for a number of reasons including typographical errors, OCR errors, units errors, and even denomination errors. Outliers are also common among commodity and energy-trading entities, as well as businesses organized in terms of holding companies. To address these issues, for the revenue continuers subset we apply the following filters:

1. We drop observations with labor productivity (revenue divided by employment) above the 99.9th percentile and below the 0.1th percentile for both years $t-1$ and t.

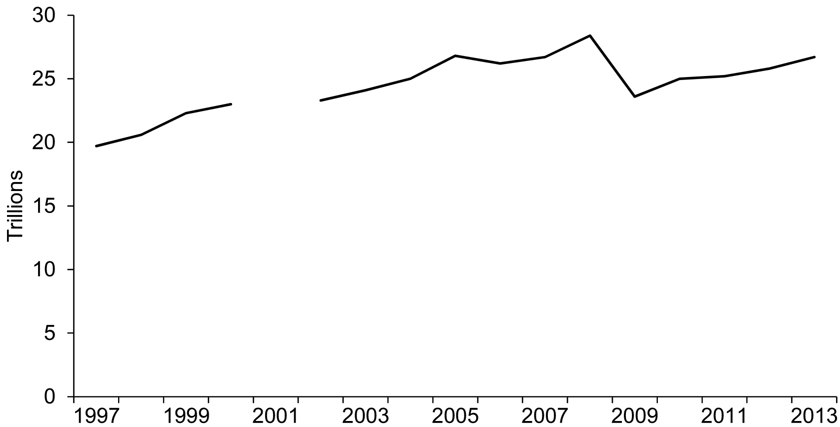

Fig. 1A.1 Real revenue (2009 dollars)

2. We drop observations reporting over $1 billion in average revenue and a DHS revenue growth rate of less than –0.5 or greater than 0.5.

3. We drop observations reporting over $100 million in average revenue and a DHS revenue growth rate of less than –1.5 or greater than 1.5.

4. We drop any observations reporting $1 trillion in average revenue or more.

These filters are designed to narrowly target specific problems such as unusually high or low labor-productivity values, unusually high revenue values, and unusually high changes in revenue, all the while minimizing the number of records we exclude from the data. Overall, this procedure excludes 0.14 percent of the total universe of revenue continuers.

For the revenue deaths and births we apply the same labor productivity filter for the relevant year of revenue. Because all revenue deaths and births have DHS revenue growth rates of –2 or 2, application of the additional filters amounts to a restriction on the DHS revenue denominator of $100 million. Overall, this procedure excludes 0.08 percent and 0.13 percent of the total universe of revenue deaths and births, respectively. Then, so that only true employment deaths and births are counted, the revenue death and revenue birth files are restricted to observations that represent employment deaths for the former and employment births for the latter. The remaining observations from each subset are then combined to form the revenue-enhanced LBD subset.

Missing Observations, Selection, and Propensity-Score Matching

Firms typically use the same EINs when filing payroll and income tax reports. This facilitates linking employment and revenue activity for a given firm at the Census Bureau. However, this is not always the case. About 20 percent of businesses file their payroll and income reports under different EINs. When this happens, the Census Bureau has no direct way of linking the two records. These revenue EINs become orphan records to payroll EINs, although they are never identified as such. Revenue records without a corresponding payroll record are considered nonemployer EINs.[27] The practical consequence of this is that for 21.8 percent of the revenue-enhanced LBD subset, we are missing revenue data. Further, it is often the case that employers will consistently use different EINs when filing their payroll and income, so many of these firms are missing all of their revenue data, making it difficult to impute their records. In addition to potential selection resulting from the examination of only observations that have revenue data, the additional filters and restrictions placed on the data may create problematic selection effects, particularly in the case of deaths and births.

Given that selection effects may differ for continuers, deaths, and births, we developed separate propensity-score models for employment continuers with revenue data, employment deaths with revenue data, and employment births with revenue data. Each of these partitions constitutes the set of firms for which the dependent variable equals 1 in a propensity score model that is run on the universe of LBD employment continuers, LBD employment deaths, and LBD employment births, respectively. For the employment continuers, the propensity score is inverse probability-weight calculated from the predicted values from a logistic regression including firm size, firm size squared, firm age, firm age squared, an indicator variable for firms of age sixteen and older, employment growth rate (seven classes), broad industry (twenty classes), and a multiunit status indicator. For deaths, we employ the same model, except we exclude the growth-rate classes. Finally, for births the model includes firm size, firm size squared, broad industry, and the indicator for multiunit status. Figures 1A.2, 1A.3, and 1A.4 examine the performance of our propensity-score model in terms of total net job creation, job destruction from exit, and job creation from births. Although these figures indicate some modest selection effects present in the revenue-enhanced LBD subset, the propensity-score model yields patterns of employment growth dynamics for continuers, births, and deaths for the enhanced-revenue subset of the LBD that closely mimic those for the full LBD. Figures 1A.2–1A.4 also show that even without weighting the enhanced-revenue subset does a reasonable job of capturing the employment dynamics from the full LBD.

27. For example, corporations file Form 1120 for their income taxes and Form 941 for their employment taxes (http://www.irs.gov/Businesses/Small-Businesses-&-Self-Employed/Corporations).

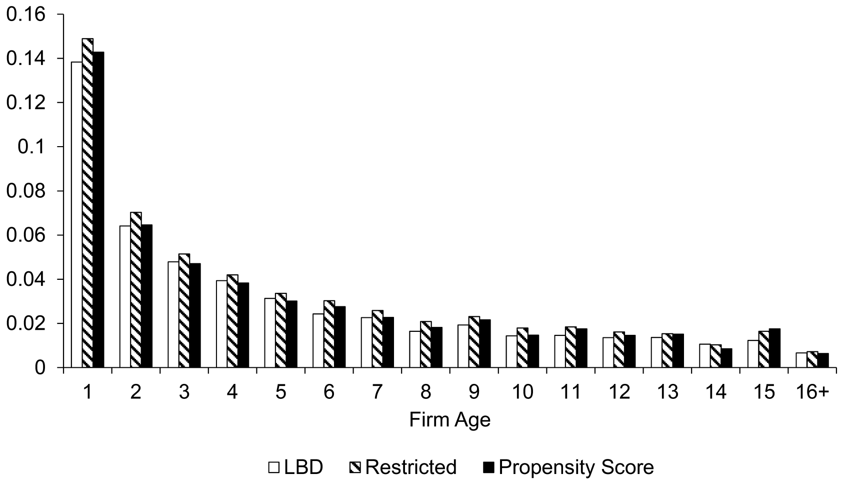

Fig. 1A.2 Net employment growth for surviving firms by sample, 1996–2013

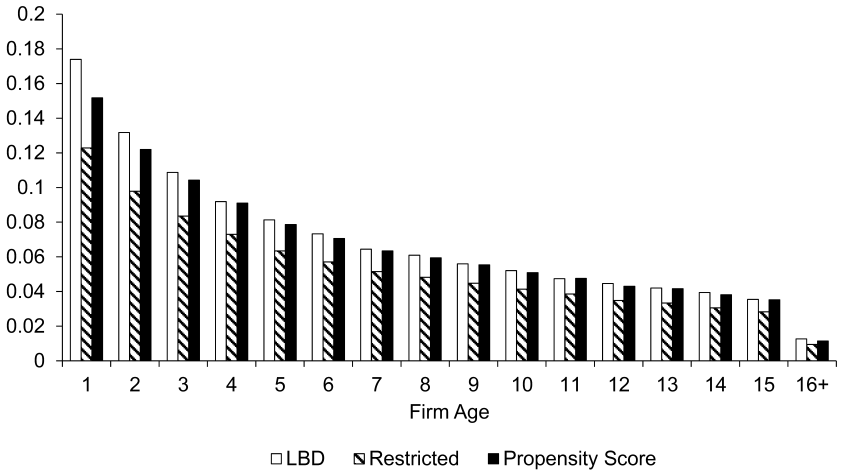

Fig. 1A.3 Job destruction from exit by sample, 1996–2013

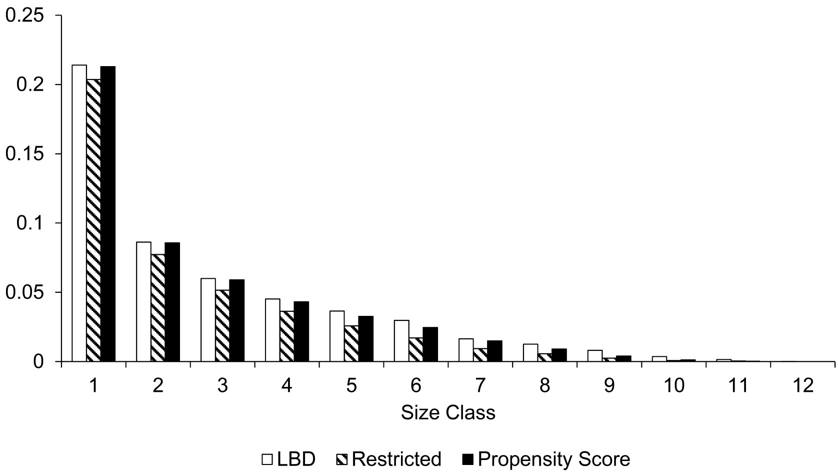

Fig. 1A.4 Job creation from births by sample, 1996–2013

Appendix B

Supplemental Results

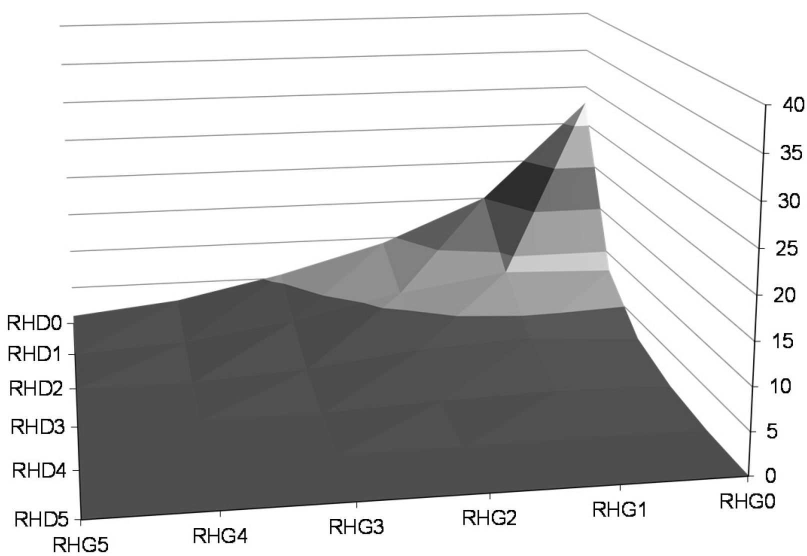

Fig. 1B.1 Percentage of RHG and RHD events for five-year-old firms (revenue weighted)

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: RHG = revenue high growth; RHD = revenue high decline. Reported shares are revenue weighted.

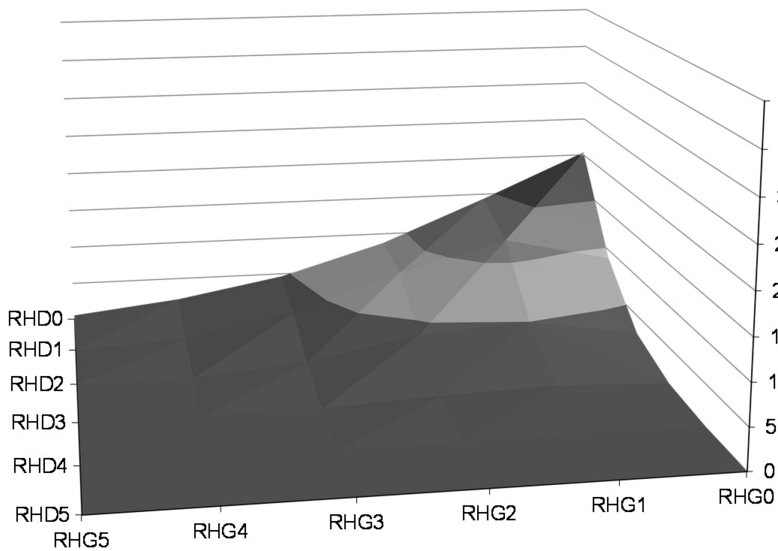

Fig. 1B.2 Percentage of EHG and EHD events for five-year-old firms (employment weighted)

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: EHG = employment high growth; EHD = employment high decline. Reported shares are employment weighted.

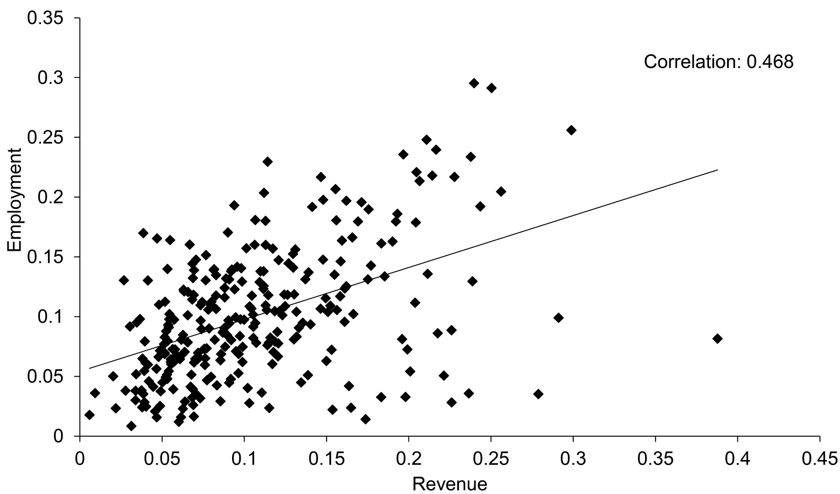

Fig. 1B.3 Industry effects revenue versus employment

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on industry effects.

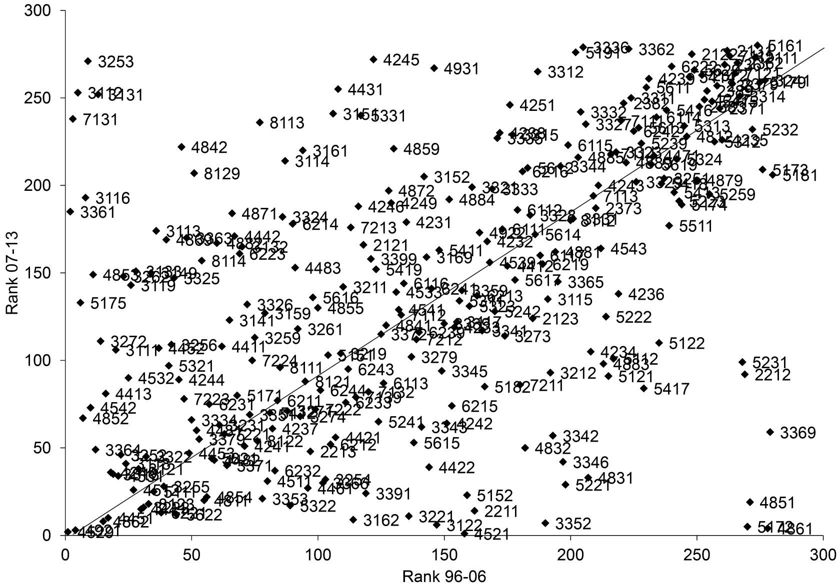

Fig. 1B.4 Industry rankings in revenue HGF, change from 96–06 to 07–13

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: The rankings for 1996–06 use the estimates from the 1996–06 period (except for 2001 and 2002) and the rankings of 2007–13 use the estimates from the 2007–13 period. Reported are estimated effects of linear probability models on industry effects.

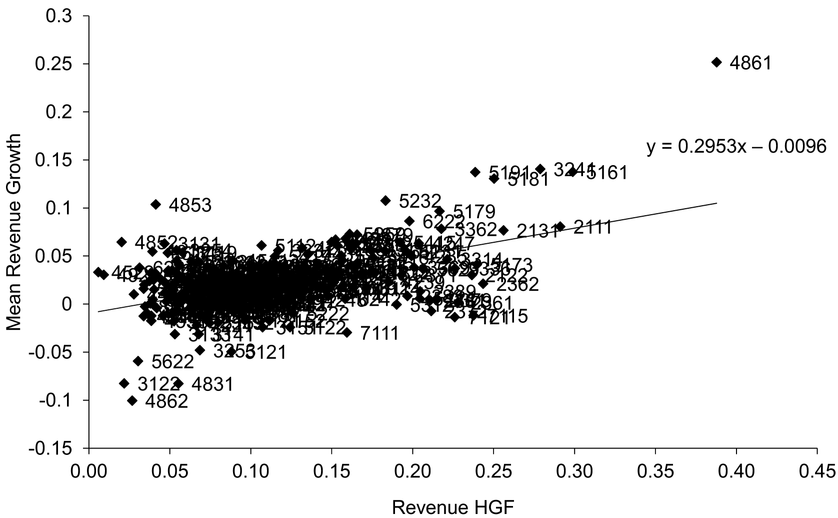

Fig. 1B.5 Revenue HGF versus revenue growth by industry

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported revenue HGF are estimated effects of linear probability models on industry effects. Mean revenue growth is revenue-weighted mean revenue growth for firms.

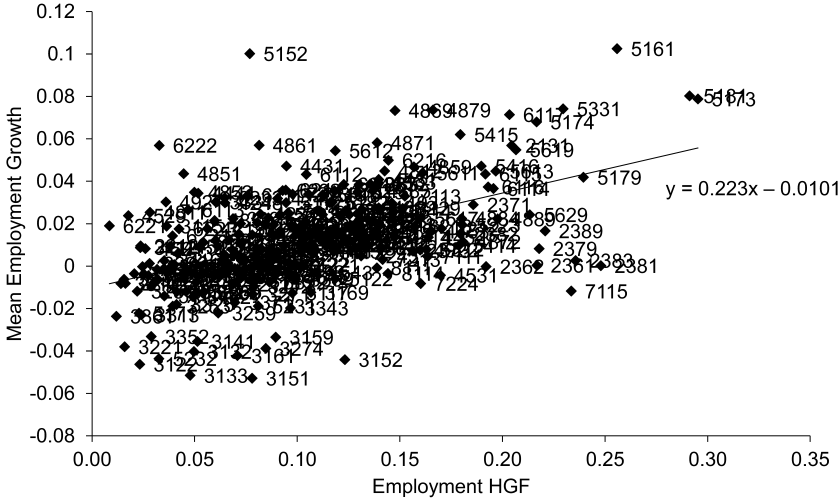

Fig. 1B.6 Employment HGF versus employment growth by industry

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported employment HGF are estimated effects of linear probability models on industry effects. Mean employment growth is employment-weighted mean employment growth for firms.

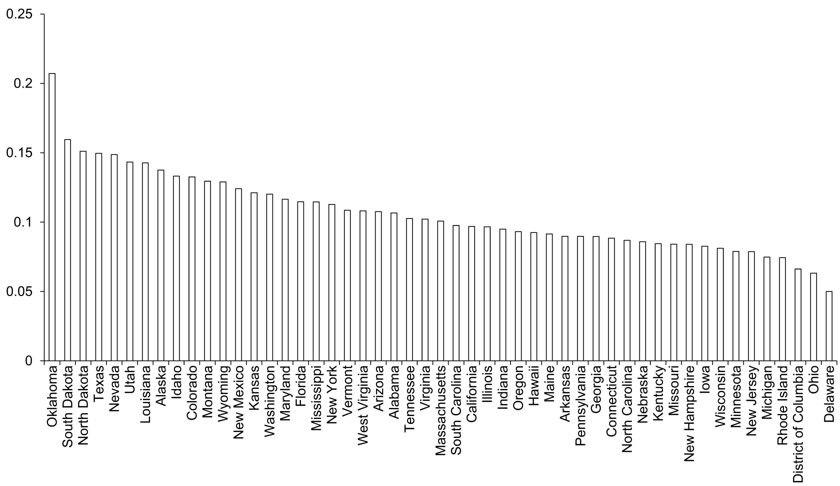

Fig. 1B.7 High-growth firm state effects (revenue)

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on industry effects.

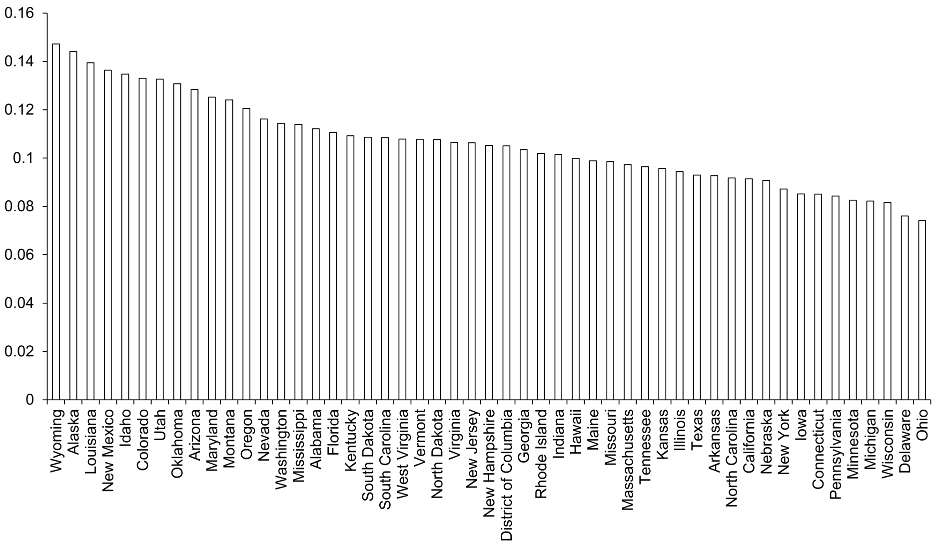

Fig. 1B.8 High-growth firm state effects (employment)

Source: Statistics computed from the revenue-enhanced LBD subsets 1998–2000 and 2003–2011.

Note: Reported are estimated effects of linear probability models on industry effects.

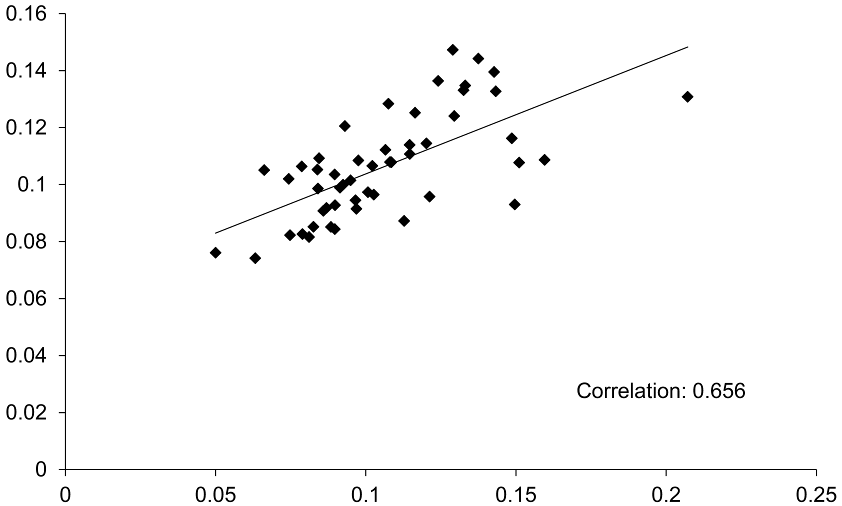

Fig. 1B.9 State effects revenue versus employment

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported are estimated effects of linear probability models on state effects.

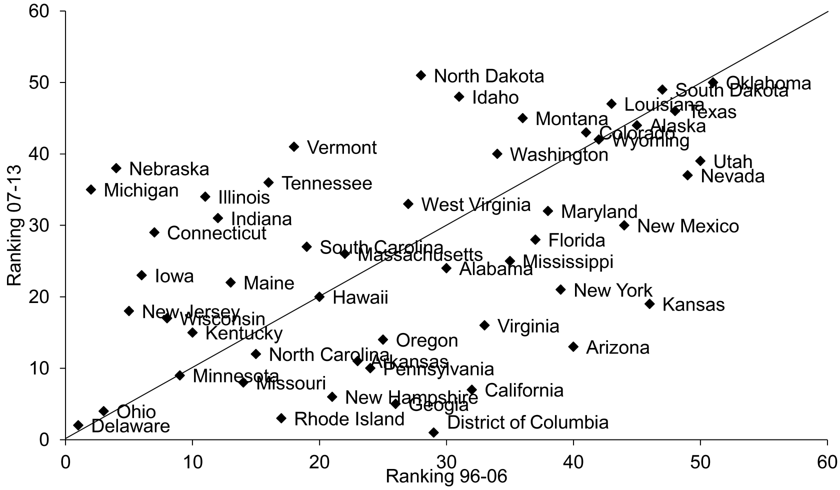

Fig. 1B.10 State rankings in revenue HGF, change from 96–06 to 07–13

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: The rankings for 1996–06 use the estimates from the 1996–06 period (except for 2001 and 2002) and the rankings of 2007–13 use the estimates from the 2007–13 period. Reported are estimated effects of linear probability models on state effects.

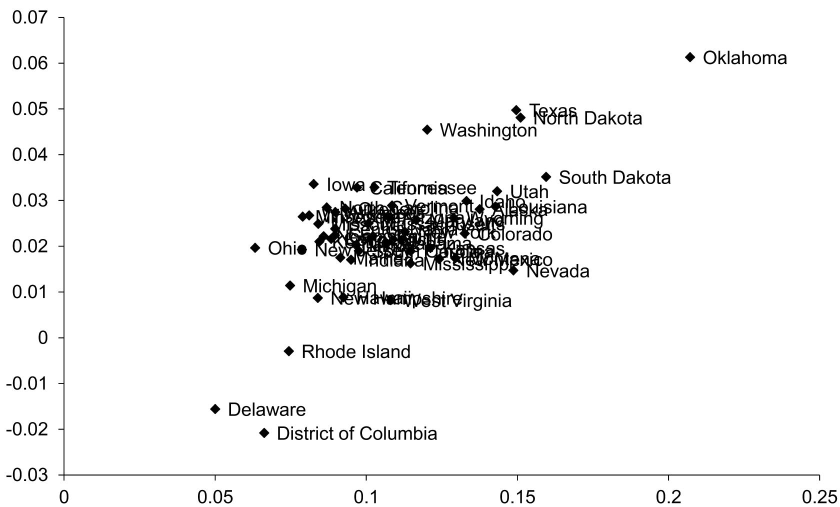

Fig. 1B.11 Revenue HGF versus revenue growth by state

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported revenue HGF are estimated effects of linear probability models on state effects. Mean revenue growth is revenue-weighted mean revenue growth for firms.

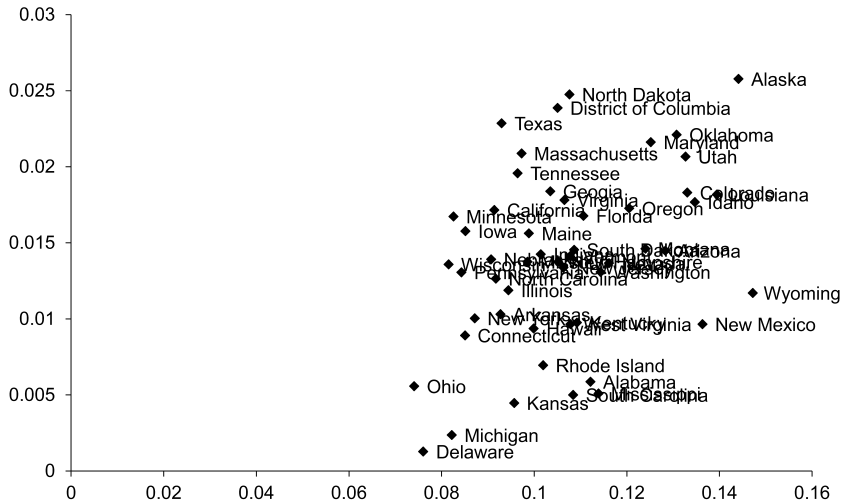

Fig. 1B.12 Employment HGF versus employment growth by state

Source: Statistics computed from the revenue-enhanced LBD subsets 1996–2000 and 2003–2013.

Note: Reported employment HGF are estimated effects of linear probability models on state effects. Mean employment growth is employment-weighted mean employment growth for firms.

Table 1B.1 Correlations of high-growth industry effects with summary measures of first, second, and third moments of industry distributions

	Rev. (GR)	Rev. (HG)	Emp. (GR)	Emp. (HG)	Rev. (HD)	Emp. (HD)	Rev. (90–10)	Emp. (90–10)	Rev. (skew)	Emp. (skew)
Rev. (GR)	1.00	0.52	0.39	0.06	–0.40	–0.16	–0.02	–0.07	0.40	0.16
Rev. (HG)		1.00	0.28	0.47	0.43	0.36	0.76	0.44	0.15	0.15
Emp. (GR)			1.00	0.52	–0.05	0.03	0.08	0.25	0.24	0.54
Emp. (HG)				1.00	0.44	0.81	0.54	0.92	–0.02	0.35
Rev. (HD)					1.00	0.53	0.81	0.52	–0.44	–0.02
Emp. (HD)						1.00	0.57	0.94	–0.15	–0.07
Rev. (90–10)							1.00	0.61	–0.24	0.05
Emp. (90–10)								1.00	–0.11	0.13
Rev. (skew)									1.00	0.16
Emp. (skew)										1.00

Note: Rev. = revenue, Emp. = employment, GR = net growth, HG = high-growth industry effect, HD = high-decline industry effect, 90–10 = activity-weighted 90–10 differential (employment weights for Emp. and revenue weights for Rev.). Skew = (90–50) –(50–10) (activity weighted).

Table 1B.2 **Correlations of high-growth state effects with summary measures of first, second, and third moments of state distributions**

	Rev. (GR)	Rev. (HG)	Emp. (GR)	Emp. (HG)	Rev. (HD)	Emp. (HD)	Rev. (90–10)	Emp. (90–10)	Rev. (skew)	Emp. (skew)
Rev. (GR)	1.00	0.64	0.46	0.21	–0.36	0.06	0.20	0.17	0.61	0.15
Rev. (HG)		1.00	0.45	0.66	0.37	0.54	0.82	0.63	0.43	0.31
Emp. (GR)			1.00	0.33	–0.05	0.02	0.20	0.22	0.28	0.61
Emp. (HG)				1.00	0.58	0.91	0.75	0.98	0.07	0.35
Rev. (HD)					1.00	0.63	0.79	0.60	–0.46	0.22
Emp. (HD)						1.00	0.73	0.97	–0.01	0.05
Rev. (90–10)							1.00	0.77	0.01	0.25
Emp. (90–10)								1.00	0.05	0.19
Rev. (skew)									1.00	0.01
Emp. (skew)										1.00

Note: Rev. = revenue, Emp. = employment, GR = net growth, HG = high-growth industry effect, HD = high-decline industry effect, 90–10 = activity-weighted 90–10 differential (employment weights for Emp. and revenue weights for Rev.). Skew = (90–50) –(50–10) (activity weighted).

References

Acemoglu, Daron, Ufuk Akcigit, Nicholas Bloom, and William R. Kerr. 2013. "Innovation, Reallocation, and Growth." NBER Working Paper no. 18993, Cambridge, MA.

Baily, Martin Neil, Charles Hulten, and David Campbell. 1992. "Productivity Dynamics in Manufacturing Plants." In *Brookings Papers on Ecnomic Activity: Microeconomics*, edited by Clifford Winston and Martin Neil Baily. Washington, DC: Brookings Institution Press.

Becker, Randy A., John Haltiwanger, Ron S. Jarmin, Shawn D. Klimek, and Daniel J. Wilson. 2006. "Micro and Macro Data Integration: The Case of Capital." In *A New Architecture for the US National Accounts*, Studies in Income and Wealth, vol. 66, edited by Dale W. Jorgenson, J. Steven Landefeld, and William D. Nordhaus, 541–609. Chicago: University of Chicago Press.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–85.

Cooper, Russell, John Haltiwanger, and Jonathan Willis. 2007. "Search Frictions: Matching Aggregate and Establishment Observations." *Journal of Monetary Economics* 54:56–78.

Davis, Steven J., and John Haltiwanger. 1990. "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications." In *NBER Macroeconomics Annual 1990*, edited by Olivier Jean Blanchard and Stanley Fischer, 123–68. Cambridge, MA: MIT Press.

———. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Quarterly Journal of Economics* 107:819–64.

Davis, Steven J., John Haltiwanger, Ron Jarmin, C. J. Krizan, Javier Miranda, Alfred Nucci, and Kristin Sandusky. 2009. "Measuring the Dynamics of Young and Small Businesses: Integrating the Employer and Nonemployer Universes." In *Producer Dynamics: New Evidence from Micro Data*, edited by Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts. Chicago: University of Chicago Press.

Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda. 2007. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." In *NBER Macroeconomics Annual 2006*, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 107–80. Cambridge, MA: MIT Press.

Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*. Cambridge, MA: MIT Press.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. "The Role of Entrepreneurship in US Job Creation and Economic Dynamism" *Journal of Economic Perspectives* 28 (3): 3–24.

———. 2015. "Where Has All the Skewness Gone? The Decline of High Growth (Young) Firms in the US" NBER Working Paper no. 21776, Cambridge, MA.

Dunne, Timothy, Mark Roberts, and Larry Samuelson. 1989. "Plant Turnover and Gross Employment Flows in the US Manufacturing Sector." *Journal of Labor Economics* 7 (1): 48–71.

Ericson, Richard, and Ariel Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62 (1): 53–82.

Fort, Teresa, John Haltiwanger, Ron S Jarmin, and Javier Miranda. 2013. "How Firms Respond to Business Cycles: The Role of the Firm Age and Firm Size." *IMF Economic Review* 61:520–59.

Foster, Lucia, John Haltiwanger, and C. J. Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*, Studies in Income and Wealth, vol. 63, edited by Charles R. Hulten, Edwin R. Dean, and Michael J. Harper, 303–72. Chicago: University of Chicago Press.

———. 2006. "Market Selection, Reallocation and Restructuring in the US Retail Trade Sector in the 1990s." *Review of Economics and Statistics* 88 (4): 748–58.

Haltiwanger, John, Ron S Jarmin, and Javier Miranda. 2013. "Who Creates Jobs? Small vs. Large vs. Young." *Review of Economics and Statistics* 2013:347–61.

Hsieh, Chang-Tai, and Peter Klenow. 2014. "The Life Cycle of Plants in India and Mexico." *Quarterly Journal of Economics*. doi: 10.1093/qje/qju014.

Hopenhayn, Hugo A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5): 1127–50.

Hurst, Erik, and Benjamin Pugsley. 2012. "What Do Small Businesses Do?" *Brookings Papers on Economic Activity* 43 (2): 73–142.

Jovanovic, B. 1982. "Selection and the Evolution of Industry." *Econometrica* 50 (3): 649–70.

Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–97.

Törnqvist, Leo, Pentti Vartia, and Yrjö O. Vartia. 1985. "How Should Relative Changes be Measured?" *American Statistician* 39 (1): 43–6.