

The Role of Entrepreneurship in US Job Creation and Economic Dynamism[†]

Ryan Decker, John Haltiwanger, Ron Jarmin, and Javier Miranda

The United States has long been viewed as having among the world's most entrepreneurial, dynamic, and flexible economies. It is often argued that this dynamism and flexibility has enabled the US economy to adapt to changing economic circumstances and recover from recessions in a robust manner. While the evidence provides broad support for this view, the outcomes of entrepreneurship are more heterogeneous than commonly appreciated and appear to be evolving in ways that could raise concern. Evidence along a number of dimensions and a variety of sources points to a US economy that is becoming less dynamic. Of particular interest are declining business startup rates and the resulting diminished role for dynamic young businesses in the economy.

We begin by describing how the concept of entrepreneurship is reflected in existing data on firm age and size. The recent addition of firm age to official statistics represents a dramatic improvement in the information available to entrepreneurship researchers. We then turn to a discussion of the role of startup firms in job creation. Business startups account for about 20 percent of US gross (total) job creation while high-growth businesses (which are disproportionately young)

■ *Ryan Decker is a PhD student in Economics, University of Maryland, College Park, Maryland. John Haltiwanger is the Dudley and Louisa Dillard Professor of Economics and Distinguished University Professor, University of Maryland, College Park, Maryland, and also a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. Ron Jarmin is Assistant Director for Research and Methodology, US Census Bureau, Washington, DC. Javier Miranda is a Principal Economist, US Census Bureau, Washington, DC. The authors' email addresses are decker@econ.umd.edu, haltiwang@econ.umd.edu, ron.s.jarmin@census.gov, and javier.miranda@census.gov. Haltiwanger is the corresponding author.*

[†]To access the data Appendix and disclosure statements, visit <http://dx.doi.org/10.1257/jep.28.3.3>

account for almost 50 percent of gross job creation. Startups and young businesses are small, the underlying reason many commentators described small businesses as the engine of US job growth prior to availability of data by firm age (for summaries, see Haltiwanger, Jarmin, and Miranda 2013; Haltiwanger 2012).

The contribution of startups and young businesses to job creation involves rich dynamics. Most business startups exit within their first ten years, and most surviving young businesses do not grow but remain small. However, a small fraction of young firms exhibit very high growth and contribute substantially to job creation. These high-growth firms make up for nearly all the job losses associated with shrinking and exiting firms within their cohort. The implication is that each entering cohort of startups makes a long-lasting contribution to net job creation.

The contribution of startups and young firms to job creation is part of an overall rapid pace of reallocation of productive resources across firms in the US economy. Young firms exhibit rich post-entry dynamics: specifically, low-productivity young firms contract and exit, while high-productivity young firms rapidly expand. In addition, young firms appear to play a critical role in innovative activity that also contributes to productivity growth (including within-firm productivity growth).

An optimal pace of business dynamics—encompassing the processes of entry, exit, expansion, and contraction—would balance the benefits of productivity and economic growth against the costs to firms and workers associated with reallocation of productive resources. It is difficult to prescribe what the optimal pace should be, but evidence accumulating from multiple datasets and methodologies suggests that the rate of business startups and the pace of employment dynamism in the US economy has fallen over recent decades and that this downward trend accelerated after 2000 (Haltiwanger, Jarmin, and Miranda 2011; Reedy and Litan 2011). A critical factor in accounting for the decline in business dynamics is a lower rate of business startups and the related decreasing role of dynamic young businesses in the economy. For example, the share of US employment accounted for by young firms has declined by almost 30 percent over the last 30 years.

These trends suggest that incentives for entrepreneurs to start new firms in the United States have diminished over time. We do not identify all the factors underlying these trends in this paper but offer some clues based on the empirical patterns for specific sectors and geographic regions. We conclude with reflections on ripe research topics in this area.

Measuring Entrepreneurship

Measuring entrepreneurship and its economic effects is difficult. Available government data on US firms do not have a specific entry for “entrepreneurs.” These data have traditionally contained information about the size of firms, and thus some observers have written or spoken as if small businesses are synonymous with entrepreneurs. However, we believe entrepreneurial activity is better represented by *new* businesses—that is, by age rather than by size. Indeed, using recently

available data with information on both firm *size* and firm *age*, Haltiwanger, Jarmin, and Miranda (2013) show that most of the job-creating prowess of small businesses is due to the contribution of startups and young businesses, which are also small at that stage in their lifecycle, as we discuss in the next section.

When thinking about “new” businesses, the distinction between new *firms* and new *establishments* is critical. Many US firms operate multiple establishments (that is, specific locations of business activity). New establishments of existing firms can take many forms, including simply replacing outdated existing establishments, thus new establishments often do not conform to standard notions of entrepreneurial behavior. For this reason, entrepreneurial research should focus on startups and young firms. It should exclude “new” businesses emerging from reorganizations such as mergers and acquisitions.

The Business Dynamics Statistics and the Longitudinal Business Database maintained by the US Census Bureau track the universe of employer firms—that is, firms in the private, nonagricultural sector with at least one employee.¹ The Longitudinal Business Database includes annual observations beginning in 1976 and currently runs through 2011. It provides information on detailed industry, location, and employment for every establishment in the private sector in a nationally comprehensive and integrated manner. With these data, it is possible to distinguish between true “startups,” new establishments of existing businesses, and “new” firms formed by combining pre-existing establishments through merger and acquisition activity. The Longitudinal Business Database identifies the relationship between firms and establishments both cross-sectionally and over time. Thus, researchers can link establishment and firm information, compute characteristics such as firm size and firm age from the establishment information, and track patterns over time.

However, even within the category of startups, we should expect to find various types of entrepreneurs. Schoar (2010) argues for distinguishing between “subsistence” entrepreneurs and “transformational” entrepreneurs. Her distinction was intended primarily for emerging economies where many entrepreneurs have limited prospects for growth, but we think this distinction is useful for the US economy as well. Subsistence entrepreneurs can be thought of as those that create small businesses that provide employment for the entrepreneur and perhaps a few others (often family members), which do not usually grow. For example, Hurst and Pugsley (2011) find that many young and small business owners in the US economy state they do not have aspirations for high growth, but rather often started businesses for nonpecuniary reasons like time flexibility or personal goals. Transformational

¹ The Business Dynamic Statistics are recently available public domain data derived from the Longitudinal Business Database, and can be found at <http://www.census.gov/ces/dataproducts/bds/>. A fuller description of these data and the measures we use are provided in the online Appendix available with this paper at <http://e-jep.org>. We note that the Longitudinal Business Database employment and job creation numbers track closely those of the County Business Patterns and Statistics of US Business programs of the US Census Bureau (Haltiwanger, Jarmin, and Miranda 2009), as they all share the Census Bureau’s Business Register as their source data.

entrepreneurs, on the other hand, create small startup businesses with the intention to innovate and grow, thus creating employment for other workers and value added for the economy. Clearly, only a subset of these transformational entrepreneurs are likely to succeed in the marketplace and grow. When people discuss the importance of entrepreneurs in job creation and productivity growth, they are envisioning transformational entrepreneurs, not subsistence entrepreneurs.

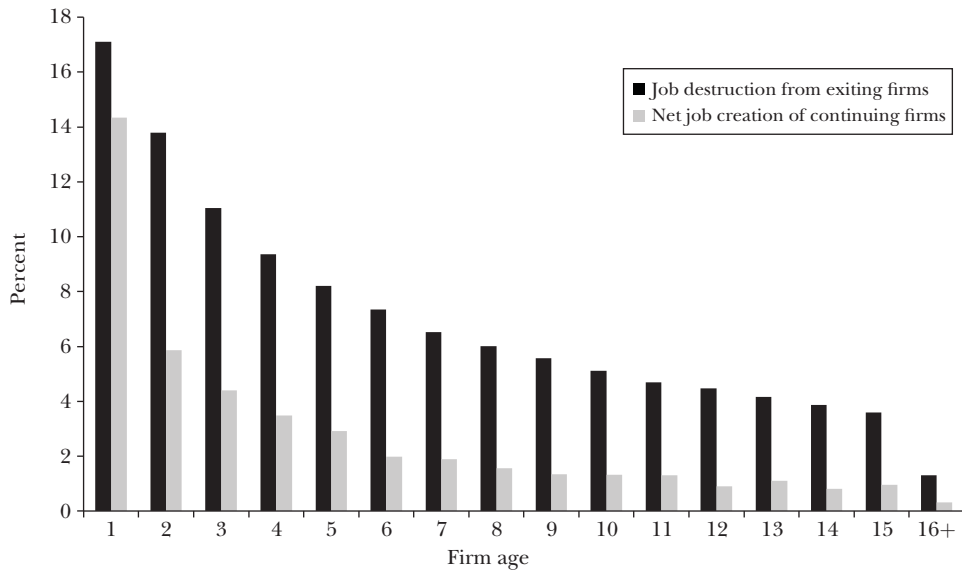
Other useful sources of data on startups include the Business Employment Dynamics from the US Bureau of Labor Statistics, which is a rich source of quarterly information on job creation and destruction. These data are more timely than the Census data but are less useful for studying firm size and firm age effects because firms are defined based on all activity operating under a single Employer Identification Number. However, many large multi-establishment firms have multiple Employer Identification Numbers. By contrast the Census Bureau uses survey and administrative data to create an enterprise-wide representation of all the establishments in the US economy that are under common ownership and control. Other important sources of data describing firm dynamics for the US economy include COMPUSTAT (which tracks publicly traded firms), the National Establishment Time Series (based on Dun and Bradstreet data), and the Kauffman Survey tracking a sample of recent entrants.

Startups and US Jobs

Startups and young businesses clearly play an important role in job creation. Between 1980 and 2010, the gross number of jobs created annually by all establishments averaged about 18 percent of the workforce—an average of 16.3 million jobs per year—according to our calculations from the Business Dynamics Statistics. About one-sixth of this amount, an average of 2.9 million jobs annually, can be traced to new firms, and another one-sixth can be traced to new establishments of existing firms. The net job creation statistics are even more striking. For new firms (that is, those with age equal to zero in the Business Dynamics Statistics), “net” and “gross” job creation are the same, because they have no previous jobs to lose, and so their net job creation is also 2.9 million jobs per year. Over these 30 years, average net job creation in the entire US private sector was approximately 1.4 million jobs per year. The implication is that cohorts of firms aged one year or older typically exhibit net job declines.

Taken at face value, these statistics might be interpreted as implying that all net job creation is due to startups. This interpretation is misleading for two reasons. First, gross job creation is an order of magnitude greater than net job creation in any given period. This implies that in any given period there are many different groupings of growing firms that could be used to account for net job creation. Second, it is important to examine the post-entry dynamics of startups. If, for example, all startups failed after a short period of time, then startups would not be making any long-lasting contribution to net job creation.

Figure 1
Up or Out Dynamics for Young Firms



Source: Annual averages of statistics computed from the Longitudinal Business Database from 1992–2011. Notes: Figure 1 shows patterns of net employment growth of continuing firms and job destruction from firm exit for firms age 1 and older. Startups have firm age equal to zero, so this figure reports on the post-entry dynamics of firms. (See footnotes 2 and 3 and online Appendix for details.)

To explore post-entry dynamics, we need to track firm growth and survival as a function of firm age. We rely here on the methodology developed by Davis, Haltiwanger, Jarmin, and Miranda (2007) and Haltiwanger, Jarmin, and Miranda (2013). Firm age is measured using the age of the oldest establishment in the firm. For startups, all of the establishments of the new organization are entrants so firm age is zero. In this methodology, continuing firms age “naturally,” one year at a time, as long as the organization stays in existence.² Consistent with this approach, firm exits represent legal entities that cease to exist and in which all of their associated establishments shut down. Thus, firm exits do not reflect legal entities that cease through organizational change or buyout activity and where at least some establishments continue operation in subsequent years.

Using this approach, Figure 1 shows patterns of net employment growth for continuing firms and job destruction from firm exits for firms age 1 and older.³

² As part of this same methodology for assigning firm age, Haltiwanger, Jarmin, and Miranda (2013) develop a method for capturing firm growth that focuses on organic growth rather than growth from merger and acquisition activity and other related changes in organization. We use that methodology here, and explain it in more detail in the Appendix available with this paper at <http://e-jep.org>.

³ Reported statistics in this figure reflect the net employment growth for the cell using the growth rate methodology developed by Davis, Haltiwanger, and Schuh (1996). The growth rate concept at any level

Recall that startups have firm age equal to zero, and so Figure 1 provides insights on the post-entry dynamics of firms. The statistics are based on tabulations of pooled data from 1992–2011 from the Longitudinal Business Database. Conditional on survival, younger firms have much higher rates of job growth than more mature firms. Indeed, even the typical 10 year-old firm has net growth that is 1.4 percentage points higher than the typical 16+ year-old firm. The exit dynamics look different as well. Young firms have a substantially higher exit rate (job destruction from exit is an employment-weighted exit rate): 50 percent of the jobs generated by an entering cohort of firms will have been lost to exits by age five.

Thus, young firms exhibit a strong “up or out” dynamic (Haltiwanger, Jarmin, and Miranda 2013). For any given cohort, jobs lost due to the high failure rate of young firms are almost offset by the growth of the surviving firms. Five years after the entry of a typical cohort, total employment is about 80 percent of the original employment contribution of the cohort—in spite of losing about 50 percent of the original employment to business exits. In this sense, cohorts of startups do make a long-lasting contribution to net job creation. But the high volatility of young firms highlights that the contribution to net job creation from startups can’t be simply understood from the immediate contribution upon starting up.

The high mean net employment growth of surviving young firms masks enormous heterogeneity among young surviving firms. Figure 2A shows the 90th, 50th, and 10th percentiles net job growth of surviving firms by firm age. Figure 2B shows the mean and median of net employment growth of surviving firms by firm age. The mean is the aggregate net growth rate for each age group. The 90th, 10th, and median are the employment-weighted percentiles, so they reflect the employment-weighted distribution underlying the aggregate mean. Young firms have very high dispersion of growth and also very high skewness. The skewness is seen in the relative magnitudes of the 90th and 10th percentiles where the growth rates of younger firms are much more skewed to the right (positive) compared to more mature firms. This accounts for the high mean growth rate of young firms relative to older firms that is evident in the lower panel of Figure 2. In short, the typical young firm (as captured by the median) exhibits little or no growth even conditional on survival (many of these are presumably the “subsistence” entrepreneurs discussed earlier), however, among all the young firms, a few do exhibit very high rates of growth which yields a high mean growth rate.

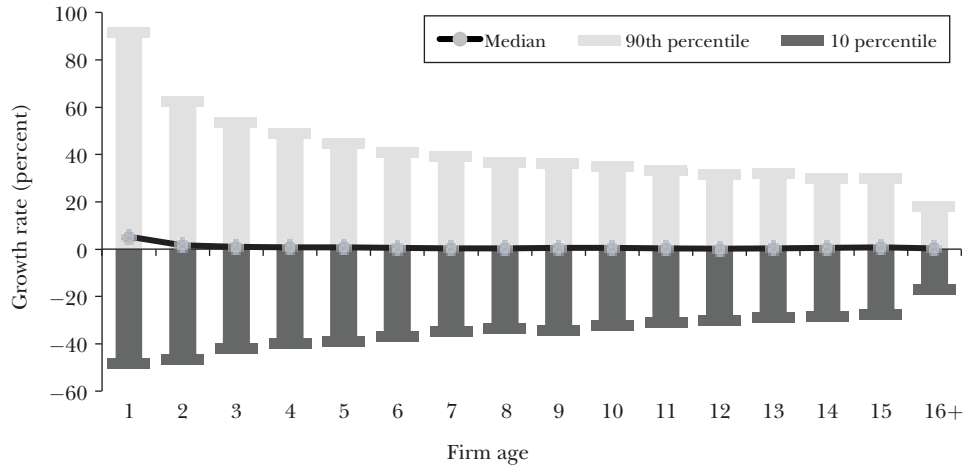
The skewed right tail of the growth rate distribution of young firms is important for understanding the contribution of startups and young firms to overall job creation. Startups account for less than 10 percent of firms and about 20 percent of firm-level gross job creation. But the contribution of startups to job creation does not stop at entry, at least for some firms. High-growth firms—defined here as firms expanding

of aggregation is based on the change in the number of jobs for a cell from period $t - 1$ to t divided by the average number of jobs in periods $t - 1$ and t . This growth rate measure was developed originally by Törnqvist, Vartia, and Vartia (1985). Like log changes, this growth rate measure is symmetric but it has the advantage of accommodating zeroes (entry and exit). It is a second order Taylor approximation to the log first difference. See the online Appendix with this paper at <http://e-jep.org> for more details.

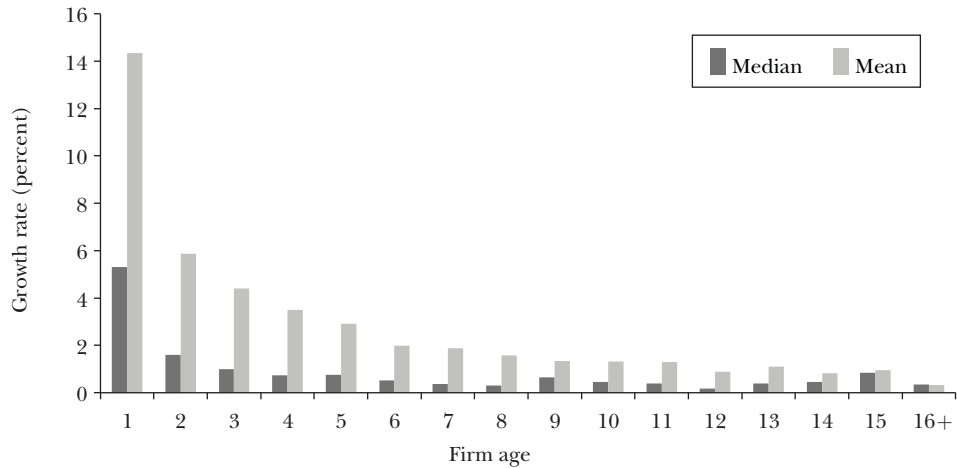
Figure 2

Net Employment Growth and Growth Rates for Surviving Firms

A: 90th, Median, and 10th Percentiles of Net Employment Growth for Surviving Firms



B: Mean and Median Net Employment Growth Rates for Surviving Firms



Notes: Annual averages of statistics computed from Longitudinal Business Database, from 1992–2011. The 90th, 10th, and median are all based on the employment-weighted firm level growth rate distribution for each firm age cell. The mean is the aggregate net growth rate of the firm age cell which is equivalent to the employment-weighted average of the firm level net growth rate in each cell.

their employment by more than 25 percent per year—account for about 15 percent of firms and 50 percent of firm-level gross job creation. Together, startups and high-growth firms (which are disproportionately young as seen in Figure 2) account for about 70 percent of firm-level gross job creation in a typical year. Balancing this positive contribution is the sharp job loss that occurs for many firms in the first several

Table 1

Estimates of Net Growth for Continuing Firms by Firm Size and Firm Age
(in percent)

Firm age	Firm size (number of employees)								
	1–4	5–9	10–19	20–49	50–99	100–249	250–499	500+	All sizes
1–2	6.7	9.0	11.5	12.9	14.9	14.5	14.3	17.9	11.8
3–4	2.0	2.5	4.4	5.6	7.4	6.8	6.7	10.1	4.5
5–6	–0.5	–0.2	1.8	2.9	4.6	3.9	3.7	7.0	3.0
7–8	–2.0	–1.2	0.8	2.1	3.7	3.0	2.9	6.0	2.4
9–10	–3.7	–2.4	–0.2	0.9	2.5	1.8	1.5	4.5	1.8
11–12	–2.1	–0.9	0.9	2.4	3.8	3.0	2.5	5.5	1.9
13–15	–0.8	0.0	3.0	3.5	4.4	4.0	3.2	5.7	1.8
16+	–2.4	–1.7	–0.9	–0.1	0.6	1.3	1.6	0.9	0.7
All ages	0.2	1.2	1.7	2.1	2.5	2.9	3.1	1.4	

Notes: Tabulations from the US Census Bureau. Details of methodology and results can be found in the online appendix of Haltiwanger, Jarmin, and Miranda (2013). The analysis covers the period 1992–2005.

years after startup entry. Overall, the evidence shows that most startups fail, and most that do survive do not grow. But among the surviving startups are high-growth firms that contribute disproportionately to job growth. These high-growth young firms yield the long-lasting contribution of startups to net job creation.

To understand the population of startups and their dynamics, it helps to realize that their size distribution is quite different from that of the overall distribution of firms. Startups are small; more than 90 percent of all startups have fewer than 20 employees, and these small startups account for about half of all startup employment. For purposes of comparison, about 50 percent of employment in the US private sector is accounted for by the less than 1 percent of US firms with more than 500 employees. Large startups are almost nonexistent and account for a very small share of startup activity; however, many small firms are old, accounting for about 30 percent of employment. This image is enriched by examining the average net employment growth rate by firm age and firm size class in Table 1. Table 1 focuses on continuing firms, to help link back to the evidence in Figure 1 regarding the high average net growth rate for young surviving firms. Average net job growth falls monotonically with firm age for all firm size groups. However, the average net growth for young firms is substantially higher for firms that are larger than 20 employees. Such patterns highlight that rapid employment growth among young surviving firms is especially present among larger—or at least not micro-sized—young firms.

Startups, Reallocation, and Productivity

The high pace of labor market reallocation with a critical role for startups and young firms raises a question: What is driving this reallocation? Considerable

evidence suggests that entry, exit, expansion, and contraction of firms are closely related to measures of productivity and profitability (for surveys and summaries, see Syverson 2011; Foster, Haltiwanger, and Krizan 2001; Bartelsman and Doms 2000).

Producers, even those within the same industry, display enormous differences in measured productivity and profitability. For example, within the same four-digit US manufacturing industry, the establishment at the 90th percentile of the productivity distribution produces almost twice as much output (measured by real revenue) with the same measured inputs as the 10th percentile establishment (Syverson 2004, 2011). These large differences are persistent, but not permanent, indicating that firms are subject to ongoing idiosyncratic productivity and profitability shocks.

These large differences in productivity and profitability are connected to the growth and survival dynamics of firms. As such, they underlie the reallocation dynamics we have discussed and documented above. For example, Foster, Grim, and Haltiwanger (2013) estimate, using data for the entire US manufacturing sector for the period 1980–2011, that the probability of exit for the establishment at the 90th percentile of the productivity distribution within an industry is 4 percentage points lower than the establishment at the 10th percentile (the average annual exit rate for manufacturing establishments is 8 percent per year). Conditional on survival and holding initial size constant, estimates from the same data imply that an establishment at the 90th percentile of productivity has a growth rate that is 3 percentage points higher than a plant at the 10th percentile. Finally, estimates from Syverson (2011) also show that entrants tend to start with productivity that is on average similar to incumbents. These estimated effects of productivity on growth and survival are very similar to those found in the literature. As stated in that paper: “[A] robust finding in the literature—virtually invariant to country, time period or industry—is that higher productivity producers are more likely to survive than their less efficient counterparts.”⁴

When focusing on the relationship between reallocation of resources across the economy and firm-level measures of productivity, several issues arise. One issue is that much of the micro empirical literature in this area focuses on establishment-level rather than firm-level measures of productivity. Part of the reason for this focus reflects the complexity of measuring firm-level productivity for large, mature firms operating across several sectors. However, when the research looks just at establishments owned by young firms, as in Foster, Grim, and Haltiwanger (2013), the marginal effects of revenue-based measures of productivity on growth and survival (discussed above) are, if anything, larger. This is consistent with the up-and-out dynamics of young firms discussed earlier.

Another issue is that much of the literature in this area measures total factor productivity in terms of revenue. This choice is primarily due to data limitations as most sources of firm- and establishment-level data that are used to measure productivity do not include micro level measures of output and input prices. Instead, real

⁴ We discuss the robustness of the findings in the literature further in section II of the online Appendix available with this paper at <http://ejep.org>.

output and real inputs are measured as establishment- or firm-level revenues and costs divided by a price deflator. Such revenue-based measures of productivity will thus reflect some combination of the technical efficiency with which inputs are translated into outputs, together with demand and cost effects on prices.

Revenue-based measures of productivity may be a reasonable approximation for measures that adjust for price dispersion across producers. When Foster, Haltiwanger, and Syverson (2008) compare the revenue-based measures of total factor productivity with total factor productivity measured in physical quantities for a sample of manufacturing industries, they find that the correlation between these two measures of productivity is .75, which suggests that broad findings in the literature based on measures of revenue productivity would hold up reasonably well with measures of productivity based on quantities.⁵ The industries for which physical quantity data are available in the US economy are limited, so there is some question as to the wider applicability of these findings across the whole economy. However, evidence for Colombia (where establishment-level price indices are available for all manufacturing establishments) suggests these patterns are robust for a much wider range of industries (Eslava, Haltiwanger, Kugler, and Kugler 2004, 2013).

To be clear, productivity growth in an economy is not only a matter of more-productive entering firms replacing less-productive exiting firms. A common finding in the literature about productivity growth in manufacturing is that about 60 percent of industry-level productivity growth happens within existing establishments and the rest comes from reallocation of productive resources resulting from entry, exit, and the expansion and contraction of existing establishments. For example, Foster, Haltiwanger, and Syverson (2008) find that entrants and young establishments have slightly higher total factor productivity (measured in quantity terms) than more mature incumbents, but the entrants have substantially higher productivity than exiting establishments. In their data, 35 percent of industry-level productivity growth is accounted for by net entry. However, their study looks over a five-year time period, and thus some of the 60 percent of productivity growth happening within existing establishments occurs in young firms. Foster, Haltiwanger, and Krizan (2001, 2006) provide evidence of such selection and learning dynamics and show that within-plant productivity growth is more rapid for surviving young establishments than more mature establishments.⁶

In sectors like retail trade, the evidence suggests that productivity growth within establishments is less important except in the case of young establishments

⁵ In section II of the online Appendix available with this paper at <http://ejep.org>, we discuss the evidence that shows the relationship between revenue productivity and growth and survival are very similar to the relationship between physical productivity and growth and survival.

⁶ Foster, Haltiwanger, and Syverson (2008) also find that the role of net entry is substantially larger using total factor productivity measured in quantity terms as opposed to measuring it in revenue terms. The net entry component accounts for 35 percent of productivity growth with quantity-based productivity and 24 percent of productivity growth with revenue-based productivity. The difference must lie, of course, in price factors that differ between entrants and younger incumbents on the one hand and more mature incumbents on the other. They find that entrants and young establishments face substantially lower prices, which may reflect the challenges of becoming established and building reputation in a market.

that show productivity growth in their early years. Instead, most of the labor productivity growth in this sector has been attributed to net entry (Foster, Haltiwanger, and Krizan 2006). In many cases, existing firms improve productivity in retail trade primarily through adding new, more-productive retail locations rather than expanding existing establishments. Moreover, much of the exit of low-productivity retail establishments in the US economy has been dominated by the exit of “mom and pop” single-establishment firms. Doms, Jarmin, and Klimek (2004) find that more-intensive investment in information technology by large producers is related to both differences in the level of productivity across establishments and differences in the within-establishment growth rates of productivity, thus highlighting the link between technology, productivity, and business dynamics.

An important related question is the role of startups and young firms in innovative activity. Among innovative firms, young and small firms have higher innovation intensities than mature firms as measured by the ratio of research and development spending to sales (Acemoglu, Akcigit, Bloom, and Kerr 2013).

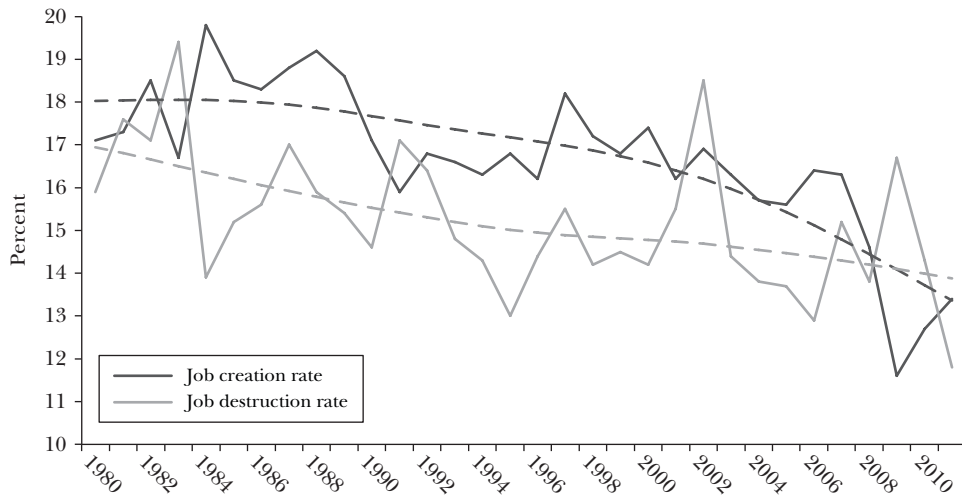
To sum up, while the evidence that startup firms are a powerful driver of job growth is quite clear, the evidence is less definitive for the contribution of startup firms to overall productivity growth. In addition to the limitations in relation to the studies discussed above, micro-based studies of productivity beyond the manufacturing and retail trade sectors are scarce. Measuring productivity in other industries is difficult, given the limitations on data on inputs other than labor in nonmanufacturing sectors. Still, the evidence that does exist for manufacturing and retail trade tends to support the existence of an up-or-out dynamic with high productivity (and high-profitability) young firms growing and low-productivity firms contracting and exiting. Moreover, young firms have a high innovation intensity, which suggests that young firms may be disproportionately important in terms of growth and productivity.

The Decline in the Startup Rate and Business Dynamism

The US economy displays a rapid pace of gross job creation and destruction, as depicted in Figure 3.⁷ Even in 2009, when the economy experienced a large net employment decline due to the recession, expanding and new businesses added jobs at a rate equal to 12.4 percent of total employment. In absolute terms, the

⁷ The online Appendix available with this paper at <http://e-jep.org> offers formal definitions of job creation and destruction and other related measures of dynamism in the job market. In this literature, “job reallocation” is the sum of job creation and destruction—it is a summary measure of all the changes in the location of jobs across producers. We also use that measure in the analysis in this section. Another summary measure often used in the literature is the “excess reallocation rate,” which is job reallocation less the absolute value of net growth. The excess measure captures the reallocation over and above that needed to accommodate net growth. For the sake of brevity, we do not show patterns by excess reallocation here, but note that our patterns of declining trends carry over to excess reallocation. Davis et al. (2007) show patterns by excess reallocation.

Figure 3
US Annual Job Creation and Destruction Rates, 1980–2011



Source: Author calculations from the US Census Bureau's Business Dynamics Statistics.

Notes: The filter is Hodrick–Prescott with multiplier 400. The vertical axis does not begin at zero.

US private sector created more than 14 million new jobs between March 2008 and March 2009. The nature and pace of reallocation does change over the business cycle: specifically, job creation and job destruction tend to move in opposite directions during expansions and contractions. As Figure 3 shows, the decline in job creation during the 2007–2009 period was especially large.

But abstracting from the business cycle, the rates of job creation and destruction exhibit a general downward trend during the last few decades. Figure 3 illustrates this point by also showing smoothed trends. In particular, the job creation rate averaged 18.9 percent in the late 1980s and decreased in what appears to be a roughly stepwise pattern following recessions to an average of 15.8 percent in the 2004–2006 period just before the Great Recession. Meanwhile, the job destruction rate fell from an average of 16.1 percent in the late 1980s to an average of 13.4 percent in the 2004–2006 period. Moreover, the decline in the pace of job creation and destruction appears to have accelerated since about 2000.⁸

These declining trends appear in a wide range of data and settings. Using multiple measures of business dynamics from the Longitudinal Business Database, Davis, Haltiwanger, Jarmin, and Miranda (2007) found that the trend decline

⁸ The annualized decline increases in magnitude from -0.15 to -0.22 of a percentage point for job creation and from -0.14 to -0.19 of a percentage point for job destruction.

is present in both firm-level and establishment-level measures.⁹ Similarly, Davis, Faberman, Haltiwanger, Jarmin, and Miranda (2010), using the Bureau of Labor Statistics' Business Employment Dynamics data, show a declining pace of job flows over time. They find that the declining trend in the pace of job destruction is closely linked to the secular decline in the inflow rate to unemployment (at both the national and sectoral level). Davis, Faberman, and Haltiwanger (2012) show that this declining pace of job flows is matched by a declining pace of worker flows in the Job Openings and Labor Turnover (JOLTS) data. Specifically, the decline in the pace of hires exceeds that of the decline in job creation, and similarly the decline in the pace of separations exceeds that of the decline in job destruction. In other words, there is a decline in the pace of excess worker reallocation or worker churn. Lazear and Spletzer (2012) report similar findings using the JOLTS and BED data. Hyatt and Spletzer (2013) find similar patterns using the worker and job flows data from the Quarterly Workforce Indicators.

Job reallocation measures the change in the allocation of jobs across producers. In contrast, excess worker reallocation measures the change in the allocation of workers over a given set of jobs. Total worker reallocation is the sum of job reallocation and excess worker reallocation. It is striking that there have been secular declines in these related but distinct components of overall worker reallocation.

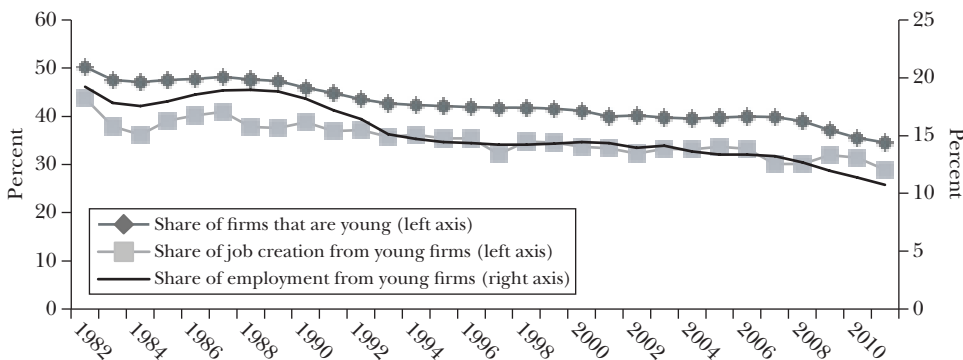
Secular Shifts: The Changing Distribution of Firm Age

A decline in the startup rate is one of the factors affecting the pace of reallocation in labor markets. The firm startup rate is measured by the number of new firms divided by the total number of firms. Our calculations based on the Business Dynamic Statistics data show that the annual startup rate declined from an average of 12.0 percent in the late 1980s to an average of 10.6 percent just before the Great Recession, when it plummeted below 8 percent. We also find that the startup rate has declined in all major sectors. We note, however, that in high-tech sectors (Haltiwanger, Hathaway, and Miranda 2014), the startup rate only began to decline in the post 2000 period.

Meanwhile, the average size of startups, as measured by employment, has either remained approximately the same over this time period as measured by the Census Bureau's Business Dynamics Statistics data (Haltiwanger, Jarmin, and Miranda 2013) or has declined as measured by the Bureau of Labor Statistics' Business Employment

⁹ Davis et al. (2007) also show that the trend decline is present in within-firm and within-establishment measures of volatility as well as cross sectional dispersion measures such as the pace of job creation and destruction. The decline in the pace of overall firm volatility does mask an increase in the pace of firm volatility among publicly traded firms through 2000 (Comin and Philippon 2005). Davis et al. (2007) confirm the Comin and Philippon findings using data that have both privately held and publicly traded firms. They show that the decline in the pace of business volatility among privately held firms overwhelms the rise in firm volatility for publicly traded firms. Their findings suggest that the difference in patterns between publicly and privately held firms through 2000 primarily reflects a change in the composition of publicly held firms. In particular, the 1980s and 1990s cohorts of new publicly traded firms are younger when going public and also grow more rapidly after going public than earlier cohorts.

Figure 4

Declining Share of Activity from Young Firms (Firms Age 5 or Less)

Source: Author calculations from the US Census Bureau's Business Dynamics Statistics.

Note: Employment shares in each period based on the average of employment in period $t - 1$ and t (the denominator of the Davis, Haltiwanger, and Schuh (1996) growth rate).

Dynamics (Reedy and Litan 2011; Choi and Spletzer 2012). Either way, the lower startup rate is not being offset by a larger size of startup firms.

A consequence of the declining startup rate and flat or declining startup size is that the share of young firms in the economy, and the share of activity for which they account, is declining. Figure 4 shows that firms aged five years or less made up about 47 percent of all firms in the late 1980s, but this number declined to 39 percent of all firms before the start of the Great Recession, and has declined further since then. Similarly, the share of employment at firms less than five years of age declined from an average of 18.9 percent in the late 1980s to an average of 13.4 percent at the cyclical peak before the Great Recession. Finally, firms that were five years old or less contributed 39 percent of all new jobs in the late 1980s, but about 33 percent of all new jobs before the Great Recession. Because rates of job creation and destruction are much higher at young firms, these changes will contribute to the decline in both job creation and job destruction rates (Fort, Haltiwanger, Jarmin, and Miranda 2013).

To quantify the contribution of the changing age structure of firms on the change in job reallocation, we apply a standard shift-share decomposition for seven unique firm age groups (ages 0 through 5, and then 6 or more). To abstract from business cycle effects, we focus on the long-run change between two cyclical peaks. Specifically, we measure the difference between three-year averages of job flows for the periods of 1987–89 and 2004–06. Using a three-year average allows us to avoid issues that can arise in looking at particular years.

Our shift-share approach decomposes the change in job flows between 1987–89 and 2004–06 into three components: 1) the contribution from within-age group changes in job flows holding constant the employment-weighted firm age distribution at the initial levels; 2) the contribution from changes in the employment

distribution across firm age groups holding constant the job flows by firm age at the initial levels; and 3) the contribution of a cross term relating changes in shares with changes in flows. Doing this, we find that the shifting age composition of firms towards older businesses accounts for 32 percent of the observed decline in job creation, 20 percent of the decline in job destruction, and 26 percent of the decline in job reallocation.

Secular Shifts: The Changing Industrial Structure of Firm Activity

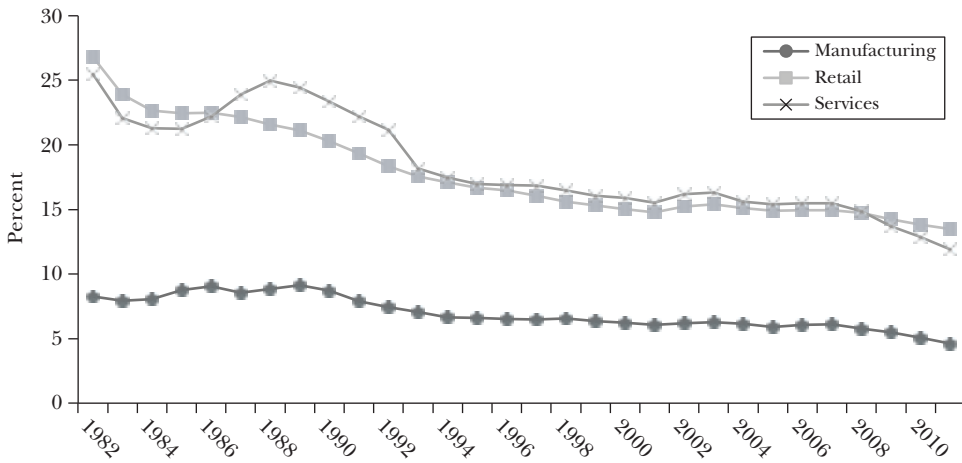
Unlike the changing age distribution of firms, shifts in the industrial composition of the economy would be expected to lead to increases in the pace of job reallocation. For example, the sectoral shift from manufacturing (relatively low pace of reallocation) to services (relatively high pace of reallocation) suggests that we should have expected an increase in the overall pace of business dynamics rather than a decline. In 1980, manufacturing accounted for 28 percent of all US jobs, while retail and services accounted for 24 percent. By 2011, manufacturing accounted for only 11 percent of all jobs, while retail and services accounted for 43 percent (using the broad sector definitions in the Business Dynamic Statistics).

To quantify the contribution shifts in the industry structure of the US economy have on job reallocation, we use the same type of shift-share analysis, but now look at 282 unique four-digit NAICS (2002) industries.¹⁰ We find that the changing industrial share of firm activity indeed dampened the decline in the pace of job reallocation. If the industrial structure had remained the same from 1987–89 to 2004–06, the decline in job creation rates would have been greater by 21 percent, the decline in job destruction would have been greater by 3 percent, and the decline in job reallocation would have been greater by 12 percent.

Given the offsetting effect of the change in industrial composition, we considered a richer shift-share decomposition where we take into account the changes in industrial composition, the age structure, and other firm characteristics that have been associated with systematic differences in the pace of reallocation. In particular, a common finding is that establishments belonging to large firms are less volatile than those that belong to small firms. We consider a shift share decomposition with fully interacted cells defined by 282 unique four-digit NAICS (2002) industries, seven unique firm age groups (ages 0 through 5, and 6+), and eight firm-size groups (1–9, 10–19, 20–49, 50–99, 100–249, 250–499, 500–999, and 1000+ employees). In combination, the changing composition by industry, firm age, and firm size account for 9 percent of the decline in job creation, 18 percent of the decline in job destruction, and 14 percent of the decline in job reallocation. These findings suggest that unexplained within-group variation must account for most of the observed decline in job flows.

¹⁰ We thank Teresa Fort for the development of a methodology that reclassifies all establishments in the Longitudinal Business Database to a consistent NAICS (2002) industry classification system. See Fort (2013) for details. Having a consistent classification system for our entire panel is critical for our analysis. These consistent NAICS codes have not yet been incorporated into the Business Dynamics Statistics.

Figure 5

Share of Employment from Young Firms (Firms Age 5 or Less), Selected Sectors

Source: Author calculations from the US Census Bureau's Business Dynamics Statistics.

Notes: Sector definitions are on an Standard Industrial Classification basis. Employment shares in each period are based on the average of employment in period $t - 1$ and t (the denominator of the Davis, Haltiwanger, and Schuh (DHS) growth rate).

Looking for Clues in Sectoral and Geographic Changes in Startups

Insights into the unexplained within-group declines in job flows can be gained from exploring entrepreneurial activity within groups. Figure 5 shows the share of total employment for young firms (five years or less in age) for the retail, services, and manufacturing sectors. The share of employment in young firms has declined in all three sectors, suggesting that factors that are not sector-specific are causing the decline in entrepreneurial activity. Consistent with this view, we find this decline in the share of employment at young firms in every major sector. Figure 5 also shows that the drop is much larger for young retail and services firms than it is for young manufacturing firms. For whatever reason, incentives to start new businesses appear to be declining in all sectors, but disproportionately so in certain sectors such as retail trade, and this has contributed substantially to the declines in the pace of business dynamics.

The especially large decline in entrepreneurial activity in retail trade is consistent with the fundamental transformations ongoing in this sector over many decades. The expansion of “big box” retailers and, more generally, large national firms has dramatically changed the characteristics of the firms and establishments in the industry. Jarmin, Klimek, and Miranda (2005) report that the share of US retail activity accounted for by single-establishment (“mom-and-pop”) firms fell from 70 percent in 1948 to 60 percent in 1967, and further still to 39 percent in 1997. If the decline in reallocation were confined to retail trade only, then we might guess that changing retail business models play a dominant role in accounting for

observed decreases in dynamism. But since we observe decreased dynamism in virtually all sectors, additional factors must be at work.

Additional clues might be found by exploring the changing pattern of entrepreneurial activity by US regions and states. We find that the share of employment by young firms has declined in all 50 states. The decline is more pronounced in states in the South and West compared to the declines in the Midwest and the Northeast. Such geographic variation can potentially capture different changes in the business climate, as states differ in regulations across a range of dimensions including occupational licensing requirements, banking regulations, tax burden for businesses and households, employment protection regulations, minimum wages, and others. However, we have found that, for example, states with business climates as different as California and Texas exhibit large and similar declines in entrepreneurial activity. More generally, since all states exhibit a decline in entrepreneurial activity, this suggests that factors other than state-specific business climate effects must be at work.

Understanding the Causes and Consequences of the Decline in Entrepreneurial Activity

We do not yet fully understand the causes of the decline in indicators of business dynamism and entrepreneurship, nor in turn, their consequences. Improving our understanding of the causes and consequences should be a high priority. A straightforward way of considering these issues is to ask whether either the structure of shocks affecting firms or changes in how firms are responding to these shocks can account for the findings. While both possibilities deserve inquiry in future research, we think the latter line of inquiry is likely to prove of greater value. Our prior is that the variance of idiosyncratic shocks affecting firms due to factors like technological change and globalization has, if anything, increased rather than decreased over time.¹¹ To the extent that such shocks have increased, then the decline in the pace of labor market reallocation is an even greater puzzle.

If the magnitude of economic shocks has not changed in a way that would tend to reduce the amount of labor market reallocation, then any such declines are driven by the way firms respond to shocks. One possibility is that the business climate, broadly defined, has changed in ways that impede job reallocation—that is to say, by impeding entry, exit, expansion, and contraction. Moreover, if the cause of the decline is an increase in the costs of adjustment on one or more of these margins, this can imply adverse consequences for growth, productivity, and welfare. The seminal work of Hopenhayn and Rogerson (1993) offers guidance here. They show that if an economy experiences an increase in adjustment costs for job destruction (for example, due to increased regulation), then not only will there be a decline in job destruction but also a decline in job creation (including a decline in startups)

¹¹ It is the trend in the pace of second moment shocks impacting firms that matters here and not so much the pace of first moment shocks. The period of the Great Moderation of business cycles is within our sample period, but that should have primarily impacted aggregate volatility and not firm-level volatility.

and ultimately in both productivity and welfare. The loss in welfare and productivity arises because the increase in adjustment frictions reduces the pace at which resources move away from less-productive to more-productive businesses. The same logic applies to changes in regulations or institutions that affect the costs of starting up or expanding a business, including regulations that raise the costs associated with expanding beyond some threshold of size.

A very different possibility is that firms are increasingly able to respond to shocks without as much churning of jobs and firms. After all, the churning of jobs and firms have no social value per se—it only has value to the extent that churning facilitates allocating outputs and inputs to their highest-valued use. Technological changes combined with globalization may have changed how businesses are organized and respond to shocks. For example, information and communications technology has arguably provided greater advantage for large, multinational firms in all sectors since these technologies can facilitate the coordination of production and distribution networks in multiple locations.¹² This change could help to explain the shift away from young firms to large, mature firms. In addition, it might help to explain the especially large declines in the share of startups and the decline in the pace of reallocation in specific sectors and geographic regions where information and communication technology has been especially relevant. Even with this more benign view, there might be a tradeoff between economies of scale induced by information and communications technology and flexibility in terms of how quickly the economy can adjust to changing economic conditions.

Yet another possibility is a hybrid of the first two explanations. Perhaps technological changes have induced changes in costs of hiring and training workers. For example, Cairó (2013) offers evidence that the training requirements of jobs have increased due to both the changing occupational mix of jobs and increases in training requirements within jobs. She develops a model where increased training requirements reduce the pace of job reallocation in the economy. She models this change in training requirements as an increase in adjustment costs, so that it is related to the first explanation above. However, it is an open question as to whether the increased training requirements of jobs is due to changes in technology, as Cairó hypothesizes, or to changing regulations or institutions, such as increased requirements for occupational licensing examined in Kleiner and Krueger (2013).

This discussion illustrates a few of many possible alternative causes, each with different consequences for trends in business dynamism. If the more sluggish pace of adjustment is due to increasingly burdensome regulation and institutions, this has potentially large adverse consequences for intermediate and long-run US job and productivity growth. Alternative explanations, like the notion of information and communications technology favoring large, multi-national firms suggest more benign consequences. Evaluating the productivity and welfare implications

¹² Evidence in support of the hypothesis that information and communications technology has favored large, multi-establishment firms for retail trade can be found in Doms, Jarmin, and Klimek (2004) and Foster, Haltiwanger, and Krizan (2006).

of the change toward lower levels of business dynamics and labor market reallocation will depend critically on the underlying causes of these changes.

Evidence about the role of startups from the rest of the world offers some useful perspective, too. Both the academic literature and international organizations such as the World Bank discuss the importance of entrepreneur-friendly business climates for economic growth; the 2013 *World Development Report* from the World Bank, focused on the theme of jobs, offers an excellent summary of the recent work in this area. But these issues are complicated and sometimes counterintuitive. For example, the evidence suggests that poorly performing emerging economies have plenty of entry—perhaps too much. Most startups in emerging economies are informal micro-enterprises with few if any paid employees. A high startup rate has little value per se. Instead, the problem in many emerging economies is that there is little or no evidence of post-entry growth (Hsieh and Klenow 2012). Using Schoar's (2010) characterization, poorly performing economies seem to have too many subsistence entrepreneurs and too few high-growth transformational entrepreneurs. Translating this issue back to the US context, the open question is whether the observed decline in startups and the associated decline in indicators of business dynamism imply fewer transformational entrepreneurs in the US economy.

Concluding Remarks

Startups and young firms are important contributors to job creation and productivity growth. The patterns for the US economy are roughly consistent with canonical heterogeneous firm models that place an important role for the reallocation of resources away from low-productivity (or low-profitability) businesses to high-productivity businesses. Entry and exit are important components of that ongoing reallocation. However, the contribution of startups and young businesses to jobs and productivity is a noisy and complex process. While startups contribute substantially to jobs immediately, most startups fail or, even if they survive, do not grow, while a small fraction of high-growth young firms contribute disproportionately to job creation in the United States. These findings pose challenges to policymakers seeking to promote job creation by encouraging entrepreneurship, because most young and small businesses are not in fact primary creators of jobs.

The rate of business startups in the US economy has been declining in recent decades, and business dynamism, as measured by the pace of job creation and job destruction, has declined as well. We do not have an explanation for the decline in the pace of entrepreneurship. Some of the structural changes we have observed in the economy such as the shift in the industrial composition actually increase the puzzle rather than provide an explanation since the shift has been towards sectors with higher paces of reallocation historically. The decline in entrepreneurial activity is present in all broad industrial sectors and in all US states. This finding suggests there are some common factors that are not sector- or region-specific accounting for the decline in entrepreneurial activity.

The decline in startups and accompanying decline in the pace of reallocation is a legitimate concern, although the causes of the decline and whether it is having or will have adverse consequences remain open questions for research. In considering these issues, it is worth noting that the declining pace of startups, job creation, and job destruction is mirrored in other measures of the dynamism of American society. In this journal in 2011, Molloy, Smith, and Wozniak noted that internal migration has been declining since a peak in the early 1980s. Like us, they failed to find simple explanations for these trends in changing demographics or broad economic factors. In a related fashion, the churn of workers over and above the job reallocation that we have emphasized here has exhibited a pronounced decline. The building of a career path and finding a good match has traditionally involved a high pace of job switching for workers (especially young workers). The evidence on worker churning implies that the process of building careers through job switching has slowed down. Taken together, there appears to be less scope for the US economy to adjust to changing economic conditions through the migration of workers, the reallocation of jobs across producers, and through the switching of workers across a given allocation of jobs.

■ *We thank the Kauffman Foundation for financial support, Kristin McCue and Lucia Foster for helpful comments, and Sinan Hastorun for excellent research assistance. We also thank the editors of this journal for many useful comments. Haltiwanger notes that the discussion in this paper draws on joint work with Steven Davis, Marcela Eslava, Jason Faberman, Teresa Fort, Lucia Foster, Ian Hathaway, C. J. Krizan, Adriana Kugler, Maurice Kugler, Scott Schuh, and Chad Syverson. Any opinions and conclusions expressed herein are those of the authors 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.*

References

- Acemoglu, Daron, Ufuk Akcigit, Nicholas Bloom, and William R. Kerr.** 2013. "Innovation, Reallocation, and Growth." NBER Working Paper 18993.
- Bartelsman, Eric J., and Mark Doms.** 2000. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature* 38(3): 569–94.
- Cairó, Isabel.** 2013. "The Slowdown in Business Employment Dynamics: The Role of Changing Skill Demands." At <http://www.econ.upf.edu/eng/graduates/gpem/jm/cairo.html>.
- Choi, Eleonor J., and James R. Spletzer.** 2012. "The Declining Average Size of Establishments: Evidence and Explanations." *Monthly Labor Review*, March, pp. 50–65.
- Comin, Diego A., and Thomas Philippon.** 2005. "The Rise in Firm-Level Volatility: Causes and