

WHO CREATES JOBS? SMALL VERSUS LARGE VERSUS YOUNG

John Haltiwanger, Ron S. Jarmin, and Javier Miranda*

Abstract—The view that small businesses create the most jobs remains appealing to policymakers and small business advocates. Using data from the Census Bureau's Business Dynamics Statistics and Longitudinal Business Database, we explore the many issues at the core of this ongoing debate. We find that the relationship between firm size and employment growth is sensitive to these issues. However, our main finding is that once we control for firm age, there is no systematic relationship between firm size and growth. Our findings highlight the important role of business start-ups and young businesses in U.S. job creation.

I. Introduction

A common popular perception about the U.S. economy is that small businesses create the most private sector jobs. This perception is popular among politicians of different political persuasions, small business advocates, and the business press.¹ While early empirical studies (see Birch, 1979, 1981, 1987) provided support for this perception, a variety of subsequent empirical studies have highlighted (see, in particular, Davis, Haltiwanger, Schuh, 1996) statistical and measurement pitfalls underlying much of the evidence in support of this perception. These include the lack of suitable data to study this issue, the failure to distinguish between net and gross job creation, and statistical problems associated with size classification methods and regression to the mean.² From a theoretical perspective, the notion of an inverse relationship between firm size and growth runs

counter to that described by Gibrat's law (see Sutton, 1997). But in spite of these questions from the academic literature, given the lack of definitive evidence to the contrary, the popular perception persists.

Neumark, Wall, and Zhang (2011; hereafter NWZ) recently performed a careful analysis where they avoid the misleading interpretations of the data highlighted by Davis, Haltiwanger, and Schuh (1996; hereafter DHS). Using the National Establishment Time Series (NETS) data including coverage across the U.S. private sector from 1992 to 2004, they find an inverse relationship between net growth rates and firm size. Their analysis indicates that small firms contribute disproportionately to net job growth.

In this paper, we demonstrate that an additional critical issue clouds the interpretation of previous analyses of the relationship between firm size and growth. Data sets traditionally employed to examine this relationship contain limited or no information about firm age. Our analysis emphasizes the role of firm age and, especially, firm births in this debate using comprehensive data tracking all firms and establishments in the U.S. nonfarm business sector for the period 1976 to 2005 from the Census Bureau's Longitudinal Business Database (LBD).³ As will become clear, the LBD is uniquely well suited to study these issues on an economy-wide basis.

Our main findings are summarized as follows. First, consistent with NWZ, when we do not control for firm age, we find an inverse relationship between net growth rates and firm size, although this relationship is quite sensitive to regression-to-the-mean effects. Second, once we add controls for firm age, we find no systematic inverse relationship between net growth rates and firm size. A key role for firm age is associated with firm births. We find that firm births

³ An important early study that also emphasized the role of firm age for growth dynamics is Evans (1987), who found an inverse relationship between firm growth and firm size (holding firm age constant) and between firm growth and firm age (holding firm size constant) using firm-level data for U.S. manufacturing firms. As Evans points out, the work is based on data with substantial limitations for tracking start-ups and young firms, but, interestingly, some aspects of his findings hold for our data which do not suffer from the same limitations. Specifically, the departures from Gibrat's law are primarily for young and small firms. A variety of other studies have also examined the role of employer age for employer dynamics and employment growth, including Dunne et al. (1989), Haltiwanger and Krizan (1999), and Acs, Armington, and Robb (1999). These studies focused on the establishment-age establishment-growth relation, including patterns of growth and failure, as well as the volatility of new establishments. All of these studies with the exclusion of Acs et al. (1999) are limited to the manufacturing sector.

Received for publication August 25, 2010. Revision accepted for publication December 21, 2011.

* Haltiwanger: University of Maryland and NBER; Jarmin and Miranda: U.S. Census Bureau.

We thank Philippe Aghion, Peter Huber, Harald Oberhofer, Michael Pfaffermayr, an anonymous referee, and conference and seminar participants at the NBER 2009 Summer Institute Meeting of the Entrepreneurship Working Group, CAED 2009, World Bank 2009 Conference on Small Firms, NABE Economic Policy Conference 2010, OECD Conference on Entrepreneurship 2010, Queens University and the 2010 WEA meetings for helpful comments. We thank the Kauffman Foundation for financial support. Any opinions and conclusions expressed here are our own and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00288.

¹ Policymakers regularly state that small businesses create most net new jobs. One of these common claims is that small businesses create two-thirds or more of net new jobs. Every president since President Reagan has included such statements in major addresses (often in the State of the Union addresses to Congress), and many other leaders in the U.S. House and Senate have made similar remarks. A list of selected quotes from speeches is available on request.

² Brown, Hamilton, and Medoff (1990) raise many related statistical issues in considering statistics by firm size but focus more on the impact of measurement issues for the employer size wage differential.

contribute substantially to both gross and net job creation. Importantly, because new firms tend to be small, the finding of a systematic inverse relationship between firm size and net growth rates in prior analyses is entirely attributable to most new firms being classified in small size classes.

Our findings emphasize the critical role start-ups play in U.S. employment growth dynamics. We document a rich up-or-out dynamic of young firms in the United States. That is, conditional on survival, young firms grow more rapidly than their more mature counterparts. However, young firms have a much higher likelihood of exit, so job destruction from exit is also disproportionately high among them. More generally, young firms are more volatile and exhibit higher rates of gross job creation and destruction than their older counterparts.

These findings highlight the importance of theoretical models and empirical analyses that focus on the start-up process—both the entry process and the subsequent post-entry dynamics, especially in the first ten years or so of a firm's existence. This is not to deny the importance of understanding and quantifying the ongoing dynamics of more mature firms but to highlight that business start-ups and young firms are inherently different.

Using the rich data available from the LBD and its public use version, the Business Dynamics Statistics (BDS), we highlight how the complex dynamics underlying firm formation, growth, decline, and exit combine to determine net job creation in the economy. The formation and execution of effective policies intended to increase net job creation require a rich and nuanced understanding of these processes. A natural conclusion from our findings on the role of firm size and age is that policies that target businesses of a certain size, while ignoring the role of age, will likely have limited success in improving net job creation. Our findings show that small, mature businesses have negative net job creation, and economic theory suggests this is not where job growth is likely to come from. Alternatively, our findings show that start-ups and young firms are important sources of job creation but that young firms are inherently volatile, with a high exit rate. It may be that even if the latter patterns are qualitatively consistent with healthy business dynamics, the challenges that start-ups and young firms face (such as regulatory challenges and market failures) warrant policy intervention. Exploring the latter is beyond the scope of this paper, but our findings highlight that effective policymaking in this area requires a rich understanding of such business dynamics. We return to this theme in our concluding remarks.

The rest of the paper proceeds as follows. In section II, we provide further background on the literature. Section III describes the data. Section IV presents the main empirical results. Section V provides concluding remarks. In several places, we point interested readers to an online appendix containing several analyses not discussed in detail here.⁴

⁴ Available at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00288.

II. Background

Much of the support for the hypothesis of an inverse relationship between employer size and growth comes from interpreting patterns observed in public use data products. An example is the Census Bureau's Statistics of U.S. Business (SUSB) that is released in partnership with the Small Business Administration.⁵ However, as demonstrated by NWZ and confirmed in this paper, this finding can also be obtained from a careful analysis of business microdata. In this section, we review the data and measurement issues in prior studies of firm size and growth and describe the characteristics of data sets suited to such analyses. We then briefly highlight findings from the Census Bureau's Business Dynamics Statistics (BDS). This new public use product gives data users a much richer window on the interactions of size, age, and growth that was previously available only to those with access to restricted-use data.

A. Review of Data and Measurement Issues

Analyses of the relationship between firm size and growth have been hampered by data limitations and measurement issues. As a consequence, these studies fail to emphasize a much richer description of the firm dynamics associated with the creative destruction process prevalent in market economies. Results from the new public use BDS, as well as from its underlying source data, the LBD, reveal a more accurate picture of firm dynamics with a more limited role for firm size. This section describes the basic characteristics of these data and how we address some of the limitations of prior analyses.

The analytical power of the LBD and data products constructed from it for understanding firm dynamics comes from its ability to accurately track both establishments and their parent firms over time.⁶ This is a critical feature of the data since it is difficult to discern the relationships of interest using only firm- or establishment-level data. Measures of job growth derived solely from establishment-level data have the virtue that they are well defined; when we observe an establishment grow, we know there are net new jobs at that establishment. In contrast, job growth observed in firm-level data may simply reflect changes in firm structure brought about by mergers, acquisitions, and divestitures. These activities clearly have an impact on observed employment at firms engaging in them and are ubiquitous features of market economies. For the purposes of allocating employment growth across different classes of firms (for example, by size, age, and industry), we clearly want to abstract from changes that reflect only a reallocation of employment across firms due to mergers and acquisitions.

⁵ SUSB data are available at <http://www.census.gov/econ/susb/index.html>.

⁶ For purposes of this discussion as well as the subsequent empirical analysis, we use the definitions of establishments and firms as defined by the U.S. Census Bureau. Specifically, an establishment is a specific physical location where business activity occurs, while a firm reflects all the establishments under common operational control.

Having only establishment-level data is inadequate as well. If the only data available are at the establishment level, the relationship between growth and the size and age of the establishment may not provide much information about the relevant firm size and firm age. A large, national retail chain is a useful example. In retail trade, a firm's primary margin of expansion is opening new stores rather than the expanding existing stores (see Doms, Jarmin, & Klimek, 2004; Foster, Haltiwanger, & Krizan, 2006; Jarmin, Klimek, & Miranda, 2009). This implies that there are many new establishments of existing firms, and for the core issues in this paper, the growth from such new establishments should be classified based on the size and age of the parent firm, not the size and age of the establishment. Much of the literature on employer size and net growth has been based primarily on establishment-level or firm-level data but not both.⁷ Tracking the dynamics of both firms and their constituent establishments permits clear and consistent measures of firm growth, as well as firm entry and exit.⁸

Even with rich source data, a key challenge in analyzing establishment and firm dynamics is the construction and maintenance of high-quality longitudinal linkages that allow accurate measurement of establishment and firm births and deaths. Given the ubiquitous changes in ownership among U.S. firms, a common feature observed in business microdata is spurious firm entry and exit caused by purely legal and administrative actions. Early versions of the D&B data Birch used were plagued with these limitations, which hampered the ability of researchers to distinguish between real business dynamics and events triggered by legal actions or business transactions such as credit applications (see Birley, 1984, and Aldrich et al., 1988, for detailed discussion). The NETS data used by NZW is based on a much improved version of the D&B data, although there are some open questions about the nature of the coverage in NETS.⁹ For our ana-

lysis, we minimize the impact of these data quality issues by using the LBD's high-quality longitudinal establishment linkages and its within-year linkages of establishments to their parent firms.

DHS recognized the statistical pitfalls in relating employer size and growth. One issue they highlight is the role of regression-to-the-mean effects. Businesses that recently experienced negative transitory shocks (or even transitory measurement error) are more likely to grow, while businesses recently experiencing positive transitory shocks are more likely to shrink. This effect alone will yield an inverse relationship between size and growth. Friedman (1992) states that this type of regression fallacy "is the most common fallacy in the statistical analysis of economic data." This issue is particularly relevant when studying the business size-growth relationship and is manifest in the method used to classify businesses into size classes in many commonly used data sources. The early work by Birch and others classified businesses into size classes using base year employment, a method now known to yield results that suffer from regression to the mean.

DHS propose an alternative classification method to mitigate the effects of regression to the mean. They note that while base-year-size classification yields a negative bias, using end-year-size classification yields a positive bias. To avoid the bias, negative or positive, DHS propose using a classification based on current average size, where current average size is based on the average of employment in years $t - 1$ and t . Using current average size is a compromise between using year $t - 1$ (base) or year t (end) size to classify firms. In what follows, we refer to current average size as simply average size.

Although average size is a compromise, it has limitations as well. Firms that are affected by permanent shocks that move the firm across multiple size class boundaries between $t - 1$ and t will be classified into a size class that is between the starting and ending size classes. Recognizing this potential limitation, the Bureau of Labor Statistics has developed a dynamic size classification methodology (see Butani et al., 2006).¹⁰ Specifically, the methodology attributes job gains or losses to each of the size classes that the firm passes through in its growth or contraction. Interestingly, comparisons across size-classification methods show that the average (DHS) and dynamic (BLS) size classification methodologies yield very similar patterns. This is not surprising since both are a form of averaging over time to deal with transitory shocks.

We prefer the average size class methodology because it is inherently more robust to regression-to-the-mean effects. However, we also report results using the base-year metho-

⁷ DHS analysis is restricted to U.S. manufacturing establishments, although they were able to construct a measure of firm size at the manufacturing level. Dunne et al. (1989) examine the role of establishment size and age for the growth and failure of U.S. manufacturing plants. Evans (1987) used firm-level data for a sample of firms in the U.S. manufacturing sector in continuous operation between 1976 and 1980. Birch (1979, 1981, 1997) uses the D&B data, which have both firm and establishment-level information although subject to the limitations of the D&B data. NZW use the NETS data with both firm- and establishment-level information.

⁸ In our analysis, firm entry is defined when all of the establishments at that firm are de novo establishment entrants. Firm exit is defined when all of the establishments at that firm cease operations.

⁹ NWZ report about 13.1 million firms and 14.7 million establishments in a typical year. The LBD (and the closely related County Business Patterns) report about 6 million firms and 7 million establishments in a typical year that have at least one paid employee. The Census Bureau also reports more than 15 million additional nonemployer businesses in a typical year. It appears that NETS is some combination of employer and nonemployer businesses but does not reflect the universe of businesses. For our purposes, we focus on employer businesses. For discussion of the importance of nonemployer businesses and the relationship between nonemployer and employer businesses, see Davis et al. (2009). There also remain questions about how well NETS captures start-ups especially for small businesses. These questions about coverage also raise questions about whether the type of analysis we conduct here focusing on the role of firm age would be feasible with NETS. We provide a table comparing the major characteristics of the principal data sets available to study the dynamics of U.S. businesses in our online appendix.

¹⁰ Related evaluation work on alternative methodologies by BLS is found in Okolie (2004). We also note that the BLS BED series releases net and gross job quarterly flows by this firm size measure. The firm size measure they use is based on a taxpayer ID definition of the firm, so for multiunit establishment firms that have multiple taxpayer IDs, their firm definition is somewhere between the establishment and overall firm.

dology for our core results and also in order to explore the sensitivity of the results to this methodological issue.¹¹

DHS also emphasize avoiding inferences that arise from the distinction between net and gross job creation. Policy analysts are tempted to want to make statements along the lines that “small businesses account for X% of net job creation.” The problem with statements like this is that many different groupings of establishments can account for a large share of the net job creation since gross job flows dwarf net job flows. For example, the annual net employment growth rate for U.S. nonfarm private sector business establishments between 1975 and 2005 averaged 2.2%. Underlying this net employment growth rate were establishment-level average annual rates of gross job creation and destruction of 17.6% and 15.4%, respectively (statistics from the BDS, which we describe below). Decomposing net growth across groups of establishment or firms is problematic (at least in terms of interpretation) when some shares are negative. We elaborate on these issues in section II.B by taking a closer look at the Census Bureau’s new BDS data.

B. Overcoming Data and Measurement Issues with the BDS

To help illustrate these points before proceeding to the more formal analysis, we examine some tabular output from the BDS on net job creation by firm size and firm age. The precise definitions of firm size and firm age are discussed below (and are described on the BDS website, <http://www.census.gov/ces/dataproducts/bds/index.html>). Table 1 shows the number of net new jobs by firm size and firm age class in 2005. The upper panel shows the tabulations using the base-year-size method and the lower panel the average-size method. The table yields a number of interesting observations. About 2.5 million net new jobs were created in the U.S. private sector in 2005. Strikingly, firm start-ups (firms with age 0) created about 3.5 million net new jobs. In contrast, every other firm age class except for the oldest firms exhibited net declines in employment in 2005.¹² However, it would be misleading to say that it is only firm start-ups and the most mature firms that contributed to job gains. In both panels there are large positive numbers in many cells but also large negative numbers in other cells. It is also clear that there are substantial differences in these patterns depending on using the base-year or average-size method, although some common patterns emerge. For example, excluding start-ups, firms that employ between 5 and 99 workers consistently exhibit declines in net jobs.

The patterns reflect two basic ingredients. Obviously, whether the size or age class contributes positively or negatively depends on whether that class has a positive or negative net growth rate. In addition, the magnitude of the positive or negative contribution depends, not surprisingly, on how

much employment is accounted for by that cell. That is, a size or age class may have a large positive number not so much because it has an especially high growth rate but because it accounts for a large fraction of employment (for example, a 1% growth rate on a large base yields many net new jobs).

Figure 1 summarizes these patterns in the BDS over 1992 to 2005 by broad size and age classes.¹³ The figure shows the fraction of job creation and job destruction accounted for by small (fewer than 500 workers) and large firms (500 workers and above) broken out by whether they are firm births, young firms (less than 10 years old), or mature firms (10 years and above). Also included is the share of employment accounted for by each of these groups. We focus on gross job creation and destruction at the establishment level but classified by the characteristics of the firms that own them.

Several observations emerge. First, for the most part, the fraction of job creation and destruction accounted for by the various groups is roughly proportional to the share of employment accounted for by each group. For example, it is the mature and large firms that account for most employment (about 45%) and most job creation and destruction. This observation, while not surprising, is important in the debate about what classes of businesses create jobs. The basic insight is that the firms that have the most jobs create the most jobs, so workers looking for the places where the most jobs are being created should go where the jobs are: large and mature firms. This is not the whole story, of course, as what we are primarily interested in is whether any identifiable groups of firms disproportionately create or destroy jobs. The rest of the paper is a rigorous examination of this issue. However, figure 1 nicely previews some of our primary findings. Young firms disproportionately contribute to both job creation and job destruction. Included among young firms are firm births, which by definition contribute only to job creation. Nearly all firm births are small.¹⁴ Before the BDS, all publicly available data that could be used to look at the role of firm size in job creation were silent on the age dimension. As such, it is easy to see how analysts perceived an inverse relationship between size and growth in the data.

Before proceeding, it is instructive to discuss briefly the implications of focusing on March-to-March annual changes of employment at the firm and establishment level in our analysis of firm dynamics and job creation. One implication is that we neglect high frequency within year firm and establishment dynamics—for example, changes that are transitory and reverse themselves within the year. We think that for the most part, neglecting such high-frequency variation is not important for the issues of con-

¹³ We use the base-year size method in figure 1. The results in figure 1 are robust to using either of the size classification methods discussed in the analysis that follows. Precise definitions of job creation and destruction are provided below.

¹⁴ These large firm births are often associated with the appearance of a new U.S. affiliate of a foreign-owned firm or changes in employment arrangements, for example, the use of employee leasing firms.

¹¹ The online appendix includes all results by base-year size methodology.

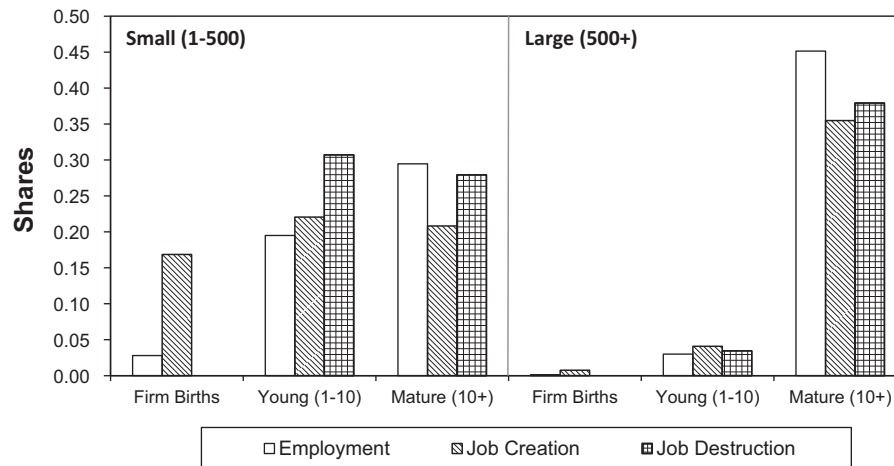
¹² Note that LBD processing uses longitudinal edits that can alter values in the BDS tables.

TABLE 1.—NET JOB CREATION BY FIRM SIZE AND FIRM AGE, U.S. PRIVATE SECTOR, 2005

Firm age	Firm Size (Base Year)																All
	1-4	5-9	10-19	20-49	50-99	100-249	250-499	500-999	1,000-2,499	2,500-4,999	5,000-9,999	10,000+					
0	731,515	503,644	498,317	553,181	313,511	292,348	157,120	151,518	186,087	131,178	D	D	3,518,419				
1	79,759	-12,547	-20,836	-47,837	-41,006	-57,188	-48,830	-5,476	-14,532	-20,131	-408	-408	-188,821				
2	26,506	-24,840	-31,883	-44,488	-26,738	-18,026	-9,049	-13,579	-23,615	-12,782	D	D	-178,494				
3	7,535	-22,650	-26,855	-37,824	-15,918	-14,813	-8,981	-7,548	-11,581	-12,114	D	D	-150,749				
4	20,456	-18,442	-23,212	-29,616	641	-9,816	-4,301	-5,436	-298	-4,011	D	D	-74,035				
5	4,808	-19,792	-24,392	-29,425	-14,870	-6,222	-2,449	-6,849	-293	-3,418	D	D	-102,902				
6-10	14,577	-71,332	-99,235	-110,111	-40,652	-1,324	-9,452	5,437	-20,693	-13,945	-9,903	17,928	-338,705				
11-15	15,663	-47,730	-67,923	-81,876	-40,432	-27,666	-9,530	2,179	-2,028	22,441	6,140	69,409	-161,353				
16-20	5,673	-36,856	-58,236	-71,299	-35,979	9,780	-5,725	10,200	3,204	12,615	10,491	2,158	-153,974				
21-25	2,923	-28,173	-42,609	-51,490	-22,246	-13,346	3,901	10,269	36,484	10,075	9,889	-56,563	-140,886				
26+	1,016	-38,599	-71,235	-107,390	-48,873	10,309	19,924	85,473	56,436	143,701	58,245	307,517	416,524				
All	910,431	182,683	311,901	-58,175	27,438	164,036	82,628	226,188	209,171	253,609	90,973	360,214	2,481,097				
Firm age	Firm Size (Current)																
	0	1	2	3	4	5	6-10	11-15	16-20	21-25	26+	All					
0	1,157,210	541,230	453,073	445,091	236,121	216,911	151,518	128,772	188,493	D	D	D	3,518,419				
1	-188,206	-1,242	10,705	3,028	-20,046	-28,733	20,118	14,346	-6,509	7,898	-42	-138	-188,821				
2	-102,079	-34,487	-24,132	-15,745	-5,380	3,125	5,036	-9,743	-13,282	8,392	D	D	-188,295				
3	-77,770	-30,935	-25,119	-12,259	1,824	2,215	2,572	888	-10,155	D	3,699	D	-145,040				
4	-61,216	-27,141	-19,487	-7,210	1,630	2,221	3,505	6,655	7,375	-10,228	D	D	-103,896				
5	-54,616	-28,196	-23,791	-16,205	-2,595	6,890	5,779	11,703	-4,850	3,017	D	D	-102,864				
6-10	-190,115	-112,735	-99,872	-76,025	-17,730	13,713	26,305	19,344	5,364	26,494	23,546	43,006	-338,705				
11-15	-105,596	-74,905	-75,477	-60,259	-17,677	11,166	20,401	1,617	34,591	18,886	20,201	65,699	-161,353				
16-20	-74,278	-59,389	-61,306	-60,496	-13,235	12,172	27,334	6,559	4,413	16,969	14,550	32,733	-153,974				
21-25	-49,929	-43,548	-47,143	-42,924	-16,172	4,020	23,438	22,298	30,120	34,280	-46,129	-9,197	-140,886				
26+	-89,878	-83,682	-107,356	-114,182	-40,005	42,481	63,939	69,597	93,401	110,311	38,147	433,751	416,524				
All	163,527	44,970	-19,905	42,814	106,735	286,181	349,945	272,036	328,961	253,373	71,269	581,191	2,481,097				

U.S. Census Bureau, Business Dynamics Statistics at http://www.ces.census.gov/index.php/bds/bds_home.

FIGURE 1.—SHARES OF EMPLOYMENT, JOB CREATION, AND DESTRUCTION BY BROAD FIRM (AVERAGE), SIZE AND AGE CLASSES: ANNUAL AVERAGE RATES, 1992–2005



cern in this paper but would be of more relevance in exploring cyclical volatility by firm size and age.

However, a related implication of focusing on March-to-March annual changes is that very short-lived firms that enter and exit between March of one year and March of the subsequent year are not captured in our analysis. The neglect of the latter might be important in the current context given our findings of the important role of firm births for job creation, as is evident in table 1 and figure 1. Fortunately, the LBD includes information that suggests that such short-lived firm births are not especially important. That is, the LBD also includes annual payroll for all establishments and firms. The payroll measure captures any positive activity of establishments and firms, including very short-lived firms, whereas employment is measured only as of March 12. Using the same longitudinal links as used in the BDS and LBD, we calculated the payroll-weighted firm entry rate as 1.72% of payroll. This compares to the employment-weighted firm entry rate of 2.79% of employment in figure 1. It is not surprising that the payroll-weighted entry rate is lower than the employment-weighted entry rate given that entrants are small and pay lower wages. Of more interest is how much of the payroll-weighted entry rate is accounted for by very short-lived entrants. Excluding the short-lived entrants (defined as firm start-ups that do not survive until March), the payroll-weighted entry rate is 1.64%. This negligible decline in the payroll-weighted entry rate from short-lived entrants implies that such entrants account for very little of the activity even for start-ups. Abstracting from such short-lived firms should not have a quantitatively important impact on our analysis. It does, however, remind us of the highly volatile nature of start-ups, an issue that we discuss further.

C. Data and Measurement

The LBD underlies the public use statistics in the BDS. As suggested in section II.B, many of the patterns we discuss in this paper can be readily seen in the public domain

BDS. However, we use the LBD microdata rather than the BDS since in using the LBD microdata, we can identify firms and abstract from firm growth due to ownership changes in the manner we describe below.¹⁵

The LBD (Jarmin & Miranda, 2002) covers all business establishments in the U.S. private nonfarm sector with at least one paid employee.¹⁶ The LBD begins in 1976 and currently covers over thirty years of data, including information on detailed industry and employment for every establishment. We note that the LBD (and, in turn, the BDS) employment and job creation numbers track closely those of the County Business Patterns (CBP) and Statistics of U.S. Business (SUSB) programs of the U.S. Census Bureau (see Haltiwanger, Jarmin, & Miranda, 2009) as they all share the Census Bureau's Business Register (BR) as their source data. However, due to design features and differences in processing, in particular the correction of longitudinal establishment and firm linkages, the statistics generated from the LBD diverge slightly from those in CBP and SUSB.

The unit of observation in the LBD is the establishment defined as a single physical location where business is conducted. Each establishment-year record in the LBD has a firm identifier associated with it, so it is possible to track the ownership structure of firms in any given year, as well as changes over time. Firms can own a single establishment or many establishments. In some cases, these firms span multiple geographic areas and industries. Establishments can be acquired, divested, or spun off into new firms, so the ownership structure of firms can be dynamic and complex. We use these firm-level identifiers to construct firm-level characteristics for each establishment in the LBD.

¹⁵ Future versions of the BDS will include firm-level net employment growth rates and components using the type of methodology we have developed for this paper.

¹⁶ This is one clear distinction with the NETS database, which apparently includes both employer and nonemployer businesses (but also apparently not the universe of both).

A. Measuring Firm Age and Firm Size

The construction of firm size measures is relatively straightforward. Firm size is constructed by aggregating employment across all establishments that belong to the firm. As discussed above, we measure firm size using both the base-year and average-size methodologies. For base-year firm size, we use the firm size for year $t - 1$ for all businesses except new firms. For new firms, we follow the approach used by Birch and others and allocate establishments belonging to firm start-ups to the firm-size class in year t . For average size, we use the average of firm size in year $t - 1$ and year t . We use the same approach for new, existing, and exiting firms when using average size.

The construction of firm age presents more difficult conceptual and measurement challenges. We follow the approach adopted for the BDS and based on our prior work (see Becker et al., 2006, and Davis et al., 2007). The firm identifiers in the LBD are not explicitly longitudinal. Nevertheless, they are useful for tracking firms and their changing structure over time. A new firm identifier can appear in the LBD either due to a de novo firm birth or changes in existing firms. For example, a single-location firm opening additional locations is the most common reason for a continuing firm in the LBD to experience a change in firm ID. Other reasons include ownership changes through mergers and acquisitions. When a new firm identifier appears in the LBD for whatever reason, we assign it an age based on the age of the oldest establishment that the firm owns in the first year the new firm ID is observed. The firm is then allowed to age naturally (by one year for each additional year the firm ID is observed in the data) regardless of mergers or acquisitions and as long as the firm's ownership and control do not change. An advantage of this approach is that firm births as well as firm deaths are readily and consistently defined. That is, a firm birth is defined as a new firm ID, where all the establishments at the firm are new (entering) establishments. Similarly, a firm death is defined as when a firm ID disappears and all of the establishments associated with that firm ID cease operations and exit. If a new firm identifier arises through a merger of two preexisting firms, we do not treat it as a firm birth. Rather, the new firm entity associated with the new identifier is given a firm age equal to the age of the oldest continuing establishment of the newly combined entity.

Thus, our firm size and age measures are robust to ownership changes. For a pure ownership change with no change in activity, there will be no spurious changes in firm size or firm age. When there are mergers, acquisitions, or divestitures, firm age will reflect the age of the appropriate components of the firm. Firm size will change, but in a manner consistent with the change in the scope of activity.

Before proceeding, we note that we focus on growth dynamics of establishments and firms over the 1992–2005 period. We limit our analysis to this period so that we can define firm age consistently over the period for all establishments with firm age less than fifteen years. We also include

a category for establishments belonging to firms that are sixteen years or older (in 1992, these are the firms with establishments in operation in 1976 and for which we cannot give a precise measure of firm or establishment age).

B. The Establishment-Level and Aggregate Growth Rate Concepts

This section describes the establishment- and firm-level growth rate measures we use in the paper in more detail. Let E_{it} be employment in year t for establishment i . In the LBD, establishment employment is a point-in-time measure reflecting the number of workers on the payroll for the payroll period that includes March 12. We measure the establishment-level employment growth rate as follows:

$$g_{it} = (E_{it} - E_{it-1})/X_{it},$$

where

$$X_{it} = .5 * (E_{it} + E_{it-1}).$$

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 DHS and Tornqvist, Vartia, & Vartia 1985.¹⁷)

Note that the DHS growth rate measure can be flexibly defined for different aggregations of establishments. We first discuss the measures of net growth used in the analysis. In particular, consider the following relationships,

$$g_t = \sum_s (X_{st}/X_t) g_{st} = \sum_s ((X_{st}/X_t) \sum_{i \in s} (X_{it}/X_{st}) g_{it}),$$

where

$$X_t = \sum_s X_{st} = \sum_s \sum_{i \in s} X_{it},$$

where g_t is the aggregate DHS growth rate and s indexes classifications of establishments into groups defined for any level of aggregation s where s can refer to firm, industry, firm size, or firm age classifications. Thus, the DHS net growth rates for various aggregations of interest are just properly weighted sums of establishment-level growth rates where the establishment is the lowest level of aggregation in the LBD. Important groupings for this paper include firms and firm size and age categories.

Before discussing components of the DHS net growth that we use in our analysis, it is important to discuss how computing DHS net growth rates at different levels of

¹⁷ The DHS growth rate like the log first difference is a symmetric growth rate measure but has the added advantage that it accommodates entry and exit. It is a second-order approximation of the log difference for growth rates around 0. Note that the use of a symmetric growth rate does not obviate the need to be concerned about regression-to-the-mean effects. Also, note that the DHS growth rate is not only symmetric but bounded between -2 (exit) and 2 (entrant).

aggregation can affect interpretation. We are interested in computing net growth rates at both the establishment and firm levels. In the LBD, we have access to both levels of data where the establishment structure of the firms is well specified. In other settings, however, the analyst may have access to only establishment or only firm-level data. Thus, it is critical to understand how using one or the other can affect interpretation.

An important difference in computation and interpretation arises when establishments undergo 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 or shedding continuing establishments. This occurs even if the added or shed establishments experience no employment changes themselves.

To avoid this problem, we compute firm growth rates as suggested in the expressions above. Namely, the period $t - 1$ to period t net growth rate for a firm is 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, that 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 mergers and acquisitions.¹⁸

We use the establishment- and firm-level growth rate measures to compute not only net growth but also job creation and job destruction (and the related job creation from entry and job destruction from exit). At the establishment level, job creation is measured as the employment gains from all new and expanding establishments and job destruction as the employment losses from all contracting and closing establishments. At the firm level, job creation is measured as the employment gains from all expanding and new firms and job destruction as the employment losses from all contracting and exiting firms. By construction, our methods of computing growth imply that firm-level measures of job creation and destruction are lower than establishment-level measures since the latter includes within-firm reallocation of jobs across establishments. For these measures, we follow the approach developed by DHS. Details of the measurement of these concepts are provided in the online appendix.¹⁹ A key

identity that we exploit is that the overall net employment growth rate at the cell and aggregate level can be decomposed into the appropriately weighted sum of net employment growth from continuing firms, job creation from entry, and subtracting the job destruction from exit. As we show in the online appendix, our overall net growth results can be equivalently obtained from direct estimation of the effects on overall net growth or by first estimating the effects on the components and computing the employment-weighted sum of the components.

IV. The Relationship of Employment Growth, Firm Size, and Firm Age

Our primary objective is to understand the relationship between net employment growth, and its components, and firm size and age. In this section, we use a nonparametric regression approach to quantify these relationships. In our main specification, we regress net employment growth and its components at the firm level on firm size classes by themselves, on firm age classes by themselves, and by firm size and age interacted together. The latter specification follows naturally from the tabulations in table 1, which show net growth patterns for firm size and firm age cells. All of the empirical models we consider are fully saturated dummy variable models. They are either one-way dummy variable models in firm size or firm age or two-way models with a complete set of interactions. As Angrist and Pischke (2009) highlight, estimating fully saturated dummy variable models with OLS is fully general regardless of the distribution of the dependent variable. This is intuitive since by construction, the estimated coefficients will be the cell means for each of the saturated cells. This property is important in our case since the DHS net growth rate is bounded between -2 and 2 . When we estimate specifications with the dependent variables being the job creation from entry and job destruction from exit, the specifications are equivalent to linear probability models (albeit weighted as noted below). But again, as Angrist and Pischke (2009) discuss, the fully saturated dummy variable model avoids any econometric issues with using a bounded or limited dependent variable.²⁰

²⁰ We discuss the desirable econometric properties of fully saturated dummy variable models further in the online appendix in section VII.D. We also note that earlier versions of this paper (Haltiwanger, Jarmin, & Miranda, 2010) considered two-way models with firm size and firm age dummies without interactions and were potentially subject to econometric concerns with predicted values that may lie outside the range of the dependent variable. We show in appendix VII.D that the partial effects of firm size controlling for firm age (and vice versa) are quite similar for the models with and without interactions. Thus, the results from earlier versions and the current version yield very similar results and inferences. While the specifications without interactions have potential limitations, such specifications are transparent and parsimonious in terms of generating partial effects. Moreover, the two-way model without interactions permits estimating more detailed firm size and firm age effects. We also note that in earlier versions, we considered specifications controlling for industry and year effects. As we show in the online appendix, all of our results are robust to controlling for industry and year effects. We thank Peter Huber, Harald Oberhofer, and Michael Pfaffermayr for helpful discussions on these issues.

¹⁸ In the online appendix, we provide a detailed hypothetical example to clarify how in practice we handle mergers and acquisitions. This example is useful to understand the details, as well as for practitioners who want to implement our methodology.

¹⁹ The online appendix also includes depictions of the distribution of firm- and establishment-level net growth rates underlying the job creation and destruction statistics.

We focus on employment-weighted specifications since this enables the coefficients to be interpreted in terms of the impact on net employment growth rates at the aggregate level for the specified category. Given the discussion, our estimates are equivalent to employment-weighted cell means. As such, we can replicate all of the results in what follows using a cell-based regression approach where net growth rates and components are measured at firm size, firm age, and year level of aggregation.

Our focus is on comparing the effects of firm size on net growth (and components of net growth) with and without controls for firm age (and vice versa). To quantify the effects of firm size using the two-way model with interactions, we compute partials of firm size from that model holding the age distribution of employment constant at the sample mean. The partial effect of firm age analogously holds the size distribution of employment constant at the sample mean. Our approach can be interpreted in two closely related ways. First, it is equivalent to the standard approach in regression models when one computes partials of a variable x when the model includes interaction terms such as $x \times z$ by evaluating the interaction coefficients at the mean of z . Second, our approach is equivalently interpretable as a form of a shift and share decomposition holding the age (size) composition constant when examining the firm size (age) effect. Specifically, our firm size (age) estimates are age (size) composition constant estimates based on computing the weighted average of the firm size by firm age cell means using the overall employment distribution by firm age (size) as weights. Finally, we note that our findings controlling for firm age (size) are essentially the same if we use simpler and more transparent two-way models without interactions, as we discuss in the online appendix. We focus on the results from fully saturated models in the text since they use a more general and robust econometric specification (see the online appendix for further discussion).

In estimating our fully saturated dummy variable specifications, we need to take into account that some parts of the joint firm size and firm age distribution get very sparse (for example, young firms with more than 10,000 employees are virtually nonexistent, as seen in table 1). So relative to the firm size and firm age classes in table 1, we restrict ourselves to eight size classes (1–4, 5–9, 10–19, 20–49, 50–99, 100–249, 250–499 and 500 and up) and nine age classes (0, 1–2, 3–4, 5–6, 7–8, 9–10, 11–12, 13–15, and 16 and up). The fully saturated two-way model generates a very large number of estimated coefficients, especially as we estimate specifications for both initial size and current (average) size, as well as overall net growth and the components of net growth. Given our focus on the partial effects of size controlling for firm age and vice versa, we report our results in the subsequent sections in figures.²¹ We note that the

underlying firm-level regressions include more than 70 million firm-year observations; consequently, the standard errors for the estimates are very small.²²

B. Net Employment Growth and Firm Size

We present all of our remaining results with the aid of figures. To facilitate comparisons between the one-way models and the partials from the two-way models with interactions, we focus on comparing the differences in effects relative to a baseline or omitted group. In all of the subsequent figures, this is the largest (500 and up) or oldest (16 and up) group. To facilitate the interpretation of the magnitudes, we report the baseline group at its unconditional mean from the one-way model. In turn, we simply rescale the other effects by adding the value of the unconditional mean for the baseline category (for example, the firm size class of 500 and up) from the one-way model to each effect. Adding the unconditional mean to all categories does not distort the relative differences but provides perspective about the magnitude of the effects.


Figure 2 shows the relationship between net employment growth and firm size. The upper panel displays results from the regressions for all firms. The lower panel displays the size effects from the regressions where we limited the sample to continuing firms only. Beginning with the main results in the upper panel, the plotted curve for the base-year size specification without age controls shows a strong inverse relationship between firm size and net employment growth. The average annual rate of net employment growth in the smallest size class is about 15.2 percentage points higher than that for the largest firms (500 or more employees). The effect declines more or less monotonically as the size of the firm increases. The relative net employment growth premium for being small declines to 3.3% and 1.7% for size classes 5–9 and 10–19 employees, respectively.

As argued above, however, the base-year measure of firm size has several undesirable attributes for examining firm size and growth. The curve plotting the estimated effects from our preferred average size specification with no age controls shows that the inverse relationship remains, but the quantitative relationship is substantially muted. Comparing the base and average size results suggests that the effects of regression to the mean are quite strong in the smallest size classes. In the online appendix, we show that consistent with these patterns, the negative serial correlation of firm-level net employment growth rates is especially large in absolute value for small firms. But the more general point is that in the absence of controls for firm age, we obtain similar qualitative results as those in NZW. That is, size classification methodology matters, but there still is a small inverse


²¹ In table W.2 we show the results for the net overall growth fully saturated one-way and two-way models. In table W.3 we show the results for the net growth fully saturated models (one-way and two-way) for continuing firms. Estimates for other specifications are available on request.

²² The largest standard error in any of the fully saturated models for net employment growth, net employment growth for continuers, job creation from entry, or job destruction from exit is 0.003. Most standard errors are below 0.001.

FIGURE 2.—RELATIONSHIP BETWEEN NET GROWTH AND FIRM SIZE



B. Continuing Firms only

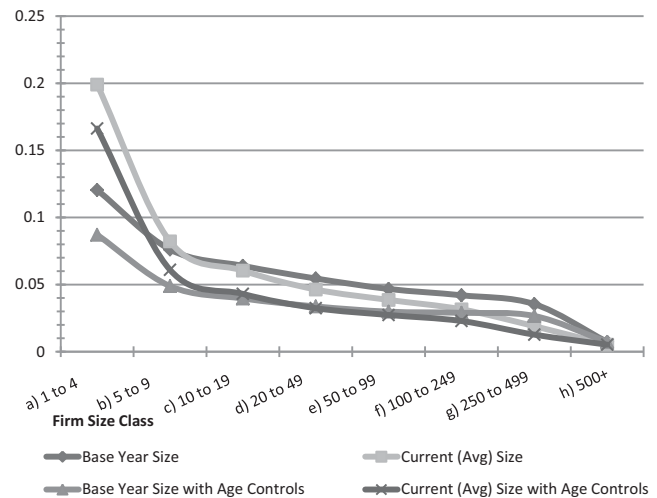


relationship between net employment growth and firm size when not controlling for firm age.

Controlling for firm age, however, has a dramatic impact on these patterns. Regardless of the size classification methodology, once we control for firm age, we observe no systematic inverse relationship between net growth and firm size. When we use base-year size, the smallest size class has the largest positive coefficient, but the size classes in the range from 5 to 499 have the most negative effects. This implies that firms in the 5 to 499 range have lower net growth rates on average than the largest businesses once we control for firm age. When we use average size, we find a positive relationship between net growth and firm size for all the size classes up through 500 workers. While the details differ nontrivially depending on which size class method we use, the main point is that once we control for firm age, there is no evidence that small firms systematically have higher net growth rates than larger businesses.

In the lower panel of figure 2, we show the results when we restrict the analysis to continuing firms only. The first thing to note is that there is a less dramatic impact of con-

FIGURE 3.—FIRM EXIT BY FIRM SIZE



trolling for firm age since there is, by construction, no role for start-ups.²³ Exploring this more deeply, we find a strong inverse relationship between net growth and firm size for continuing firms when we use the base-size methodology. This is the case whether or not we control for firm age. However, using average size, there is a positive relationship between net firm growth and firm size regardless of whether one controls for firm age. Hence, for continuing firms, it is primarily the size class methodology that matters. The stark differences for small continuing firms between the base-size and average-size results are consistent with the strong regression-to-the-mean effects for these firms.

Some of the differences between the patterns across the two panels of figure 2 reflect the role of firm exits.²⁴ We explore this further in figure 3, which shows the patterns of job destruction from firm exit by firm size with and without age controls. Job destruction from firm exit is directly interpretable as an employment-weighted firm exit rate. The firm exit rate falls monotonically with firm size regardless of size class methodology and with or without firm age controls. Controlling for firm age yields somewhat higher exit rates for small businesses, but this effect is quite modest when using average-size class methodology. Thus, a robust finding is that small firms are more likely to exit than larger firms, even controlling for age.

Combining figure 3 with the lower panel of figure 2 helps account for the patterns in the upper panel of figure 2, especially for the results controlling for firm age. The lower panel of figure 2 shows that when controlling for firm age, there is a modest but increasing relationship between net growth and average size for continuing firms. Combining

²³ NWZ briefly discuss a similar result they obtained using the NETS data when they exclude start-ups.

²⁴ Section VII.E of the online appendix shows that the overall net effects we report in table 2 and figure 2 can be generated by using the estimates for the components of net growth (continuers, job creation from entry and job destruction from exit). This property holds for all the overall net effects we report in the paper.

this effect with the patterns in figure 3 where small firms (controlling for firm age) have much higher exit rates yields that net growth rates are strongly increasing in average firm size controlling for firm age.

Figures 2 and 3 also shed light on Gibrat's law (the prediction that firm growth should be independent of size). Figure 2 suggests that Gibrat's law holds approximately if we exclude the smallest firms, especially if we use the average size measure and do not control for firm age. That is, departures arise for the smallest firms (where regression-to-the-mean effects are especially an issue) and for entering and young firms, which, as we will see below, have their own interesting dynamic not well captured by the models underlying the predictions of Gibrat's law (Sutton, 1997).

An appropriate measure of firm age is critical for obtaining the patterns in figures 2 and 3. As we noted above, these results can also be obtained by estimating employment-weighted establishment-level regressions on firm size and age characteristics. This implies that we can check the robustness of controlling for establishment as opposed to firm-level characteristics. For brevity, we summarize only the results looking at establishment characteristics and point interested readers to the online appendix. We find that controlling for establishment as opposed to firm age does not yield the same stark patterns of figures 2 and 3. That is, when controlling for only establishment age, the relationship between firm size and net growth remains strongly negative when using base size unlike the pattern in figure 2, which shows a nonmonotonic relationship between firm size and net growth when we control for firm age. Moreover, the positive relationship between average size and net growth in figure 2 when controlling for firm age becomes notably weaker when controlling for establishment age. These findings highlight the important distinction between firm and establishment age that comes about because there are many young, small establishments of large, mature firms.

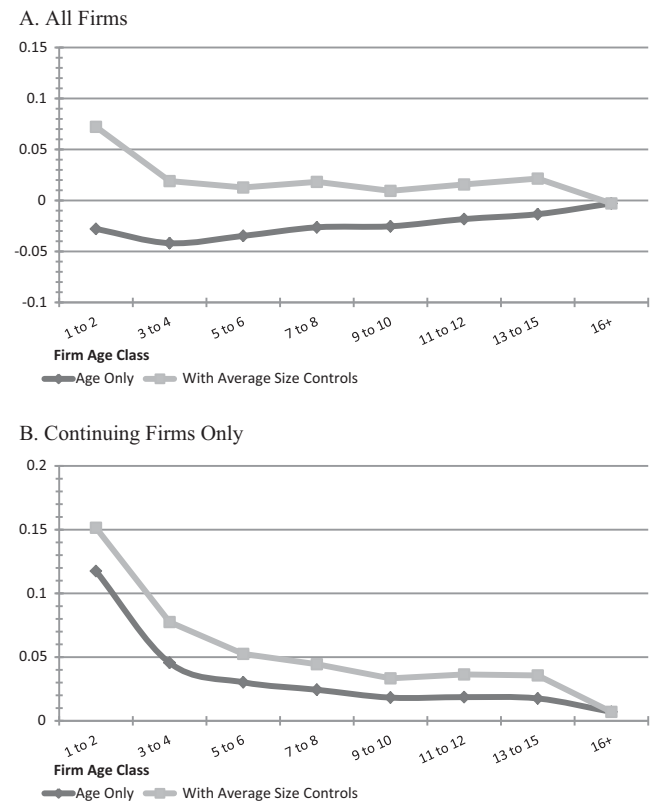
B. Net Employment Growth and Firm Age

We now turn to exploring the patterns of net employment growth by firm age. For ease of exposition, in what follows we show in figures only the results for firm age by itself and controlling for firm size using our preferred average-size measure. (For completeness, we provide the results controlling for base-year size in the online appendix.)

The top panel of figure 4 shows the results for firm age. In the figure, we omit the estimated coefficient for start-ups since it is much higher (for the fully saturated models, the predicted estimate is identically equal to 2).²⁵ The panel

²⁵ Recall that at the firm level, the net growth rate for a firm start-up is equal to 2 using the DHS methodology. The fully saturated models yield the predicted value as the cell means, so the predicted value is identically equal to 2. Note that in the appendix, we consider models where we control for industry and year effects that break this identity. Even then, we find that the predicted values for firm age = 0 are close to 2 and very similar across size classes. We also find that the results controlling for industry and year effects are very similar to those we report in the text.

FIGURE 4.—RELATIONSHIP BETWEEN NET EMPLOYMENT GROWTH AND FIRM AGE



reveals a relatively weak relationship between firm age and net growth when we exclude start-ups (there is a mild positive relationship without firm size effects and a mild inverse relationship with firm size effects). However, in the lower panel of figure 4, we find that conditional on survival, young firms exhibit substantially higher growth than more mature firms. This pattern is robust to controlling for firm size, and it clearly indicates that the fastest-growing continuing firms are young firms under the age of 5.

Reconciling the patterns of the upper and lower panel requires investigating the relationship between firm age and firm exit. That is, the firms not included in the lower panel of figure 4 relative to the upper panel are firm exits. Note that firm entrants are not driving the large difference in patterns across the upper and lower panel of figure 4 since they are not included in either panel. The relationship between firm age and job destruction from firm exit is reported in figure 5, where it is apparent that young firms have much higher firm exit rates than more mature firms do.

Taken together, figures 4 and 5 describe an up-or-out pattern for young firms that is robust to controlling for firm size (and robust to whichever size class method is used). This up-or-out pattern highlights that the net patterns by firm age depicted in the top panel of figure 4 mask the rich dynamics of young firms. This dynamic is an important feature of market-based economies and is consistent with predictions in models of market selection and learning (see Jovanovic, 1982; Hopenhayn, 1992; Ericson & Pakes,

FIGURE 5.—FIRM EXIT BY FIRM AGE




FIGURE 6.—ENTRY RATES BY SECTOR

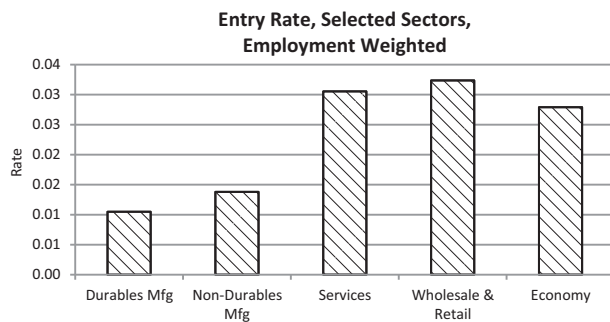
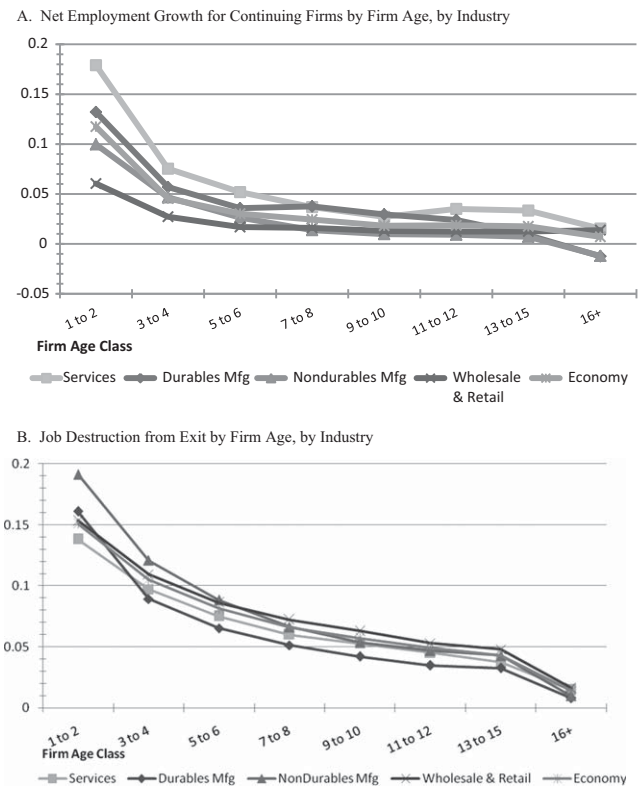


FIGURE 7.—UP-OR-OUT PATTERNS BY INDUSTRY



1995). It is also consistent with models where it takes time for firms to build up demand capital (Foster, Haltiwanger, & Syverson, 2010) or firms to build up reputation in credit markets (Evans & Jovanovic, 1989).

The up-or-out pattern of young firms also helps put the job creation from start-ups in perspective. Each wave of firm start-ups creates a substantial number of jobs. In the first years following entry, many start-ups fail (the cumulative employment weighted exit rate derived from figure 5 implies that about 47% of the jobs created by start-ups are eliminated by firm exits in the first five years), but the surviving young businesses grow very fast.²⁶ In this respect, the start-ups are a critical component of the experimentation process that contributes to restructuring and growth in the United States on an ongoing basis.

We check the robustness of the results in figure 4 by considering whether the patterns are potentially driven by large, young businesses. Although our measurement methods avoid creating new firms as the outcome of mergers and acquisitions, there are some large, young firms creating jobs, as seen in table 1. As seen in figure 1, the latter do not account for much of the contribution of firm births, but it is possible they

are high-growth firms contributing to the patterns in figure 4. To check on the contribution of such large firms to the analysis, we estimated the specifications underlying figures 2 through 5 restricting attention to firms that have fewer than 500 workers. We find the patterns in figures 2 through 5 are robust to considering only such firms in the analysis.²⁷

C. Firm Entry and Up-or-Out Dynamics by Sector

One question raised by the striking patterns in figures 4 and 5 is whether the up-or-out dynamics are driven by specific industries. It might be that the factors that yield such young firm dynamics are more important in some sectors than others. Moreover, firm entry rates vary across sectors, and the pace of entry may influence the nature of young firm dynamics.

Figure 6 shows employment-weighted firm entry rates by selected broad sectors.²⁸ Not surprisingly, sectors such as Services and Wholesale and Retail Trade have much higher entry rates than either Durable or Nondurable Goods Manufacturing. Firm entry rates are especially low in Durable Goods Manufacturing.

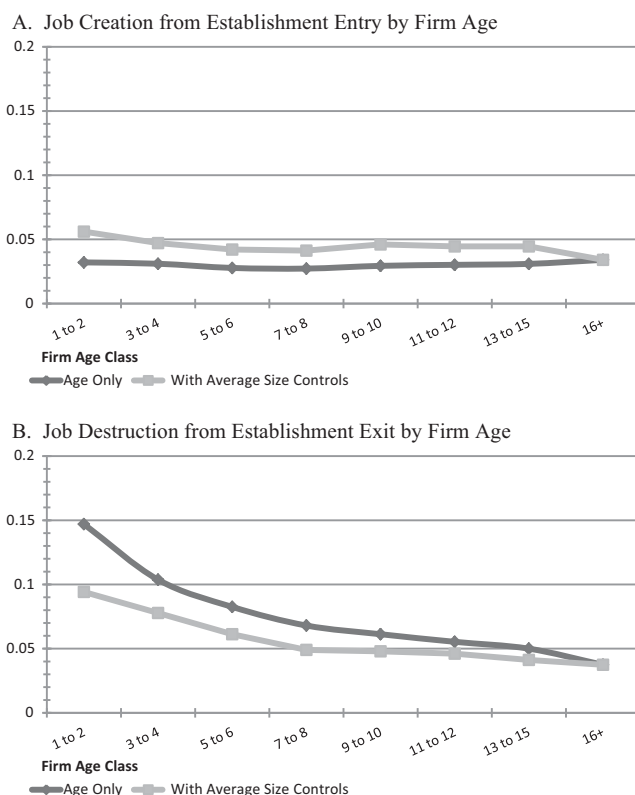
Figure 7 shows the up-or-out patterns by industry. It is striking that in spite of the large differences in entry rates, the patterns are so similar across industries. In all of the

²⁶ Without controlling for firm size effects, the growth from the survivors does not fully compensate for the exits. The cumulative net growth rate implied by figure 4 is about -16% in the first five years after entry. Note, however, that this still implies that five years after entry, a typical cohort has contributed a substantial number of jobs. When we control for firm-size effects, one can see that the cumulative net growth for firms less than 5 years old will be positive.

²⁷ The results for this robustness check are in the online appendix.

²⁸ A supplemental file available electronically includes all sectors, but we focus on selected sectors in the main text for the sake of brevity.

FIGURE 8.—ESTABLISHMENT ENTRY AND EXIT BY FIRM AGE

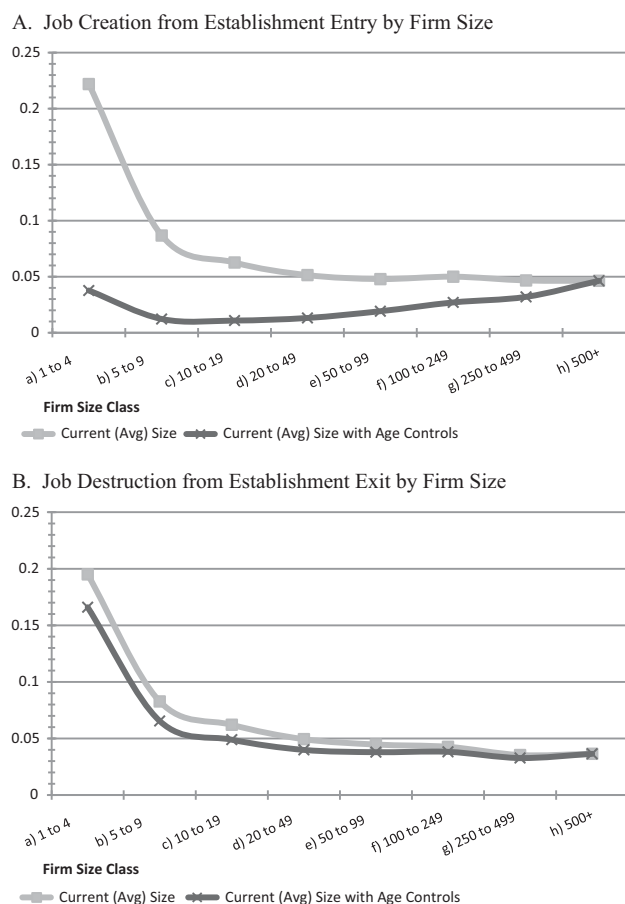


broad sectors, young businesses either grow fast or they exit. There are some interesting differences in the magnitudes of the patterns by sector. Young, continuing firms in the Service sector have especially high growth relative to firms in the Wholesale and Retail sector. In addition, there is a more pronounced decline in the exit pattern for young firms in Manufacturing Durable and Nondurable Goods. In unreported results we have found that the peak exit rate is not in the first year but in the second year after entry for these sectors. The pattern suggests that the nature of the learning and selection dynamics differs in the Manufacturing sector. Still, it is striking how similar the qualitative patterns are by sector. The implication of figures 6 and 7 is that while entry rates vary substantially across sectors, conditional on entry, the same rich up-or-out dynamics are present in every sector.

D. The Entry and Exit Margins: Establishment versus Firm

The focus thus far has been on firm entry and exit. In this section, we compare and contrast the patterns of firm entry and exit with the patterns of establishment entry and exit. Figure 8 shows establishment entry and exit by firm age. The upper panel shows that establishment entry exhibits a slight tendency to increase with firm age that is mitigated after controlling for firm size. By contrast, the lower panel shows job destruction from establishment exit falling

FIGURE 9.—ESTABLISHMENT ENTRY AND EXIT BY FIRM SIZE



monotonically with firm age. Taking these patterns together with those of figures 4 and 5 implies that young firms create jobs by expanding existing establishments rather than opening new ones.²⁹ Additionally, in comparing the establishment-level patterns in figure 8 to the firm-level patterns in figure 5, we see that job destruction from establishment exit declines less rapidly with firm age than does job destruction from firm exit. Mature firms are much less likely to exit than young firms, but establishments of mature firms have relatively high exit rates compared to establishments of young firms conditional on other observable factors (Dunne, Roberts, & Samuelson, 1989, also found this result for U.S. manufacturing).

Figure 9 shows relationships of establishment entry and exit by firm size. Establishment entry tends to fall with firm size when not controlling for firm age, and this reflects the obvious relationship between firm entry and establishment entry. However, after controlling for firm age, we observe establishment entry rising with firm size. The lower panel of figure 9 shows that job destruction from establishment

²⁹ In the online appendix, we also compare job creation by firm age with job creation by establishment entry by firm age, which makes this point more transparent.

exit tends to decline with firm size. These patterns are quite similar to those for firm exit in figure 3.

To sum up, there are notable differences in the patterns of firm versus establishment entry and exit by firm age and firm size. First, looking at establishments, we find that young firms are less likely to exhibit job creation from opening new establishments than are mature firms. Young firms, however, disproportionately create jobs by expanding existing establishments. Large firms are more likely to open new establishments when we control for firm age. Second, firm entry rates are much higher for the smallest size classes, but this simply reflects the fact that new firms tend to be small. Once we control for firm age, large firms are more likely to open new establishments than small firms are. Establishment exit is also more likely for smaller firms, and this holds even after controlling for firm age.

V. Conclusion

There is a widespread popular perception that small businesses create the most jobs in the United States. This perception has basis in empirical observation, but we demonstrate that the inverse relationship between net job growth and firm size disappears after controlling for firm age. To draw this conclusion, we take advantage of newly developed economy-wide longitudinal firm and establishment data available at the U.S. Census Bureau that permits accurate tracking of business start-ups, business exits, and continuing firms.

Our analysis focuses on measurement, not policy. However, measurement issues clearly can influence policy discussions, and our findings give those charged with policies aimed at the job creation much to consider. For example, to the extent that policy interventions aimed at small businesses ignore the important role of firm age, we should not expect much of an impact on the pace of job creation. Effective policy design in this area requires a richer understanding of business dynamics, as well as any relevant market failures. Our analysis addresses only the first issue.

We find some evidence in support of the popular perception that small businesses create most jobs along the following lines. If one looks at the simple relationship between firm size and net growth rates, there is evidence that net growth rates tend to be higher for smaller as opposed to larger businesses. This is the case using widely available data such as the Census Bureau's SUSB, but can also be obtained through careful analysis of microdata, as both NWZ and we demonstrate. Of course, the caveats raised over years, such as the role of regression to the mean, still apply. Using our preferred firm size classification that we argue is more robust to such concerns, the inverse relationship between net growth rates and size remains but is not overwhelming.

Our results show that the more important and robust finding is the role of firm age and its relationship with firm growth dynamics. We find that once we control for firm age,

the negative relationship between firm size and net growth disappears and may even reverse sign as a result of relatively high rates of exit among the smallest firms. Our findings suggest that it is particularly important to account for business start-ups. Business start-ups account for roughly 3% of U.S. total employment in any given year. While this is a reasonably small share of the stock, it is large relative to the net flow, which averages around 2.2% per year. We also find rich dynamics among young firms that help put the contribution and role of start-ups into perspective. Young firms exhibit high rates of gross job creation and destruction. Consistent with this pattern, we find that young firms have very high job destruction rates from exit, so that after five years, about 40% of the jobs initially created by start-ups have been eliminated by exit. However, we also find that, conditional on survival, young firms grow more rapidly than their more mature counterparts.

Most of our focus is on the net growth rate patterns by firm size and firm age (along with the underlying different margins of adjustment). However, we also show that large, mature businesses account for a large fraction of jobs. Firms that are over 10 years old and have more than 500 workers account for about 45% of all jobs in the U.S. private sector. In turn, we show that these large, mature firms account for almost 40% of job creation and destruction. The share of jobs created and destroyed by different groups of firms is roughly their share of total employment. An important exception in this context is the contribution of firm start-ups: they account for only 3% of employment but almost 20% of gross job creation.

We think our findings help interpret the popular perception of the role of small businesses as job creators in a manner that is consistent with theories that highlight the role of business formation, experimentation, selection, and learning as important features of the U.S. economy. Viewed from this perspective, the role of business start-ups and young firms is part of an ongoing dynamic of U.S. businesses that needs to be accurately tracked and measured on an ongoing basis. Measuring and understanding the activities of start-ups and young businesses, the frictions they face, their role in innovation and productivity growth, and how they fare in economic downturns and credit crunches are clearly interesting areas of inquiry given our findings. To the extent that market failures are found to underlie these frictions, there might be a role for well-designed corrective policies that help entrepreneurs start and grow dynamic young firms that boost overall net job creation.

More broadly, our findings suggest that the policy debate about encouraging private sector job creation should be refocused. The job-creating prowess of small businesses is often used by policymakers to motivate and justify specific policies. Our findings suggest that policies targeting firms based on size without taking account of the role of firm age are unlikely to have the desired impact on overall job creation. Taking the patterns of firm dynamics we show here into account may help identify the specific market failures

that prevent entrepreneurs from starting and growing new businesses.

In a related manner, it is important to focus not only on jobs per se but also on the role these dynamics play in productivity and earnings growth in the United States. Similarly we need to develop the data and associated analyses that will permit investigating the complex relationships between young and mature businesses. It may be, for example, that the volatility and apparent experimentation of young businesses that we have identified are critical for the development of new products and processes that are in turn used by (and perhaps acquired by) the large and mature businesses that account for most economic activity. Viewed from this perspective, our findings show that the LBD and the BDS are rich databases to track U.S. business dynamics, but it is also clear that additional information about the productivity and earnings dynamics, as well as business-to-business relationships, need to be added to these databases and related analyses.

REFERENCES

- Acs, Zoltan J., Catherine Armington, and Alicia Robb, "Measures of Job Flow Dynamics in the U.S. Economy" (Washington, DC: U.S. Small Business Administration, 1999).
- Aldrich, Howard, Arne Kalleberg, Peter Marsden, and James Casell, "In Pursuit of Evidence: Five Sampling Procedures for Locating New Businesses," *Journal of Business Venturing* 4 (1988), 367–386.
- Angrist, Josh, and Jorn-Steffen Pischke, *Mostly Harmless Econometrics* (Princeton, NJ: Princeton University Press, 2009).
- Becker, Randy A., John Haltiwanger, Ron S. Jarmin, Shawn D. Klimek, and Daniel J. Wilson, "Micro and Macro Data Integration: The Case of Capital," in D. Jorgenson, S. Landefeld, and W. Nordhaus (eds.), *A New Architecture for the U.S. National Accounts* (Chicago: NBER/University of Chicago Press, 2006).
- Birch, David L., "The Job Generation Process," unpublished report, MIT Program on Neighborhood and Regional Change for the Economic Development Administration, U.S. Department of Commerce (1979).
- "Who Creates Jobs?" *Public Interest* 65 (1981), 3–14.
- *Job Creation in America: How Our Smallest Companies Put the Most People to Work* (New York: Free Press, 1987).
- Birley, Sue, "Finding the New Firm," in *Proceedings of the Academy of Management Meetings*, 1984.
- Brown, Charles, James Hamilton, and James Medoff, *Employers Large and Small* (Cambridge, MA: Harvard University Press, 1990).
- Butani, Shail, Richard Clayton, Vinod Kapani, James Spletzer, David Talan, and George Werking, "Business Employment Dynamics: Tabulations by Employer Size," *Monthly Labor Review* (February 2006).
- Davis, Steven J., John Haltiwanger, Ron S. Jarmin, and Javier Miranda, "Volatility and Dispersion in Business Growth Rates: Publicly Traded vs. Privately Held Firms," in D. Acemoglu, K. Rogoff, and M. Woodford (eds.), *NBER Macroeconomics Annual 2006*, 21 (Cambridge, MA: MIT Press, 2007).
- Davis, Steven J., John Haltiwanger, C. J. Krizan, Ron Jarmin, Javier Miranda, Al Nucci, and Kristin Sandusky, "Measuring the Dynamics of Young and Small Businesses: Integrating the Employer and Non-Employer Businesses," in T. Dunne, B. Jensen, and M. Roberts (eds.), *Producer Dynamics: New Evidence from Micro Data* (Chicago: University of Chicago Press for the NBER, 2009).
- Davis, Steven J., John Haltiwanger, and Scott Schuh, *Job Creation and Destruction* (Cambridge, MA: MIT Press, 1996).
- Doms, Mark E., Ron S. Jarmin, and Shawn D. Klimek, "Information Technology Investment and Firm Performance in US Retail Trade," *Economics of Innovation and New Technology* 13 (2004), 595–613.
- Dunne, Timothy, Mark Roberts, and Larry Samuelson, "The Growth and Failure of U.S. Manufacturing Plants," *Quarterly Journal of Economics* 104 (1989), 671–698.
- Ericson, Richard, and Ariel Pakes, "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies* 62 (1995), 53–82.
- Evans, David S., "The Relationship between Firm Growth, Size, and Age: Estimates for 100 Manufacturing Industries," *Journal of Industrial Economics* 35 (1987), 567–581.
- Evans, David S., and Boyan Jovanovic, "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints," *Journal of Political Economy* 97 (1989), 808–827.
- Foster, Lucia, John Haltiwanger, and C. J. Krizan, "Market Selection, Reallocation and Restructuring in the U.S. Retail Trade Sector in the 1990s," this REVIEW 88 (2006), 748–758.
- Foster, Lucia, John Haltiwanger, and Chad Syverson, "The Slow Growth of New Plants: Learning about Demand?" mimeograph (2010).
- Friedman, Milton, "Do Old Fallacies Ever Die?" *Journal of Economic Literature* 30 (1992), 2129–2132.
- Haltiwanger, John, Ron Jarmin, and Javier Miranda, "Business Dynamics Statistics: An Overview," Marion Ewing Kauffman Foundation BDS briefs (2009). http://www.kauffman.org/uploadedFiles/BDS_handout_011209.pdf.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda, "Who Creates Jobs? Small vs. Large vs. Young," NBER working paper 16300 (2010).
- Haltiwanger, John, and C. J. Krizan, "Small Business and Job Creation in the United States: The Role of New and Young Businesses," in Zoltan Acs (ed.), *Are Small Firms Important? Their Role and Impact* (Norwell, MA: Kluwer Academic, 1999).
- Hopenhayn, Hugo, "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica* 60 (1992), 1127–1150.
- Jarmin, Ron S., Shawn D. Klimek, and Javier Miranda, "The Role of Retail Chains: National, Regional and Industry Results," in T. Dunne, B. Jensen, and M. Roberts (eds.), *Producer Dynamics: New Results from Micro Data* (Chicago: University of Chicago Press for the NBER, 2009).
- Jarmin, Ron S., and Javier Miranda, "The Longitudinal Business Database," CES working paper 02-17 (2002).
- Jovanovic, Boyan, "Selection and the Evolution of Industry," *Econometrica* 50 (1982), 649–670.
- Neumark, David, Brandon Wall, and Junfu Zhang, "Do Small Businesses Create More Jobs? New Evidence for the United States from the National Establishment Time Series," this REVIEW 93 (2011), 16–29.
- Okolie, Cordelia, "Why Size Class Methodology Matters in Analyses of Net and Gross Job Flows," *Monthly Labor Review* (July 2004).
- Sutton, John, "Gibrat's Legacy," *Journal of Economic Literature* 35 (1997), 40–59.
- Tornqvist, Leo, Pentti Vartia, and Yrjo Vartia, "How Should Relative Change Be Measured?" *American Statistician* 39:1 (1985), 43–46.