

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

**Business Dynamics in the National Establishment Time Series
(NETS)**

Leland D. Crane and Ryan A. Decker

2019-034

Please cite this paper as:

Crane, Leland D., and Ryan A. Decker (2019). "Business Dynamics in the National Establishment Time Series (NETS)," Finance and Economics Discussion Series 2019-034. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2019.034>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Business Dynamics in the National Establishment Time Series (NETS)

Leland D. Crane

Ryan A. Decker*

May 1, 2019

Abstract

Business microdata have proven useful in a number of fields, but the main sources of comprehensive microdata are subject to significant confidentiality restrictions. A growing number of papers instead use a private data source seeking to cover the universe of U.S. business establishments, the National Establishment Time Series (NETS). Previous research documents the representativeness of NETS in terms of the distribution of employment and establishment counts across industry, geography, and establishment size. But there exists considerable need among researchers for microdata suitable for studying business dynamics—birth, growth, decline, and death. We evaluate NETS in terms of its ability to corroborate key insights from the business dynamics literature with a particular focus on the behavior of new and young firms. We find that NETS microdata exhibit patterns of business dynamics that are markedly different from official administrative sources, limiting the usefulness of NETS for studying these topics.

*Both authors are at the Federal Reserve Board of Governors

We thank Joonkyu Choi, Sabrina Howell, Erick Sager, and conference and seminar participants at SGE 2019 and the Federal Reserve Board for helpful comments. We also thank Gray Kimbrough for making his graphing schema available.

The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors.

1 Introduction

Research based on business microdata has become increasingly important in economics in recent years. Such research can be difficult, however. The most easily available sources of business microdata, such as Compustat, do not cover the universe of private businesses (Compustat only covers publicly traded firms) and may be subject to significant selection problems (Davis et al. (2007)). The availability of comprehensive U.S. business microdata has increased significantly during the last decade due to efforts by statistical agencies; but access to these data is still costly, and prudent (and legally mandated) confidentiality restrictions limit the scope of research that can be conducted with such data. A prominent private sector data source has emerged, however, with nominal coverage of a significant fraction of the universe of U.S. business activity and without onerous publication scope restrictions. The National Establishment Time Series (NETS), a product of Walls & Associates, consists of longitudinally linked Dun & Bradstreet establishment-level data (with firm linkages) on business employment, industry, and location, as well as other variables of potential interest to researchers and policymakers.

In a companion paper, Barnatchez et al. (2017) explore the representativeness of NETS in the cross section, comparing the data to the Census Bureau's County Business Patterns (CBP) and Nonemployer Statistics (NES) and the Bureau of Labor Statistics' (BLS) Quarterly Census of Employment and Wages (QCEW). Static distributions of NETS data can be made reasonably comparable to official sources, on average, in terms of establishment size, industry, and geography cells, subject to important limitations arising from divergent counts of small establishments and, to a lesser extent, very large establishments. These differences may be due to imputation in NETS, difficulties with the firm/establishment distinction in the Dun & Bradstreet data, and the mismeasurement of employment at non-employer businesses. Barnatchez et al. (2017) conclude that NETS can be useful to cautious researchers who apply appropriate sample restrictions and investigate questions about static distributions of economic activity.

One finding of [Barnatchez et al. \(2017\)](#), however, is that the NETS data miss three key dynamic developments in the U.S. in recent years: the shale oil and gas boom that began in the mid-2000s, the post-2007 construction contraction associated with the reversal of the housing boom, and the swift decline in manufacturing employment after 2000. The failure of NETS to capture these dynamic industry events hints at potentially broader limitations of NETS along the time series dimension.

In the present paper, we investigate establishment growth and firm lifecycle patterns in NETS, including higher moments of growth distributions, motivated by key insights from the firm dynamics literature. We find that NETS data cannot replicate key empirical patterns of establishment and firm growth documented in comprehensive, official administrative data—the Census Bureau’s Longitudinal Business Database (LBD) and associated public-use product, the Business Dynamics Statistics (BDS).¹

A significant, though not sole, reason for these limitations appears to be the prevalence of data imputation in NETS, the effects of which are magnified in a dynamic setting. In 2014, employment is imputed for more than two thirds of establishments with fewer than five employees, while the imputation rate is more than one third for establishments with between five and nine employees. Even larger establishment classes have employment imputation rates close to 10 percent. Imputation is particularly prevalent among young firms: about 90 percent of firms aged zero or one have imputed employment data. Imputation can be particularly consequential in dynamic settings where multiple years of data must be relied upon for a single observation. We find that in 2014, the employment data for 10 percent of firms had been imputed for seven or more *consecutive* years. Imputation of sales data is even more prevalent, at rates around 80 percent among small firms and 95 percent among large firms. Nearly all of the sales data for establishments of multi-establishment firms are imputed.

¹We did not access LBD microdata for this project; rather, we rely on published results from the LBD as well as our own analysis of the publicly available BDS. All previously published LBD results we describe have undergone appropriate Census Bureau disclosure avoidance processes to ensure that no confidential information is disclosed.

More broadly, business-level employment data are surprisingly non-volatile in NETS. The distribution of firm employment growth rates is far less dispersed and skewed in NETS than in official data. Young firm growth, which existing literature shows is characterized by substantial dispersion and skewness, is particularly poorly captured in such a setting. NETS appears ill suited for the study of labor market flows, firm entry and exit, and business life-cycle dynamics, though careful use of NETS in case study settings may still be productive.

The paper proceeds as follows. In Section 2, we briefly describe the NETS data, related literature, and our data preparation methods. Given the prevalence of data imputation in NETS, in Section 3 we describe patterns of imputation with a focus on implications for business dynamics measurement. In Section 4 we compare NETS data with official sources in terms of aggregate patterns of firm dynamics, the geographic and industrial composition of firm growth, and the lifecycle behavior of firms. Section 5 is an argument for preferring official data to NETS when discrepancies between the two arise; some readers may prefer to start there. Section 6 concludes.

2 Data background and preparation

2.1 NETS background

[Barnatchez et al. \(2017\)](#) explain NETS data in detail; we refer the interested reader to that paper for more details while we provide a short summary here. For many years, Dun & Bradstreet (D&B) has actively sought to maintain a database of all business establishments in the U.S., which the firm uses in its business of selling marketing and other information. D&B collects these data from state secretaries of state, Yellow Pages, court records, credit inquiries, and direct telephone contact. Each year, D&B provides a snapshot of the establishment cross section to Walls & Associates, which creates longitudinal links and cleans the data for use by researchers and others. The finished data include annual establishment-level information on detailed industry, employment, sales, and other variables, with longi-

tudinal establishment linkages and firm identifiers to link the establishments of multi-unit firms. The Census Bureau's LBD, widely used by academic and government researchers, is similar in structure and aspiration to NETS, except that NETS seeks to track nonemployer businesses while the LBD is limited to employers (i.e., businesses with at least one paid employee). The NETS product to which we have access covers the years 1990-2014.

[Barnatchez et al. \(2017\)](#) review recent papers using NETS data for a variety of research questions.² Given our focus on business dynamics, here we describe just two key references. First, [Neumark et al. \(2005\)](#) evaluate the California sample of NETS through comparisons to QCEW. The authors recommend dropping establishments with one employee to approximate the employer universe; we adopt a modified (firm-level) version of this rule in our work. [Neumark et al. \(2005\)](#) also highlight the prevalence of imputation in the data and note that frequent imputation causes a low frequency of employment change at the establishment level. Most relevant to our purposes, the authors calculate employment growth at the county-by-industry level and study the correlation of employment growth between NETS and QCEW. Annual employment growth is weakly correlated between the two sources (0.528), so the authors study 3-year employment growth, which shows a correlation of 0.864. These aggregate exercises are useful and suggestive; we differ from [Neumark et al. \(2005\)](#) in focusing on the nationwide NETS sample and on a wider range of measures of business dynamics.

Separately, [Echeverri-Carroll and Feldman \(2017\)](#) evaluate the usefulness of NETS for studying entrepreneurship and firm entry by focusing on two specific case studies: the Austin-Round Rock (Texas) metropolitan statistical area and the North Carolina "Research Triangle." The authors compare NETS to Texas and North Carolina secretary of state (SOS) data (compiled by [Guzman and Stern \(2016\)](#)) and recommend restricting the data as follows: exclude known sole proprietorships (which do not appear in secretary of state data) and

²Some additional examples of recent work, not reviewed there, are [Heider and Ljungqvist \(2015\)](#), [Faccio and Hsu \(2017\)](#), [Chava et al. \(2018\)](#), and [Rossi-Hansberg et al. \(2018\)](#). [Cho et al. \(2019\)](#) match NETS to other establishment-level databases, as well as public use Economic Census files.

firms with nonprofit components, focus on headquarters establishments, and omit single-employee firms (as we do in the present paper and related work). With these restrictions, NETS data match secretary of state data for the two cities reasonably well, though there still exist significant discrepancies particularly in recent years of data. Importantly, the authors show that successive NETS vintages revise heavily for several recent years, so NETS reliability improves over time yet should be expected to be weak for the most recent years in the data (particularly the most recent four years).

A particularly notable contribution of [Echeverri-Carroll and Feldman \(2017\)](#) is that they match NETS microdata with SOS data for software startups in Austin, a painstaking process with large benefits for our questions here. They first exclude recent years of data to avoid vintage problems discussed above. They then seek to match about 3,500 NETS firms to the SOS data, first focusing on name and zip code matches, then relaxing to name matches only, using standard name generalization techniques. They successfully match about 40 percent of NETS firms to SOS firms. Among those matched, only 50 percent report the same entry year in NETS as in SOS data. About 75 percent have NETS and SOS entry years within two years of each other, and about 80 percent are within 3 years. The authors discuss reasons for the low match rate, which include missing legal form of organization data in NETS. The implications of this exercise are mixed, but the SOS data provide a degree of validation of NETS and suggest usefulness in limited exercises, particularly in case study settings similar to [Echeverri-Carroll and Feldman \(2017\)](#).

While [Echeverri-Carroll and Feldman \(2017\)](#) focus on carefully matched microdata comparisons within two specific case studies (Austin and the Research Triangle), we focus more broadly on comparisons using known average patterns of firm dynamics across the U.S. We will argue that NETS data are of limited usefulness for studying broad patterns of firm dynamics, leaving the [Echeverri-Carroll and Feldman \(2017\)](#) case study approach as the better (though more tedious and time intensive) use case for NETS.

2.2 LBD background

The longitudinal business database (LBD), housed at the U.S. Census Bureau, covers the near-universe of private nonfarm employer business establishments in the U.S. starting in the mid-1970s. It is constructed from the Census Bureau’s Business Register, the same source data as the County Business Patterns (CBP).³ The ultimate source data for the LBD, first described by [Jarmin and Miranda \(2002\)](#), draw from federal business tax records (both IRS and Social Security Administration), a variety of Census Bureau surveys, and the semi-decadal Economic Censuses (conducted in years ending 2 and 7). Importantly, the source data for the LBD include, by construction, all in-scope employer businesses in the U.S. that are known to tax authorities. The actual data consist of longitudinally linked establishment records with employment (as of March 12 of a given year), detailed industry, location, and other establishment characteristics. Establishment records also include firm identifiers that effectively group establishments under common operational control.

The LBD has become a critical resource for the study of firm dynamics. For example, [Davis et al. \(2007\)](#) first documented multi-decade declines in measures of firm-level employment volatility and gross job flows in the U.S. private sector using the LBD; the authors also linked the LBD to Compustat, a widely used dataset of publicly traded firms, and documented key differences in the behavior of publicly traded and privately held businesses. [Haltiwanger et al. \(2013\)](#) used the LBD to show that the job creation contribution that is widely attributed to small firms is more appropriately attributed to young firms. [Decker et al. \(2014\)](#) described key characteristics of young firms in the LBD, including “up-or-out” dynamics and high growth rate dispersion and skewness. [Alon et al. \(2018\)](#) show that cohort productivity growth declines with age and that high productivity growth of young firms is primarily a selection phenomenon. A further large literature exploits the LBD for studies of international trade, labor market flows, and a wide range of other topics.

While the LBD has become the primary resource for research on firm dynamics, it is

³[Barnatchez et al. \(2017\)](#) describe features of the Business Register that are relevant for comparisons with NETS. [DeSalvo et al. \(2016\)](#) describe the Business Register in exhaustive detail.

subject to strong confidentiality requirements and is therefore only accessible to sworn researchers with approved projects working in the Census Bureau or a Federal Statistical Research Data Center (FSRDC). Researchers using the LBD in FSRDCs must carefully follow rules to comply with federal law and prudent confidentiality concerns, and publishing results from statistical work on the LBD requires a lengthy process for disclosure avoidance. The process is generally costly and time consuming. Given the importance of the data, therefore, the Census Bureau publishes the publicly available Business Dynamics Statistics (BDS), which consists of aggregates of the LBD designed to track business dynamics at the level of sectors, firm age and size groups, and establishment locations. Research using the BDS has made considerable contributions to the literature. However, there are many questions that cannot be answered with the BDS, particularly questions about higher moments of the firm distribution and firm dynamics, that require microdata.

The limitations of the BDS and the tradeoffs involved with LBD access and use create demand for a public use source of business microdata like NETS. It is therefore important that researchers understand the strengths and limitations of NETS. The main purpose of this paper is to compare NETS with the LBD, with the latter serving as the benchmark against which any employer business microdata should be judged given its well-defined and near-universal scope and its wide use in the literature (see Section 5 for more discussion of official versus private data sources). We do not present any original results from LBD microdata; rather, we compare our original NETS calculations to existing LBD and BDS calculations from the literature.

2.3 NETS data preparation

2.3.1 Sample restriction

We first implement sample restrictions described in detail by [Barnatchez et al. \(2017\)](#). First, since NETS aspires to include both the nonemployer and employer universe, and since coverage beyond the employer universe is evident in the data, we restrict the sample to our best

guess of the employer universe by subtracting one employee from the employment of each firm headquarters establishment then dropping establishments with zero (post-subtraction) employment. This is a modified version of the sample restriction recommended by [Neu-mark et al. \(2005\)](#) and follows from the notion that owners are likely to be counted as employees in NETS though they may not be in official sources, where employment has a strict definition based on paycheck issuance. We restrict NETS to the employer universe to be comparable with the datasets to which we will make comparisons—the LBD and the BDS—which are both employer datasets. We then restrict the NETS sample to match the industry scope of the LBD and BDS, which is the same as the scope of CBP (see [Barnatchez et al. \(2017\)](#) for a specific industry scope list).

2.3.2 Establishment identifiers

Studying business dynamics is more complicated than studying cross-sectional snapshots of microdata. In particular, questions of business dynamics require careful attention to longitudinal linkages of business identifiers. Data problems (such as broken linkages) and real-world events like mergers and acquisitions generate challenges to longitudinal concepts and require researchers to make judgments. Given our goal of assessing the NETS data relative to official data, we attempt to treat the NETS data in a way that makes them most comparable to the LBD and the empirical firm dynamics literature based on the LBD.

The basic unit of observation in NETS is the *dunsnumber*. D&B views the *dunsnumber* as a line of business; but with respect to official sources, it is most similar to the concept of an establishment. In the LBD, an establishment is a single business operating location (identified by *lbdnum* in the LBD). In NETS, though, a single business operating location can have multiple *dunsnumbers*. This can be the case, for example, when the production operations and sales operations of a business are co-located but counted separately by D&B. [Barnatchez et al. \(2017\)](#) aggregate *dunsnumbers* to the establishment level to be consistent with CBP and QCEW establishment definitions; to do this, they identify *dunsnumbers* that

have the same reported firm headquarters (*hqduns*), 5-digit zip code, and first five street address characters (i.e., roughly speaking, same street and building number). This approach is designed to identify lines of business operating in the same location and falling under the same firm. They then sum the employment of the matched lines of business and assign the merged establishment a new identifier (termed the *locduns*) and the industry code of the largest line of business (in terms of employment). Since establishments in official data are assigned industry codes to reflect their principal activity, this method of merging D&B lines of business should roughly approximate the official concept. In practice, the line of business vs. establishment distinction seems to matter mostly for a small number of headquarters establishments.

We follow the approach of [Barnatchez et al. \(2017\)](#) for constructing establishment micro-data, but we introduce additional procedures for ensuring the longitudinal integrity of the resulting merged *locduns* establishment identifiers. A naive application of the [Barnatchez et al. \(2017\)](#) method could result in spurious changes in *locduns* establishment identifiers that reflect changes in the composition of establishment employment rather than the death of one establishment and birth of another. We first identify a *locduns* establishment as a continuer (i.e., not a birth or death) if there is a year-to-year overlap in at least one original line-of-business *dunsnumber*; that is, if a *locduns* disappears from the data we only assume the establishment has exited if all its associated line-of-business *dunsnumbers* cease to exist. We create a new identifier, the *netsnum*, that does not change from year to year even if a merged establishment's *locduns* changes due to changing employment composition of lines of business. In the (rare) case that lines of business that exist in the same location but belong to different firms (i.e., have different *hqduns*) in year $t-1$ move into the same firm (i.e., take on the same *hqduns*) in year t , we assign the year- t combined entity the *netsnum* of the year- $t-1$ *locduns* establishment that contributed the most employment (in terms of lines of business) to the new entity.⁴

⁴This is a rare occurrence because it suggests that two separate firms with establishments in the same building engaged in a merger or acquisition.

The resulting *netsnum* is a longitudinal identifier that is close in concept and spirit to the longitudinal establishment identifier in the LBD (*lbdnum*). We next focus on longitudinal firm identifiers.

2.3.3 Firm identifiers

A firm is a collection of establishments. The LBD defines the firm based on common operational control. The NETS firm concept is based on a common headquarters establishment (*hqduns*), where the *hqduns* refers to the *dunsnumber* of the headquarters establishment. NETS apparently allows for multiple levels of headquarters—perhaps including both regional and national headquarters—because we observe some cases in which an establishment record has a *dunsnumber* that is equal to other establishments' *hqduns*, but that itself refers to a different *hqduns*.⁵ That is, there are cases in which an establishment appears to be a headquarters for other establishments but does not refer to itself as its own headquarters. We attempt to unite all establishments that are related through headquarters, either directly or indirectly, under a single firm identifier by “rolling up” *hqduns* identifiers. That is, we assign all related establishments the *hqduns* of the highest level headquarters, which necessarily reports itself as its own *hqduns* (or, in rare cases, reports an *hqduns* that does not appear as a *dunsnumber* anywhere else).⁶

The firm identifier setup in NETS also presents the longitudinal challenge of determining which groups of establishments are successors to each other over time. As with the LBD's firm identifier (*firmid*), the *hqduns* can change for many reasons, including merger and acquisition activity but also simple data problems. Unlike in the LBD, in NETS the firm identifier automatically changes if the firm moves its headquarters from one establishment to another. We reassign *hqduns* firm identifiers as follows. For a given firm x in year $t - 1$, we identify all firms in year t that control at least some of firm x 's $t - 1$ establishments. We select

⁵There is some discussion of this in NETS marketing materials.

⁶In extremely rare cases, we observe headquarters linkages that are “cyclical;” for example, *dunsnumber* A reports *dunsnumber* B as its headquarters, while B reports A as its headquarters. In those cases, we arbitrarily assign an *hqduns* to apply to all related establishments.

x 's *candidate successor* as the firm which controls the plurality of employment at these continuing establishments. Very often, this firm has the same *hqduns* and essentially the same establishments as x , and there is no ambiguity. But when a firm "fractures" into several separate entities, it is sensible to match the source firm to the largest continuing fragment.

One more step is necessary to have consistent firm linkages. According to the rule above, it is possible for a single period t firm to be the candidate successor for two distinct period $t - 1$ firms. For example, a firm z in year t could include the largest continuing fragments of both firms x and y from year $t - 1$. This would be the case for an acquisition or a merger. In such a case we assume that z is the successor to whichever of x and y accounts for the largest share of employment at the new firm. The successor firm is assigned the same *hqduns* number as the source firm. Firms which lack a successor are assumed to have ceased to exist. This process is repeated year by year for the whole sample. This treatment of mergers has a number of limitations, though LBD firm identifiers are also not immune to merger problems and we accordingly follow best practice from the literature when we define firm age and growth rates.

We construct firm age to be consistent with the widely used convention from the literature (e.g., [Haltiwanger et al. \(2013\)](#)). At the first appearance of a new firm identifier (*hqduns*) in the data, we assign the firm the age of its oldest establishment (where establishment age is given by years since the first appearance of the establishment's *netsnum*, which is described above). Thereafter, the firm ages naturally. This approach abstracts from problems associated with spurious changes in the firm's headquarters identifier and is consistent with the convention used in the LBD-based literature to which we will compare NETS data.

2.3.4 Growth rate concepts

In various places below we report statistics based on firm or establishment employment growth rates. We employ the widely used growth rate concept of [Davis et al. \(1996\)](#) (the so-called "DHS growth rate"). Let $E_{e,t}$ be employment in year t for establishment e . Then

the establishment-level DHS growth rate is given by:

$$g_{e,t} = \frac{E_{e,t} - E_{e,t-1}}{0.5 \cdot (E_{e,t} + E_{e,t-1})}. \quad (1)$$

The DHS growth rate differs from standard growth rates by using average two-year employment in the denominator instead of simply employment in year $t - 1$. This growth rate concept has been widely used in the literature because it can accommodate entry (in which case, $E_{e,t-1} = 0$, $E_{e,t} > 0$, and $g_{e,t} = 2$) and exit ($E_{e,t-1} > 0$, $E_{e,t} = 0$, and $g_{e,t} = -2$).

While calculating establishment-level DHS growth rates is straightforward, calculating firm-level growth rates is more complicated due to the possibilities of mergers, acquisitions, and divestitures, which can generate changes in firm-level employment that do not reflect “organic” growth. Following [Haltiwanger et al. \(2013\)](#) and related literature, we focus on an “organic” growth concept that abstracts from such reorganizations. The firm-level organic growth rate for firm J is given by:

$$g_{J,t}^f = \frac{\sum_{e \in J} (E_{e,t} - E_{e,t-1})}{\sum_{e \in J} 0.5 \cdot (E_{e,t} + E_{e,t-1})}. \quad (2)$$

The summation subscript $e \in J$ limits the set of establishments being included to those that belong to firm J in year t , regardless of whether they belonged to firm J in year $t - 1$. That is, the firm growth rate is constructed as if all of the firm’s establishments in year t belonged to the firm in year $t - 1$ (even if they did not in reality belong to firm J in year $t - 1$), and any establishments controlled by firm J in year $t - 1$ that were divested to a different firm between $t - 1$ and t do not affect the growth rate of firm J . Establishments controlled by firm J in year $t - 1$ that exited (i.e., failed or closed) between $t - 1$ and t do contribute to measured growth, with $E_{e,t-1} > 0$ and $E_{e,t} = 0$ as mentioned above.⁷

⁷It is straightforward to show that $g_{J,t}^f$ is equivalent to the employment-weighted average of the growth rates of all of the firm’s year- t establishments (and exiters), where the employment weight is defined as average two-year employment as in the DHS denominator above.

3 Imputation

3.1 Employment imputation

NETS data include an imputation flag (*empc*) that takes on four possible values: (0) actual figure, (1) bottom of range, (2) D&B estimate, and (3) Walls & Associates estimate. The first two categories ($empc \in \{0, 1\}$) indicate values reported to D&B by survey respondents, with the “bottom of range” category indicating that the respondent reported a range rather than a specific count. D&B uses proprietary cross-sectional imputation methods in cases of non-reporters ($empc = 2$). Walls & Associates estimates employment for all non-reporters in each year using a longitudinal imputation method; in cases where this longitudinally imputed estimate differs from the D&B cross-sectionally imputed method, Walls & Associates inserts their own estimate and sets $empc = 3$. The Walls & Associates method uses regressions based on the time series of sales and employment for the establishment and its industry.⁸ We consider all values of *empc* except $empc = 0$ to be imputed, where the imputation can be done by the respondent ($empc = 1$), D&B ($empc = 2$), or Walls & Associates ($empc = 3$).

[Barnatchez et al. \(2017\)](#) report *establishment* imputation rates by establishment size. Imputation is prevalent, particularly among small establishments. For the year 2014, establishments without exact reported employment values comprise more than two thirds of establishments with fewer than 5 employees, more than one third of establishments with 5 to 9 employees, and more than 15 percent of establishments with 10 to 19 employees. Imputation rates are around 10 percent for all remaining size classes except establishments with 1000 or more employees, of which about 15 percent lack exact employment values. [Barnatchez et al. \(2017\)](#) conclude that the imputation problem is nontrivial but can be managed by omitting small establishments, which is also where NETS differs most markedly from official data.

That said, we also find evidence that data reporters implicitly impute some data by rounding their reported employment figures, leading to a potential understatement of true

⁸NETS imputation details are provided with NETS marketing materials, *Understanding Data in the NETS Database* (2009).

imputation rates in NETS. Figure 1 reports the distribution of last-digits of reported employment numbers among non-imputed ($empc = 0$) Walmart establishments in the year 2000.⁹ In that year, 88 percent of Walmart establishments' employment data are reported as not being imputed; that is, they are coded with $empc = 0$. Yet we observe overwhelming prevalence of employment figures ending in 0 or 5, suggesting that respondents are rounding. This kind of rounding by respondents is a well-known issue in the survey methodology literature. We see more reasonable last-digit distributions among establishments generally, yet within this large firm there appears to be significant rounding. This kind of rounding may have little cost in static or cross sectional settings, but below we make the case that the cost is higher in dynamic research.

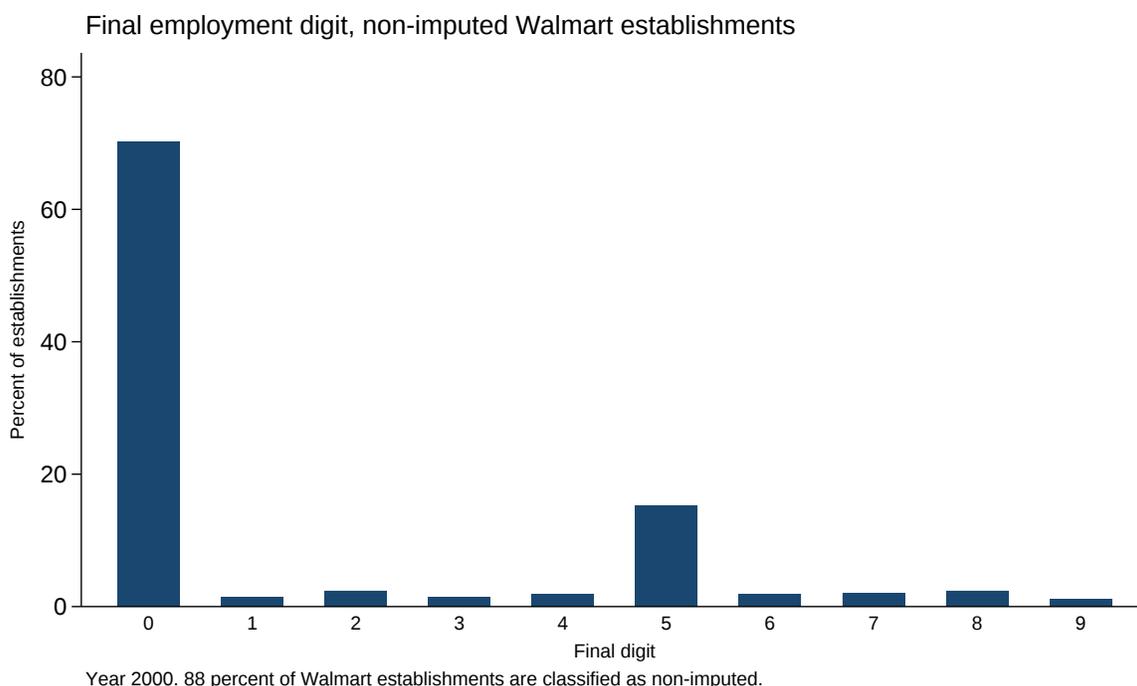


Figure 1

Given our focus on firm dynamics, we also explore *firm* imputation rates. Figure 2 reports the share of firms with imputed employment data, where the presence of *any* imputed

⁹Our graphing schemes are based on [Kimbrough \(2018\)](#).

establishments within a firm results in the firm counting as imputed (and establishments count as imputed if $empc \neq 0$). The solid blue line reports the share of firms that count as imputed, while the dashed green line reports the employment-weighted imputation rate (that is, the total employment—imputed or not—of imputed firms divided by total NETS employment). In early years, about half of NETS firms are imputed, but this share rises above two thirds by the end of the sample. Weighted imputation rates—the share of employment that is at imputed firms—are more steady, suggesting that the recent rise in unweighted imputation is primarily driven by smaller firms.

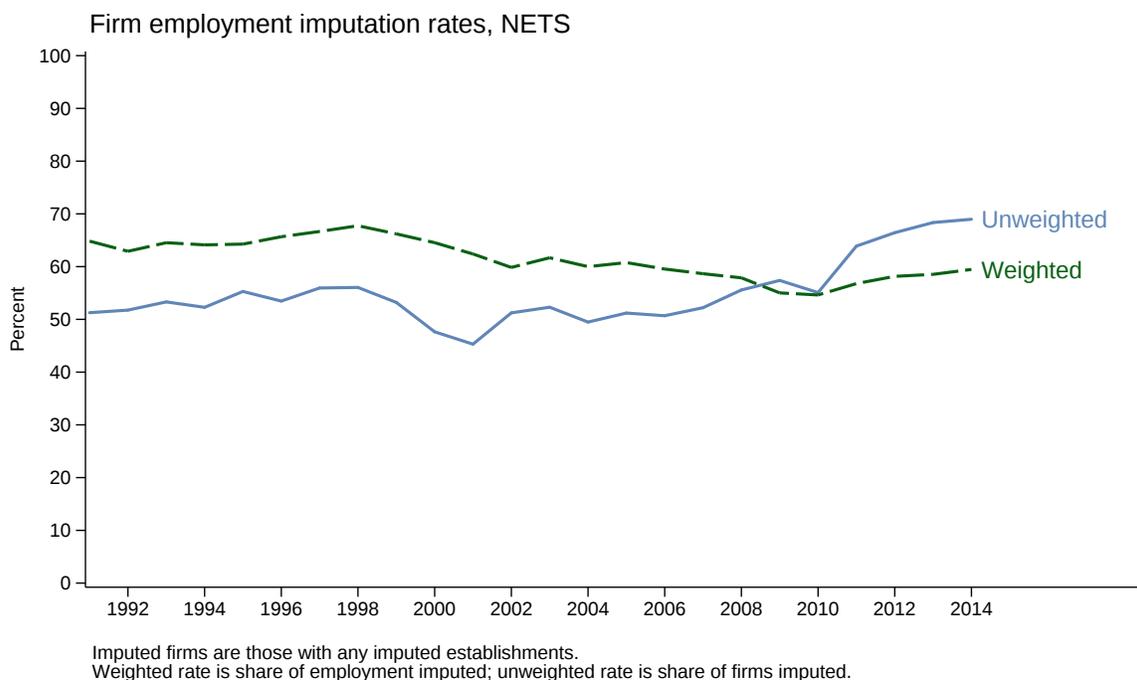


Figure 2

Figure 3 uses a more restrictive (and NETS-friendly) definition of “imputed”: we count firms as imputed if and only if at least 10 percent of their employment is at imputed establishments. This has no visible effect on the unweighted imputation rate, but the weighted rate moves down. The relaxation of the imputation definition has little effect on the un-

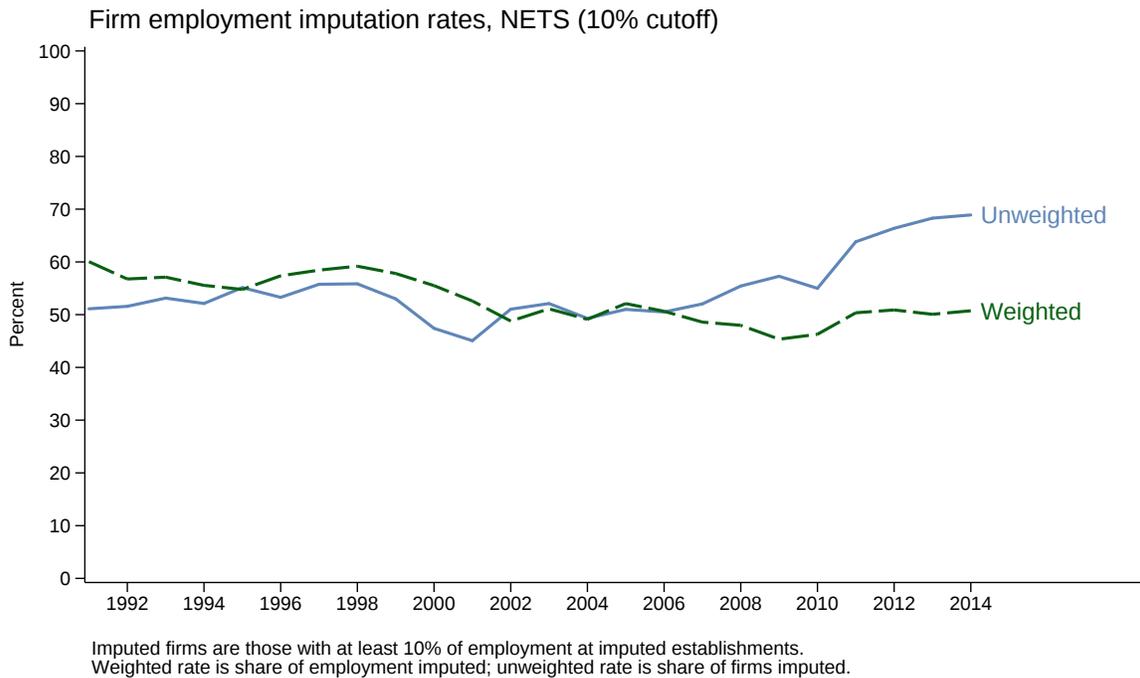


Figure 3

weighted imputation rate because it primarily affects a relatively small number of large firms. Figure 4 further relaxes the imputation standard, defining as imputed only those firms with at least half of their employment at imputed establishments. Figure 5 goes to the extreme, counting as imputed only firms with 90 percent or more of their employment at imputed establishments. Even with this excessively liberal definition, about one fifth of employment is at imputed firms.

Firm imputation is therefore nontrivial. Imputation may cause only limited problems for cross-sectional studies, but there are several reasons imputation is much more costly in research on dynamics. First, the longitudinal imputation method of Walls & Associates necessarily uses data on the establishment time series, implicitly assuming that past and future behavior is indicative of present behavior and thereby dampening dynamic volatility. Moreover, Walls & Associates rely on industry and other data that may serve to minimize

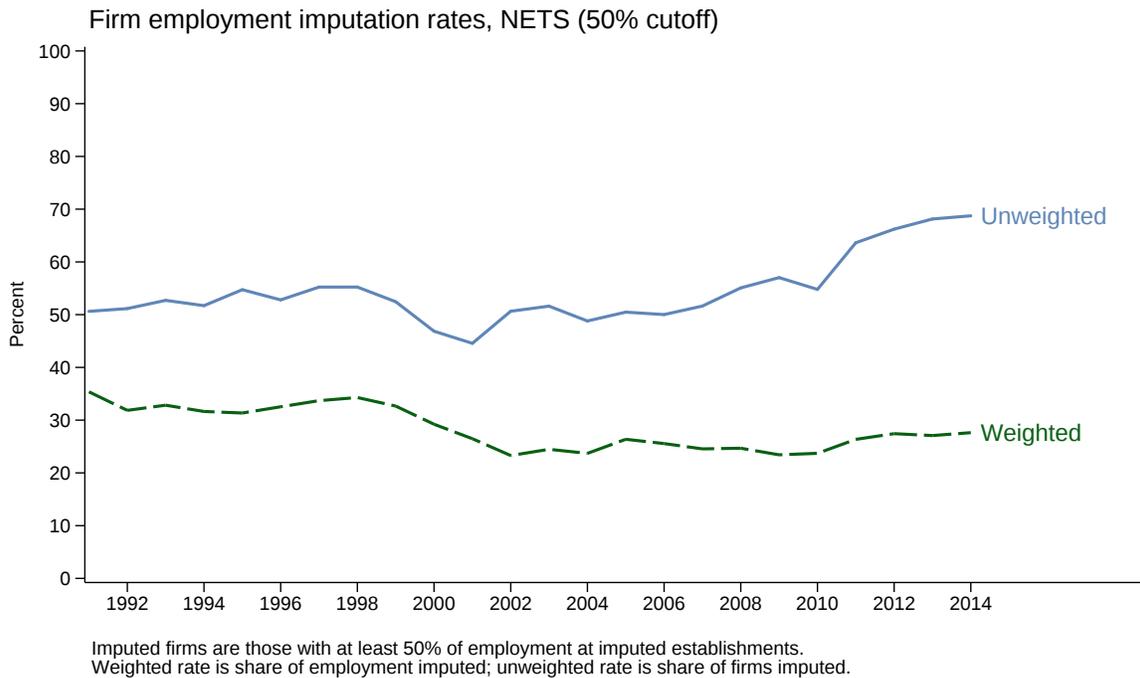


Figure 4

the dispersion of measured outcomes. Second, measures of dynamics depend on multiple consecutive data observations such that imputation is magnified. Concretely, employment growth from year $t - 1$ to year t depends on employment levels in years $t - 1$ and t ; if either year's employment value is imputed, the overall employment growth value is necessarily imputed. Third, in the case of *firm* (rather than *establishment*) dynamics, imputation of any establishments within a multi-unit firm implies that the overall firm employment value is necessarily imputed. We find that this problem is particularly salient among firms with many establishments.

A striking way to see the longitudinal costs of imputation is to consider imputation spells. We define the imputation spell as the number of consecutive years that a firm counts as imputed. For example, suppose a firm first counts as imputed in 1995. Then in 1995, the firm's imputation spell is equal to 1. If the firm is again imputed in 1996, then in that

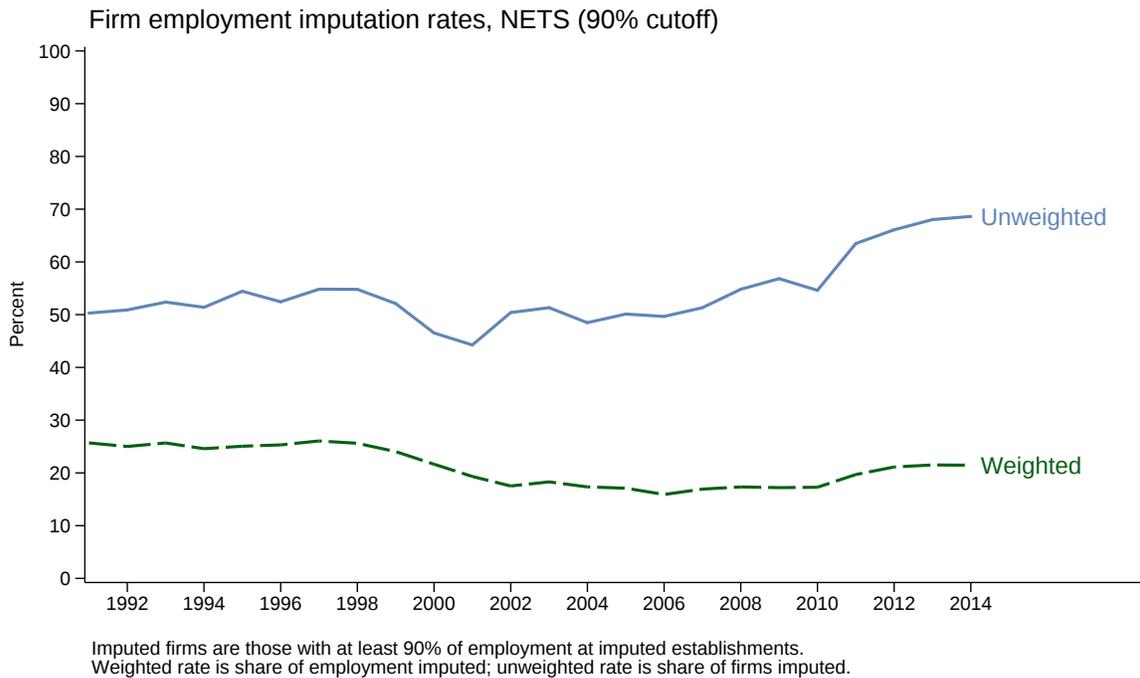


Figure 5

year its imputation spell is equal to 2. If the firm is not imputed in 1997, then its imputation spell in that year resets to 0. Figure 6 characterizes the distribution of imputation spells, where we count a firm as imputed if any of its establishments are imputed. The solid green line (the highest line) reports the 90th percentile imputation spell. For example, in 1998, the 90th percentile firm had an imputation spell of 8, meaning that 10 percent of firms had been imputed for 8 or more consecutive years. The median firm had an imputation spell of zero for most years in the sample, but by the end of the sample the median had risen to 2 years. Figure 7 reports the same exercise but restricts the sample to imputed firms in any given year; that is, the figure reports the distribution of imputation spells *conditional on* firms being imputed, rather than including non-imputed firms. Among imputed firms, even the 25th percentile reflects multiple consecutive years of imputation in many years, the median firm bounces between 2-year and 4-year imputation spells, and the 75th percentile shows

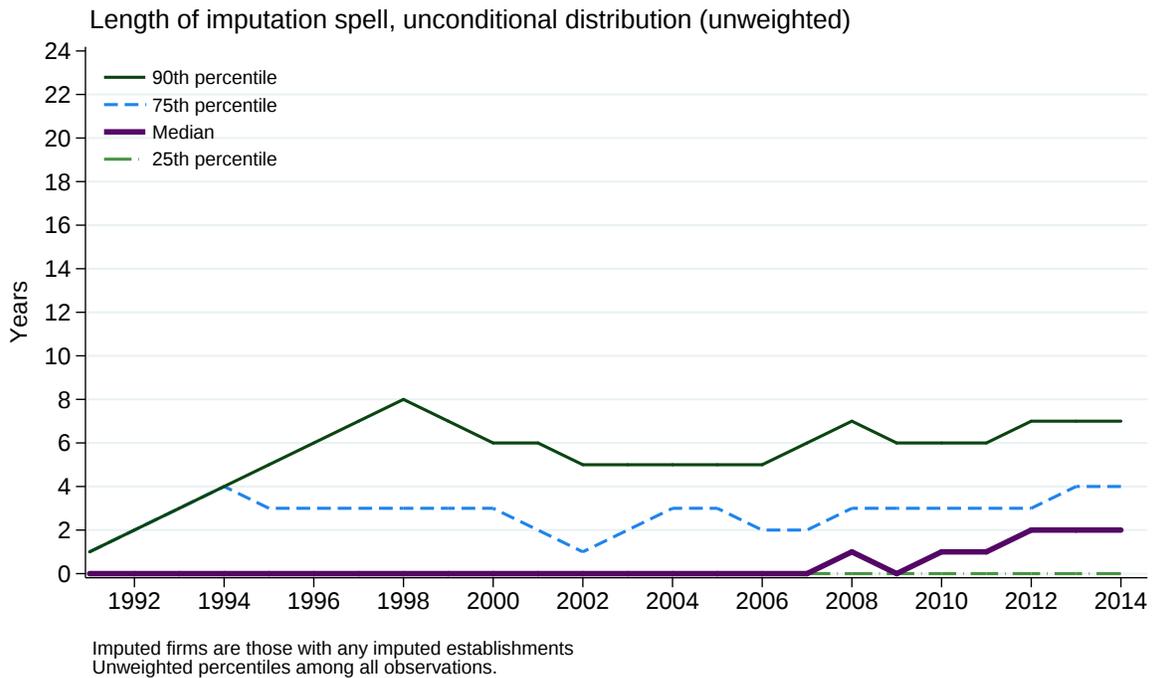


Figure 6

imputation spells between 4 and 6 years.

The problem of consecutive imputation is particularly pronounced among large firms. Figure 8 reports the employment-weighted distribution of imputation spells, again including all firms (i.e., not just conditional on imputation). The 90th percentile of the weighted distribution has the maximum possible imputation spell throughout most of the sample (i.e., a spell of imputation beginning at the origin of the sample), as does the 75th percentile. This means that 25 percent of overall employment is at firms that have been imputed for the maximum possible number of consecutive years. Figure 9 reports the same weighted, unconditional distribution with a more relaxed definition of firm imputation in which a firm counts as imputed if 25 percent or more of its employment is at imputed establishments. The picture improves some here, yet still 10 percent of employment is at firms with 7 years or more of consecutive imputation.

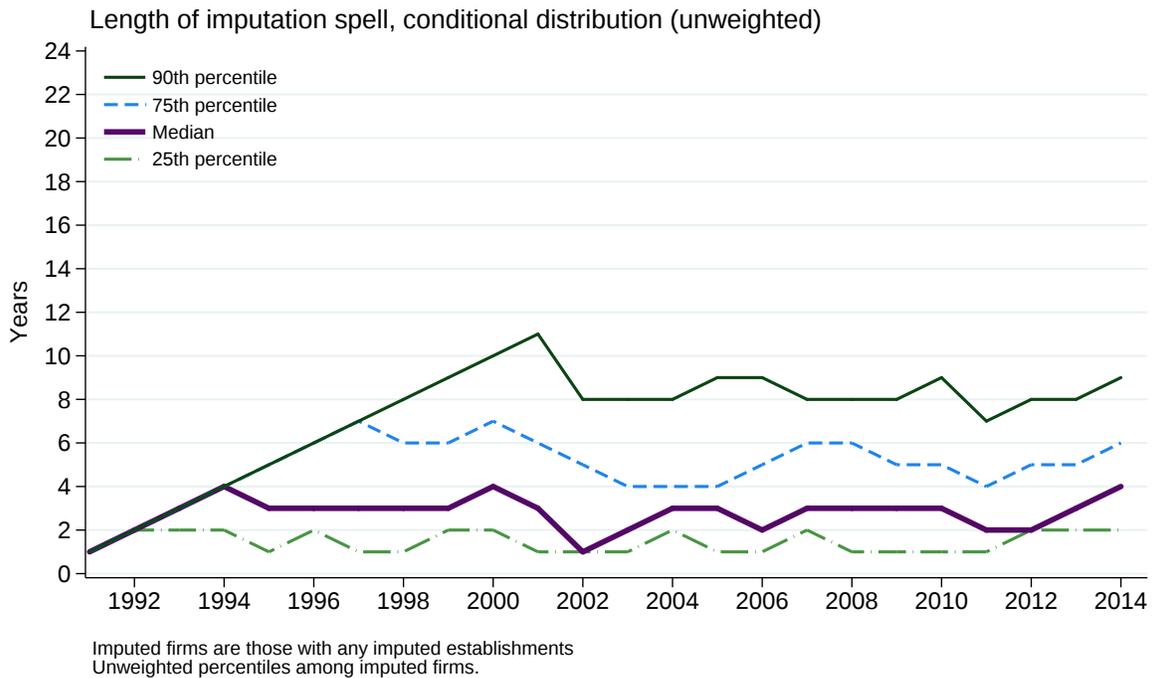


Figure 7

Needless to say, the longitudinal integrity of data in which substantial shares of activity reflect firms whose data have been imputed for multiple consecutive years is limited.

We develop one other imputation measure to track longitudinal imputation on a year-to-year basis. For the rest of the paper, we define a firm as being longitudinally imputed in year t if it counts as imputed in *either* year t or year $t - 1$. This definition is highly relevant when studying year-to-year firm-level growth or dynamics; as noted above, in a dynamic setting imputation binds in two consecutive years even if only one of the years has imputed data. We find that longitudinal imputation varies by firm age in critical ways. For example, Figure 10 reports longitudinal imputation rates by firm age for two snapshot years, 2003 and 2014. As noted elsewhere and in existing literature, the most recent year of data sees particularly acute imputation problems. But even in revised data, imputation is extremely prevalent among young firms, with rates above two thirds prior to age 3. These high imputation rates

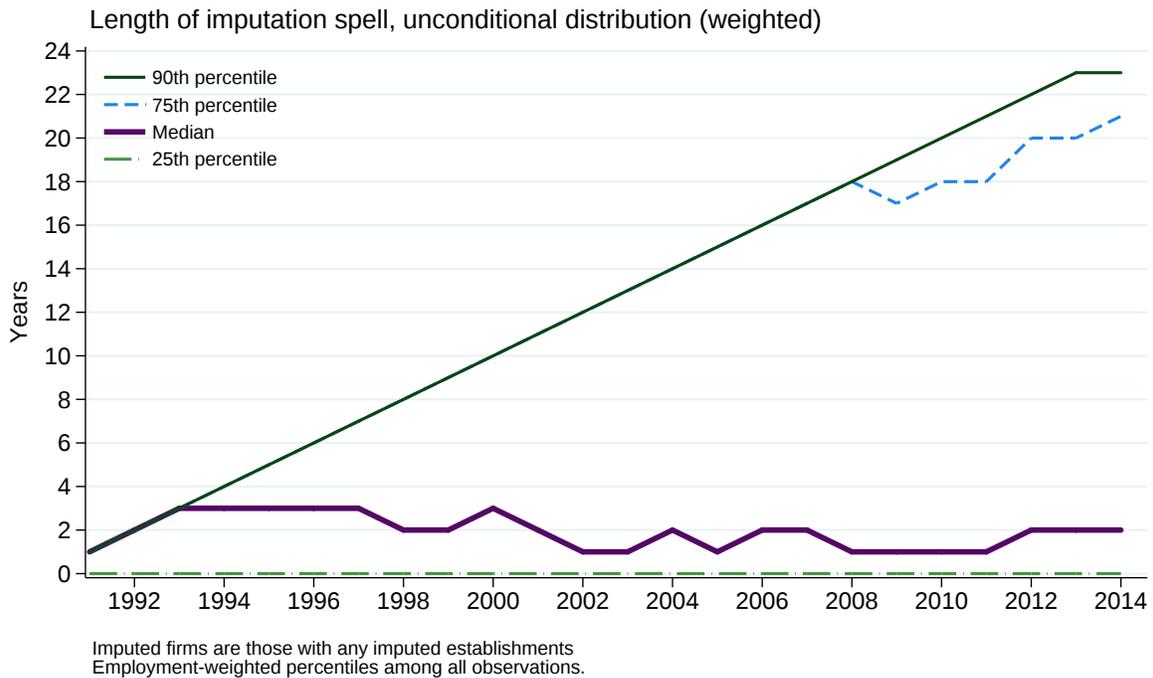


Figure 8

among young firms will prove to be problematic in the exercises below.

3.2 Sales imputation

Recent literature in firm dynamics relates firm employment growth with firm productivity (Decker et al. (2018), Alon et al. (2018)) by calculating real sales per worker at the firm level.¹⁰ While we focus primarily on employment dynamics in this paper, it is useful to briefly mention sales imputation.

The sales variable in NETS is somewhat more complicated than the employment variable.¹¹ While a respondent may be able to report current point-in-time employment to D&B surveyors at any time, a respondent is not likely to know current-year sales at the time

¹⁰This literature follows the construction of the Revenue-Enhanced LBD (RE-LBD) by Haltiwanger et al. (2017), which linked firm revenue data from the Census Bureau’s Business Register to the LBD.

¹¹This paragraph draws heavily on Walls (2008).

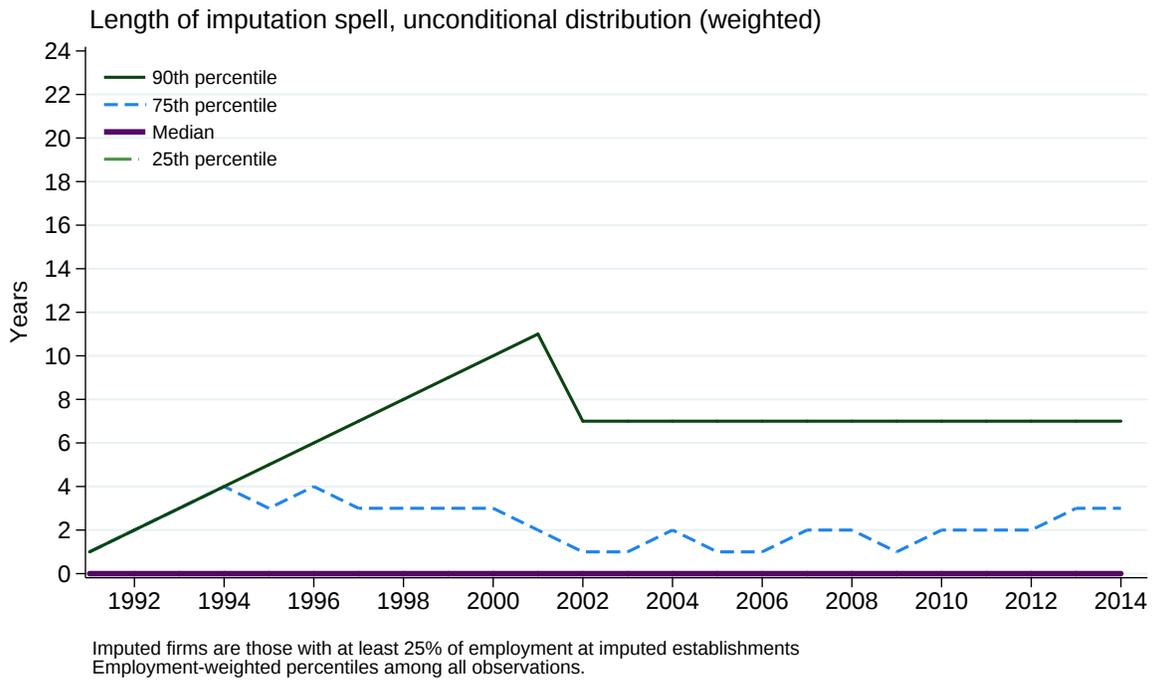


Figure 9

of contact. NETS documentation suggests that respondents are likely to report some combination of the prior year’s sales and an estimate of the current year’s sales. Moreover, establishment-level sales is a more complicated object than firm-level sales (indeed, Census Bureau researchers who bring sales data from the Business Register to the LBD study sales at the firm level). Therefore, reported establishment sales figures are estimates at best. For non-reported sales figures, D&B and Walls & Associates rely on imputation methods that are similar to those used for employment (described above), including reliance on firm or industry sales/employment ratios, with some additions. In particular, for multi-establishment firms, when firm-level sales are known (e.g., in the case of publicly traded firms), sales are allocated among establishments using employment shares. Note that sales figures are attributed even to establishments that do not sell products or services but instead produce inputs used by other establishments within the same firm; in such cases, the establishment



Figure 10

sales data provide no marginal information beyond the establishment employment data.

NETS does include an imputation flag for the sales variable, *salesc*, with the same coding as the employment imputation variable (*empc* described above). That is, *salesc* can take on the following values: 0 (actual figure or estimate provided by respondent), 1 (bottom of range reported by respondent), 2 (D&B estimate), or 3 (Walls & Associates estimate).¹² Imputation is common. In both the years 2000 and 2014, just under 20 percent of establishments report *salesc* = 0, indicating that the sales figure is a true reported value or respondent estimate. This likely overstates the accuracy of the figures, however, for the reasons above—even reported sales figures may be rough estimates. In any case, these establishments account for only about 10 percent of total (imputed and non-imputed) employment and total sales (imputed and non-imputed) in both years, indicating that imputation is *more*

¹²We also observe an extremely small number of establishments with missing sales data and sales imputation flags.

common among larger establishments. Remaining establishments are imputed, almost entirely reflecting imputation by Walls & Associates ($salesc = 3$).

| Firm size (employees) | Year | |
|-----------------------|------|------|
| | 2000 | 2014 |
| 1 to 4 | 80 | 80 |
| 5 to 9 | 78 | 85 |
| 10 to 19 | 77 | 82 |
| 20 to 49 | 79 | 84 |
| 50 to 99 | 85 | 88 |
| 100 to 249 | 89 | 91 |
| 250 to 499 | 93 | 94 |
| 500 to 999 | 94 | 94 |
| 1,000 to 2,499 | 93 | 93 |
| 2,500 to 4,999 | 95 | 92 |
| 5,000 to 9,999 | 95 | 94 |
| 10,000+ | 96 | 94 |

Source: NETS

Notes: Percent of firms with imputed establishment sales data.

Table 1: Establishment sales imputation rates

Sales imputation varies widely by *firm* size. Table 1 reports *establishment* imputation rates by *firm* size bins for the years 2000 and 2014. Small firms have establishment imputation rates around 80 percent, while around 95 percent of establishments of large firms have imputed data. The high imputation rates among large firms appear to be driven by firms with multiple establishments; in results not shown, we find that close to 95 percent of establishments of multi-establishment firms have imputed sales data, compared with about 80 percent among single-establishment firms. The interpretation of these imputation rates is not entirely clear. For example, there may be cases (particularly among publicly traded firms) where D&B receive accurate firm-level sales data, but establishment-level sales data must be imputed. Since our NETS data do not provide firm-level sales information, if we require firm sales figures we must construct them by summing across establishments within firms. So it is possible that the imputation rates we report for establishments of large firms

overstate the rate of imputation of firm-level sales among large firms; that is, there may be cases where a firm's establishments have sales data imputed from total firm sales such that summing across establishments results in true firm sales figures. However, the number of firms for which D&B receive true sales data is probably small (for example, there are fewer than 5,000 publicly traded firms in the U.S.), so if there is some overstatement, it is likely to be minimal. Moreover, the research for which establishment-level microdata like NETS would be most useful are likely to require geographic breakdowns of activity, in which case establishment imputation is the most relevant. In any case, establishment imputation rates are high across the firm size distribution, even among small firms that are likely to have only one establishment.

Sales data would be particularly useful for the study of productivity; however, we find large discrepancies between NETS and official data on this topic. For example, using the LBD, [Decker et al. \(2018\)](#) find that the within-industry dispersion (standard deviation) of sales per worker has risen in recent decades; in NETS, we find the opposite pattern. Moreover, the *level* of labor productivity dispersion is much lower in NETS than in the LBD, likely owing to the industry average rules of thumb used for NETS sales imputation. For example, [Decker et al. \(2018\)](#) find that among young (age less than five) high-tech firms, a firm that is one standard deviation more productive than its corresponding industry-by-year mean is about 2.5 times as productive as the mean in 1996 (the first year in which LBD sales data are available) and 3.0 times as productive in 2012. In NETS, this ratio is about 1.8 in 1996 and 1.7 in 2012.

The prevalence of sales imputation—which is more common than employment imputation—and the reliance of the imputation methods on employment data imply that the marginal value of the sales data is very low. Moreover, popular business dynamics topics such as productivity dispersion, decompositions of aggregate productivity growth, or the relationship between business-level productivity and growth cannot be studied with NETS.

4 Business dynamics in NETS and official data

4.1 Aggregate patterns

We first characterize the NETS data in terms of aggregate measures that are well known in the business dynamics literature. Figure 11 reports the share of firms that have age zero (often referred to as the startup rate or entry rate). The dashed green line reports the entry rate from NETS, while the blue solid line reports the entry rate from the BDS. The NETS entry rate is more volatile than the official data, though in many years the NETS rate bounces around the BDS rate. NETS sees an excess surge in entry in 2002 then again in 2011, consistent with the finding of [Barnatchez et al. \(2017\)](#) that NETS sees its establishment count surge above the levels of official data after 2000, which we believe likely reflects an expansion of D&B scope or coverage rather than true entry.

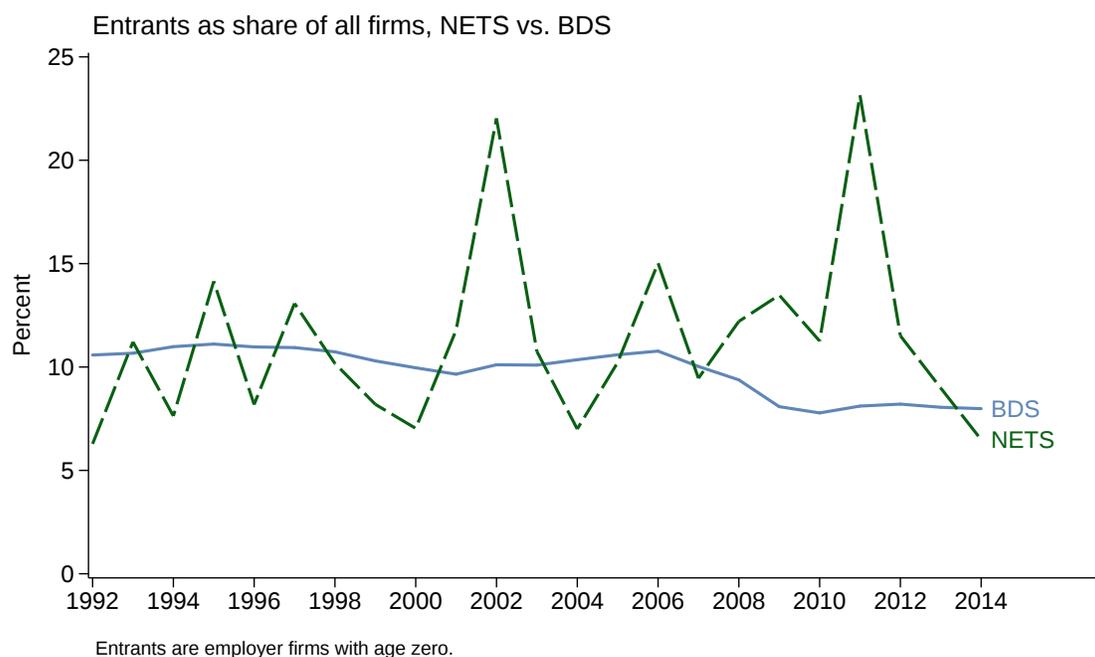


Figure 11

Figure 12 broadens our study of young firm behavior to include all firms of age less than five, a cutoff that is common in the literature. Here we report the young firm *employment* share. The surge in new businesses appearing in NETS but not in the official data is readily apparent here, with a divergence starting in 2008 and the cumulative effects of differing coverage becoming notable by the end of the sample. Importantly, the well-documented decline in young firm formation and activity described in a large and growing literature (Decker et al. (2014)) is reversed in NETS data due to this late-2000s divergence. While official data show young firm shares moving below 10 percent by 2010, young firm shares in NETS exceed 16 percent in 2012 and 2013, a level not seen in official data since the 1980s.

In short, while a large and growing literature explores the puzzling decline of young firm activity in official data, NETS data tell the opposite story. We believe this is likely due to spurious measured entry brought on by an apparent scope expansion.¹³ Barnatchez et al. (2017) plot total employment in the NETS employer universe against total employment in County Business Patterns (see their Figure 1); the shape of the gap between total NETS employment and total CBP employment documented by Barnatchez et al. (2017) closely mimics the shape of the gap between NETS young firm shares and BDS young firm shares shown on Figure 12.

We next study patterns of gross job flows; first, we define “job creation” as the number of jobs created by entering or expanding establishments, and we define “job destruction” as the number of jobs destroyed by exiting and downsizing establishments (these definitions are consistent with the literature). We express each of these as a rate by dividing by total employment, averaged over years t and $t - 1$ in usual DHS fashion. Figure 13 reports the job creation rate from the BDS (solid blue line), NETS (dashed green line), and NETS omitting firms with longitudinal imputation (dot-dashed purple line). Figure 14 reports the job

¹³NETS marketing materials point out the rise in entry and argue that this reflects growth of self employment or gig economy work brought on by changes in the nature of entrepreneurship and the weak labor markets of the Great Recession and aftermath. As noted above, we drop firms with only one reported employee, which should roughly eliminate true nonemployers from the data. Thus, the discrepancy we observe reflects apparent differences in measured employment at employer businesses.

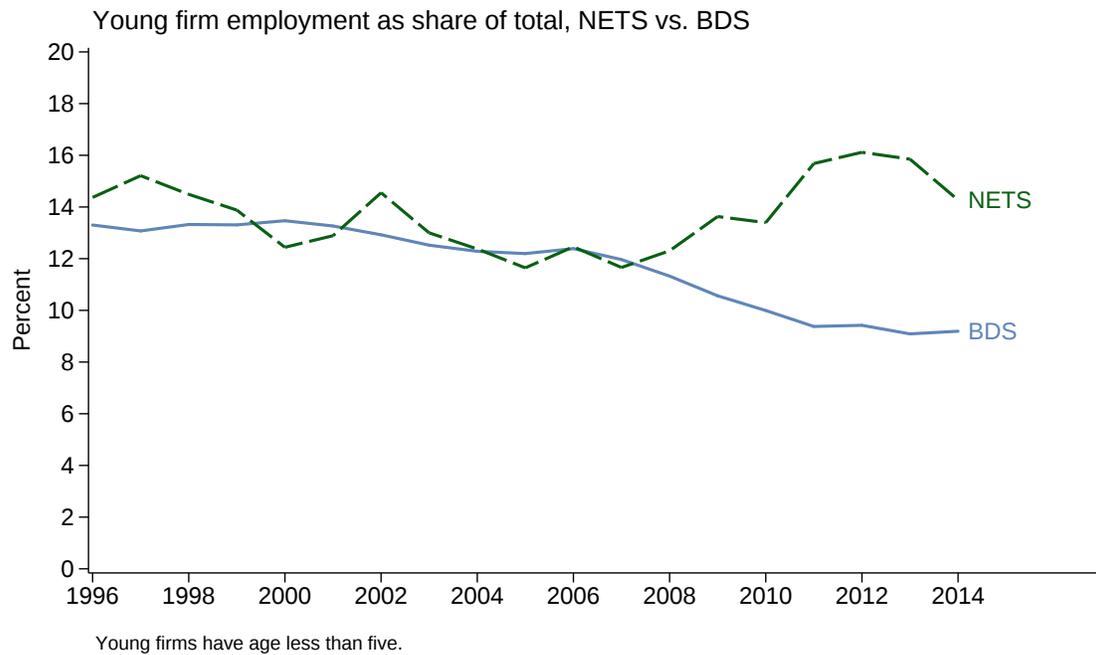


Figure 12

destruction rate.

In general, NETS exhibits much lower rates of gross job flows than the official data, as one might expect given the foregoing discussion of imputation and rounding. But it is somewhat surprising that the non-imputed NETS series are sometimes even lower than the overall NETS series, suggesting that imputation alone does not explain the low volatility of NETS firms. One likely reason is that, as shown on Figure 10, imputation is most prevalent among young firms. Dropping imputed firms means shifting the firm distribution heavily toward more mature firms that tend to have lower job creation rates. Other problems arise from the simple fact that dropping imputed firms significantly reduces the sample, and likely in a non-random way, so any statistics calculated on the residual sample are biased. The job creation rate patterns of the late years in the sample, when the NETS and NETS-without-imputation series diverge, likely reflect the surge in firm entry seen in NETS since

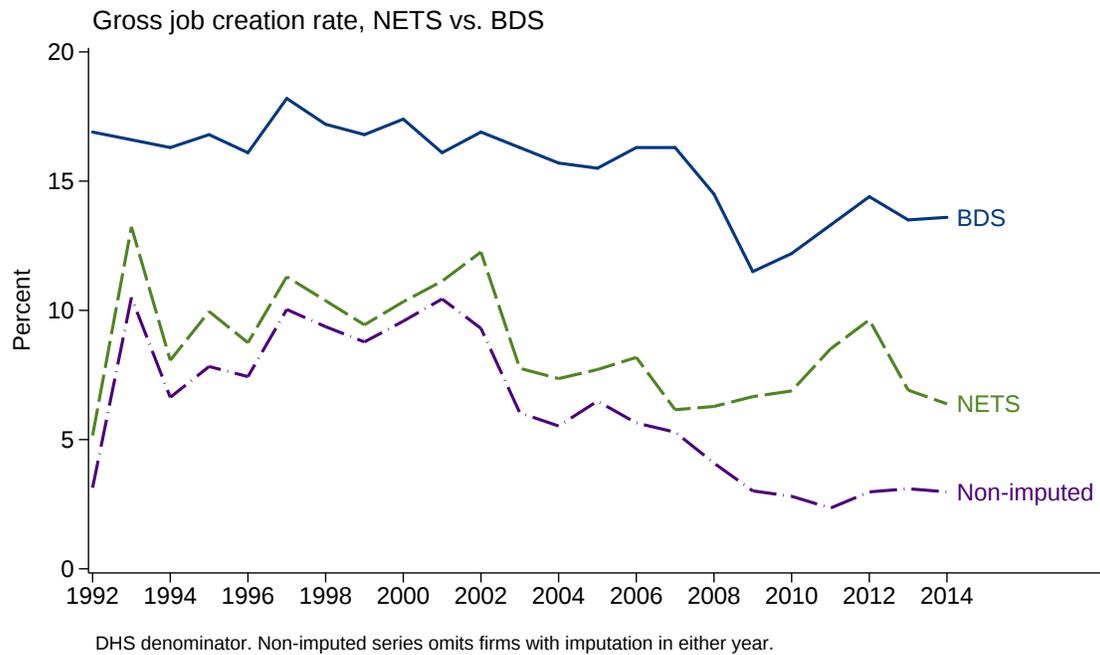


Figure 13

2000; entrants mechanically contribute to gross job creation, and the omission of imputed entrants should mechanically reduce the overall rate.

In any case, the patterns of gross job flows in NETS are substantially different from the BDS, both in terms of levels and in terms of time series behavior, and imputation alone does not account for these discrepancies.

4.2 Cell-based comparisons

We can drill down further by comparing detailed “cells” in the BDS and NETS. We focus on two disaggregations available in the publicly available BDS files: firm size by firm age by state, and firm size by firm age by industry. Comparing individual cells along these dimensions allows for a more complete picture of the two data sources. We focus on three critical measures of business dynamics: job creation rates, job destruction rates, and net

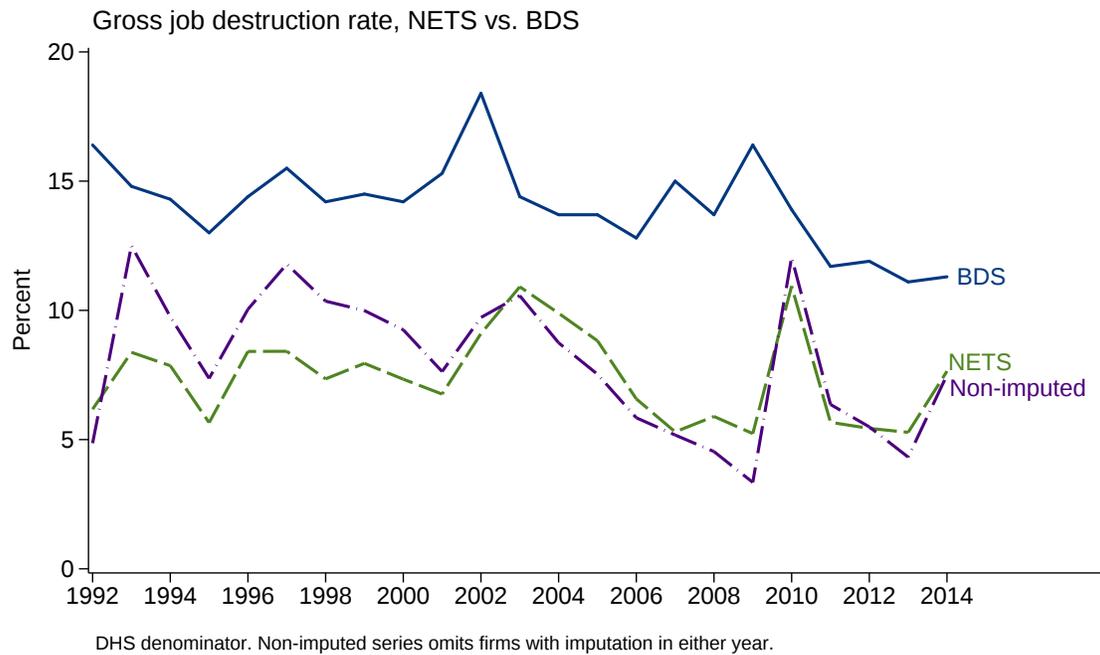


Figure 14

employment growth rates. We also study simple employment levels measured as the DHS denominator (i.e., two-year employment averages).

Firm size bins, in terms of employees (based on DHS denominator), are defined as follows: 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, and 10,000 or above; these are the narrowest size bins available in the BDS. Firm age bins are defined as follows: 0, 1, 2, 3, 4, 5, 6-10, and 11 or above. The BDS provides more age detail in the 11 and above category (11-15, 16-20, 21-25, and beyond), but given the shorter time series available in NETS, we combine the 11+ categories for better coverage. All states plus the District of Columbia are used, as are all SIC sectors available in the BDS: agricultural services, forestry, and fishing (SIC 7); mining (SIC 10); construction (SIC 15); manufacturing (SIC 20); transportation and public utilities (SIC 40); wholesale trade (SIC 50); retail trade (SIC 52); finance, insurance, and real estate (SIC 60); and services (SIC 70).

Therefore, the size by age by state disaggregation has potential for up to 4,896 cells, and the size by age by industry disaggregation has potential for up to 864 cells. When a cell in one data source is missing but that cell is not missing in the other data source, we populate each of job creation, job destruction, and DHS employment as zero in the missing source.¹⁴

Both location and industry are *establishment* characteristics, and multi-establishment firms can operate in multiple states and industries. When creating cell aggregates we implement BDS methodology, which follows.¹⁵ A single firm's activity can appear in any industry or state cell in which that firm has establishments, but only the establishments that belong to a given cell contribute data to that cell aggregate. However, *firm* characteristics are firm-wide and apply to all of a firm's establishments. That is, firm size and firm age information are the same for all establishments of a given firm. For example, consider a firm with two establishments, one in New York (first opened in 2000) and the other in New Jersey (first opened in 2002). Suppose we observe this firm in the year 2003 and find that the New York establishment has five employees and the New Jersey establishment has ten employees. Then in 2003, the firm has firm age of three and firm size of fifteen. The employment and job flows of the New York establishment will appear in the cell defined as firm size of 10-19 employees, firm age of 3, and New York state. The employment and job flows of the New Jersey establishment will appear in the cell defined as firm size of 10-19 employees, firm age of 3, and New Jersey state. That is, the New York and New Jersey establishments appear in the same firm size and age bins since they belong to the same firm, but they appear in different states. Industry is treated in the same manner.

Table 2 reports simple cell-based correlations between the BDS and NETS in terms of job creation, job destruction, net employment growth, and the DHS employment level; these correlations refer to actual levels (i.e., number of jobs created). For brevity we focus on two snapshot years, 2003 and 2014. We choose 2003 because this is the first year in which NETS is available given our firm age scheme, and we choose 2014 because it is the latest year in

¹⁴If a cell is missing in both sources, we do not generate an empty cell to populate in both sources.

¹⁵We confirmed the BDS methodology through correspondence with Census Bureau staff.

our data. The first two rows of Table 2 refer to the size-age-state cells. As can be seen from the first row, in 2003 the levels of job creation and job destruction were highly correlated between BDS and NETS, though net growth is less correlated. These correlations generally deteriorate by 2014. The correlation of employment levels, in the last column, remains extremely high throughout. The size-age-industry cell scheme shows similar results. The correlations for job creation, job destruction, and employment levels are reassuring, at least for 2003, and suggest that NETS broadly tracks the BDS along the studied dimensions. It may be that future revisions to 2014 data will improve the coverage.

| Cells | Year | Correlations | | | |
|-------------------|------|--------------|-----------------|-------|-------------|
| | | Job Creation | Job Destruction | Net | Denominator |
| Size-Age-State | 2003 | 0.891 | 0.937 | 0.651 | 0.984 |
| Size-Age-State | 2014 | 0.756 | 0.567 | 0.554 | 0.968 |
| Size-Age-Industry | 2003 | 0.893 | 0.910 | 0.685 | 0.971 |
| Size-Age-Industry | 2014 | 0.735 | 0.671 | 0.598 | 0.966 |

Source: NETS, BDS

Notes: Cross-cell, unweighted Pearson correlations of BDS and NETS levels. “Denominator” is the average of employment in years $t - 1$ and t

Table 2: Cell Correlations: Levels

However, these correlations hide an underlying divergence. Figure 15 plots job creation in BDS cells against NETS cells in 2003, and Figure 16 similarly plots job destruction. The job creation pattern illustrates how correlations can overstate the correspondence between the two data sources; a tight linear relationship is apparent, resulting in a high correlation, but the slope of the relationship is clearly steeper than the 45-degree line (dashed red line) that would indicate perfect correspondence. That is, NETS tends to understate job creation in 2003, relative to the BDS. The job destruction pattern has a less clear story but suggests that NETS may overstate job destruction relative to BDS, particularly in cells with higher job destruction levels. The high correlations shown on Table 2, therefore, partly reflect the fact that NETS and BDS show roughly similar magnitudes in an ordinal sense without always matching well in actual levels. These divergences in levels also help explain the lower cor-

relations for net job creation seen on Table 2, since net job creation is the difference between creation and destruction.

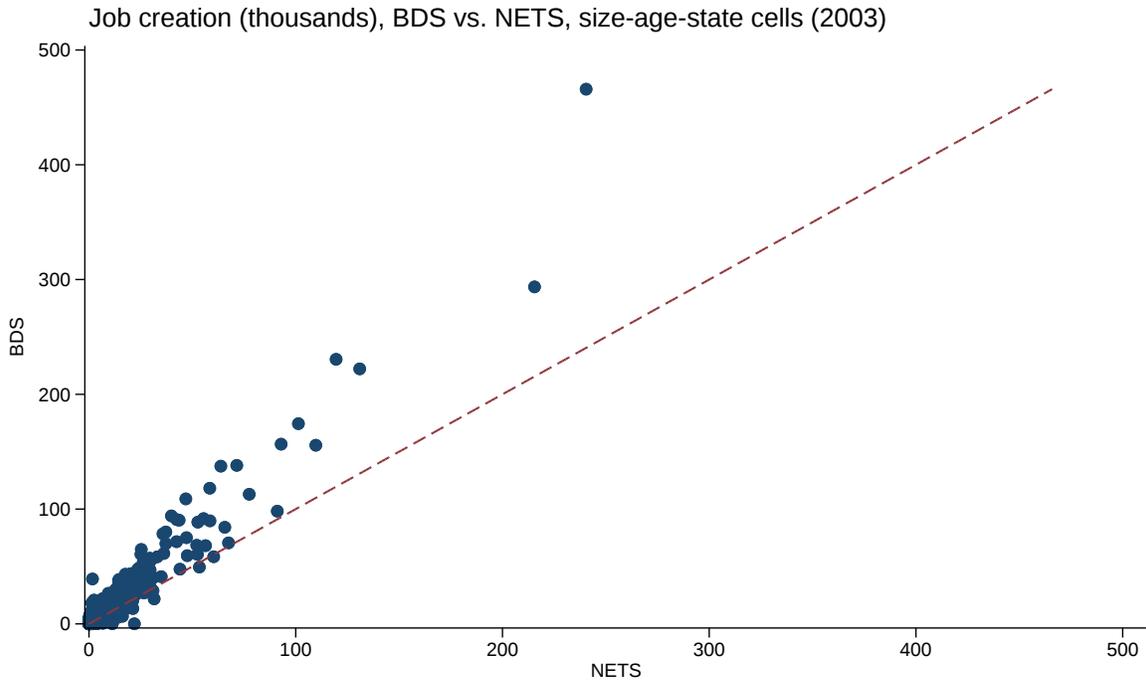


Figure 15

These results on levels of job flows may be of limited importance, however, since much research focuses on *rates* of job flows. We calculate cell-level job creation rates, job destruction rates, and net employment growth rates by dividing each level by overall DHS employment for the cell. We drop all firm age zero cells since, in both sources, these have job creation and employment growth rates of 200 percent and job destruction rates of 0 percent by construction. Table 3 reports these cell correlations, again for the two different disaggregation schemes and for the years 2003 and 2014. These correlations are generally quite low, again suggesting that the level correlations mostly reflect common employment scale effects, and that once things are normalized by employment the rates lack a close relationship across the data sources.

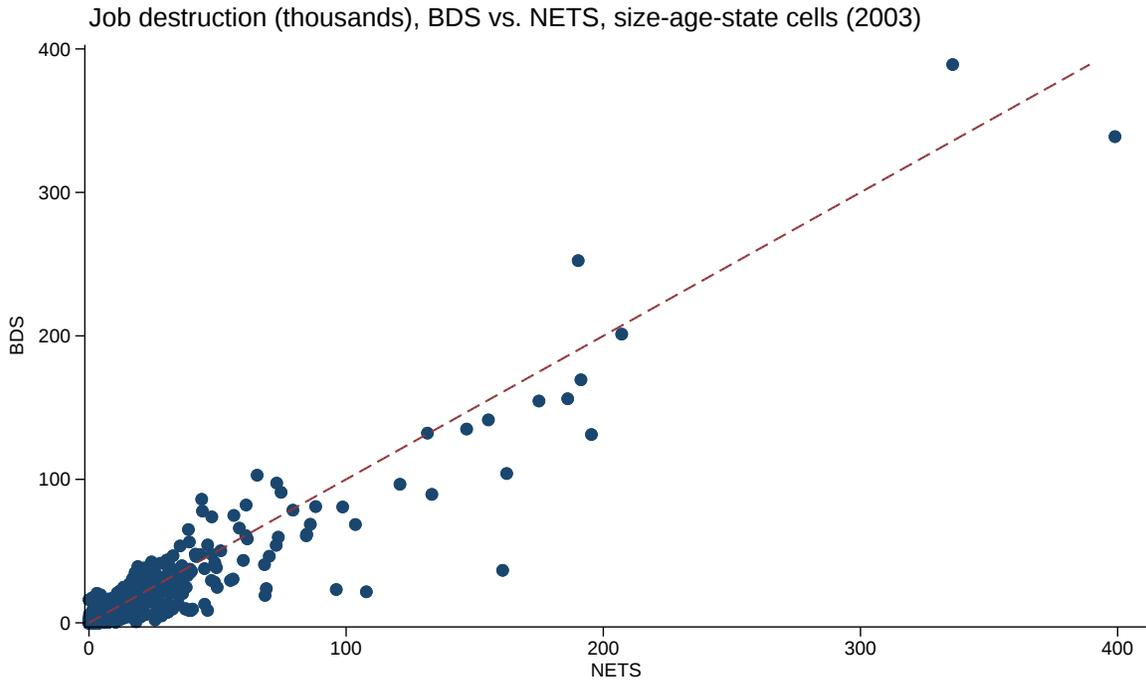


Figure 16

The cell-based comparisons generally support the concerns suggested by the aggregate analysis. NETS appears to have dampened rates of business dynamics compared with the BDS, and cell-level job flow rates are not strongly correlated between the two sources.

| Cells | Year | Correlations | | |
|-------------------|------|--------------|-----------------|-------|
| | | Job Creation | Job Destruction | Net |
| Size-Age-State | 2003 | 0.000 | 0.233 | 0.139 |
| Size-Age-State | 2014 | 0.078 | 0.158 | 0.095 |
| Size-Age-Industry | 2003 | -0.081 | 0.312 | 0.181 |
| Size-Age-Industry | 2014 | 0.134 | 0.070 | 0.045 |

Source: NETS, BDS

Notes: Cross-cell, unweighted Pearson correlations of BDS and NETS rates.

Table 3: Cell Correlations: Rates

4.3 Lifecycle dynamics

Many questions in firm dynamics focus on the firm lifecycle. Indeed, firm age and the behavior of young firms are at the center of many key firm dynamics questions because young firms play a disproportionate role in aggregate job growth (Haltiwanger et al. (2013)) and aggregate productivity growth (Alon et al. (2018); Decker et al. (2014) and references therein). As such, accurate measurement of entry and young firm behavior is critical for any dataset used to study firm dynamics. In this section, we proceed by using NETS to investigate critical firm lifecycle patterns that have been documented by LBD-based literature.

4.3.1 Average growth

A highly cited study in empirical firm dynamics is Haltiwanger et al. (2013). Using the LBD, the authors show that the widely held view that small businesses are the primary job creators—a view reinforced by NETS-based evidence (Neumark et al. (2011))—was clouded by data limitations. Rather, Haltiwanger et al. (2013) show that *young* firms are the key job creators; while small businesses do create jobs disproportionately, once the econometrician controls for firm age, the small firm advantage diminishes. The empirical regularity of small firms disproportionately creating jobs arises because young firms tend to be small. The evidence that young firms are critical for job creation has motivated a wide literature seeking to better understand young firms. Here we do not replicate the specific exercises of Haltiwanger et al. (2013) but instead illustrate the concept with a simpler exercise.

Figure 17, which relies on BDS data for 1992-2014, reports net firm employment growth rates by firm size bin, where size bins are set using *initial* firm employment and growth rates are averaged over the years in the sample.¹⁶ Exiting firms are included (which have growth of -200 percent). The light blue bars use all firms in the BDS and illustrate the view that was common prior to Haltiwanger et al. (2013): firm growth rates decline with firm size (at least among the smaller size bins) then hover near zero for larger sizes. The dark green

¹⁶Initial firm size means size in $t - 1$, where growth is calculated from $t - 1$ to t .

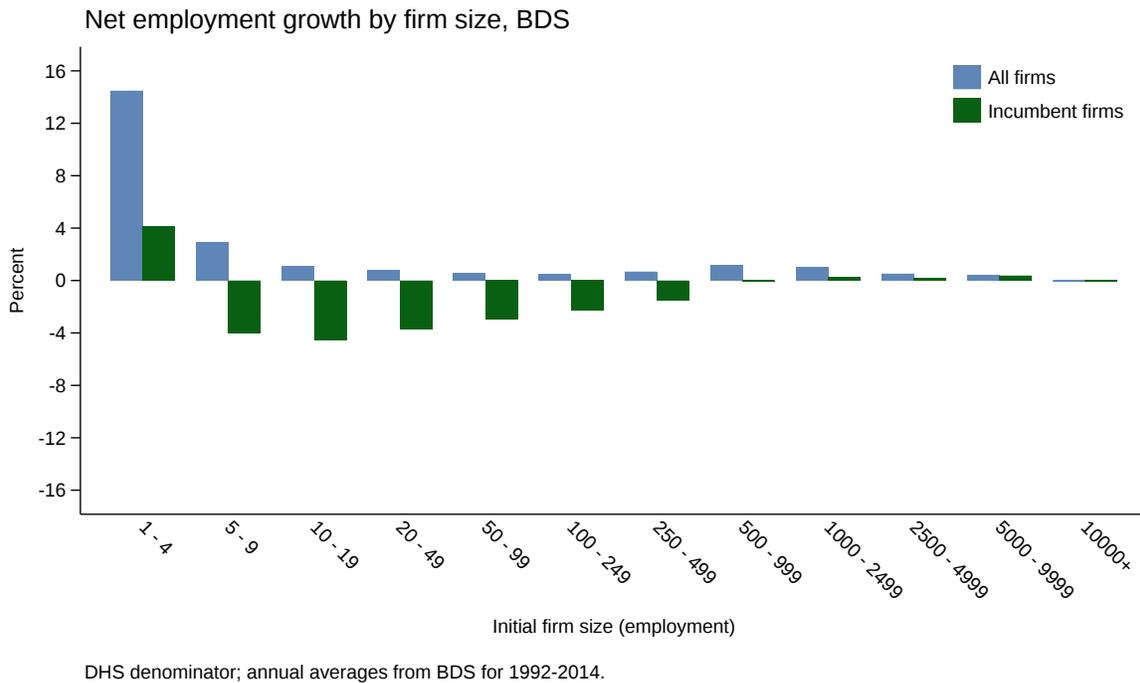


Figure 17

bars feature incumbent firms only (that is, new entrants are omitted). A starkly different picture emerges. The smallest size bin still has some growth advantage, though it is much diminished compared to the all-firm sample. Aside from the smallest class, all size classes below 500 employees actually see negative net growth on average. The figure illustrates the notion that the small-firm growth advantage is driven almost entirely by new entrants.

Figure 17 illustrates a critical stylized fact about the firm size and age distribution, so it is important that NETS data exhibit similar properties. Figure 18 reports the same exercise with NETS data. Rather reassuringly, NETS results are qualitatively (though not quantitatively) similar to those seen in the BDS. Figure 19 repeats the same exercise omitting firms in which at least 10 percent of employment is at establishments with longitudinally imputed employment figures. The result is starkly different and suggests that, oddly, the ability of NETS qualitatively to replicate Figure 17 is heavily dependent on imputed observations. In

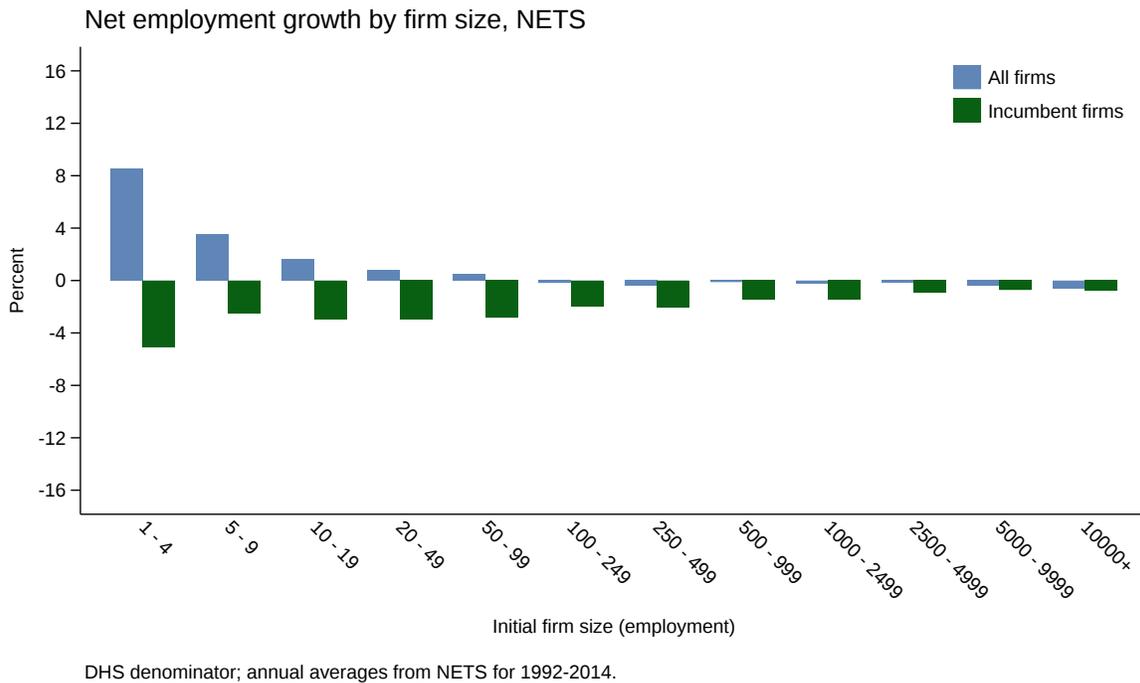


Figure 18

particular, it appears that much of entrants' contribution to the employment growth of the small firm bins reflects imputed employment data assigned to new firms. Indeed, as shown on Figure 10, close to 90 percent of new entrants (age 0) have imputed employment data. In 2014, of the new firms with imputed employment data, less than 1 percent reflect respondent imputation (i.e., "bottom of range"), while D&B and Walls & Associates estimates each comprise about half of imputations.

4.3.2 Skewness and churn

Haltiwanger et al. (2013) showed that young firms account for the high average growth rates of small firms. Decker et al. (2014) explore higher moments of the growth rate distribution over the lifecycle, documenting two key characteristics of young firm growth: skewness and churn. The growth outcomes of young firms are highly skewed, with a small number

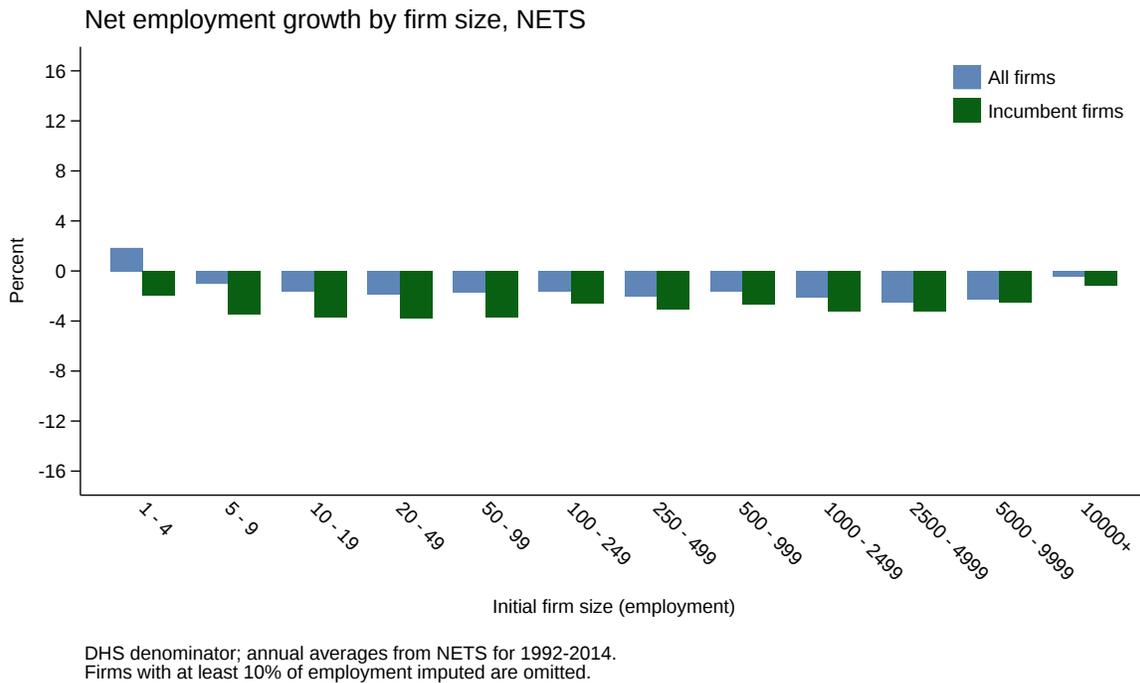
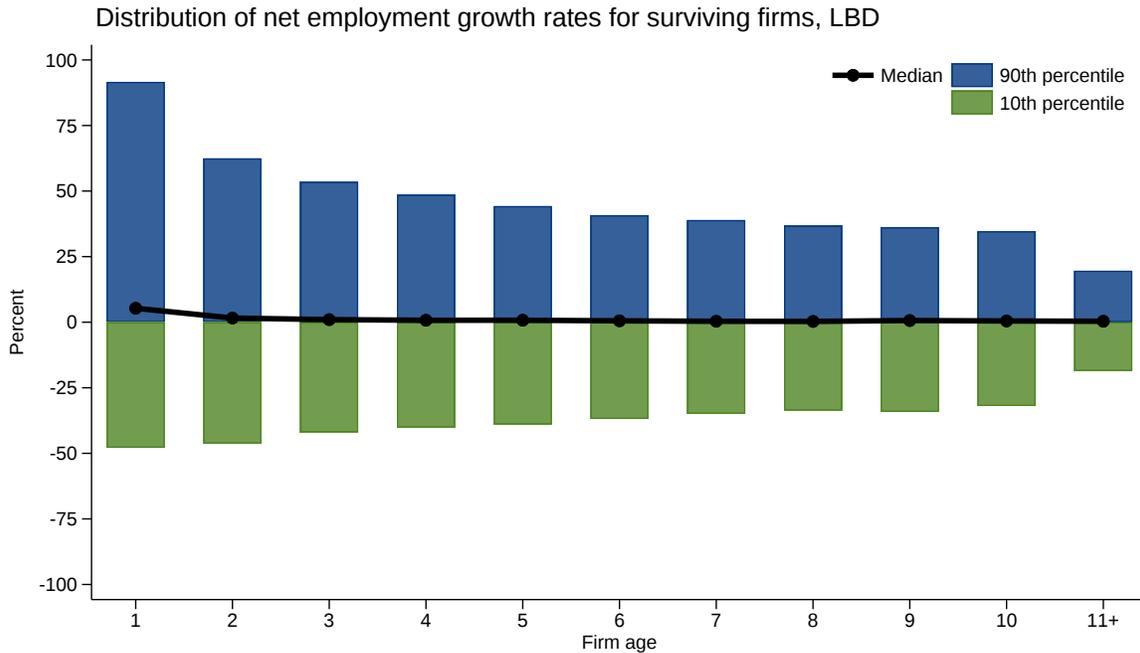


Figure 19

of extreme growth events. And young firms undergo considerable “churn”: the growth outcomes of young firms are highly dispersed, with a large amount of both very positive and very negative growth events, and young firms exhibit strong “up-or-out” dynamics as high incidence of failure among some young firms coexists with rapid growth of many survivors. These characteristics of young firms are not captured by average growth statistics but instead require study of the full distribution of growth outcomes, including outcomes of survivors and the prevalence of firm exit.

Figure 20, which is taken from Decker et al. (2014) exercises on LBD data, reports the growth rate distribution of surviving firms (i.e., those that do not exit) by age, averaged over the years 1992-2011.¹⁷ The solid line with dots is the median of the employment-weighted

¹⁷Decker et al. (2014) report 16 age bins, with the top bin including all firms age 16 and above. Given the shorter time series of NETS, to improve the comparison we report only 11 age bins. Since our project lacks access to the LBD microdata, in our reproduction of the Decker et al. (2014) figure we collapse age bins 11 and higher using simple averages of the reported percentiles.



Source: Decker et al. (2014). DHS denominator. Average across years 1992-2011.

Figure 20

growth rate distribution for the corresponding age bin; that is, for each age bin, half of all employment is at firms with growth rates at or below the black line. The top of the dark blue bars indicates the 90th percentile of the employment-weighted growth rate distribution, while the bottom of the light green bars indicates the 10th percentile of the employment-weighted growth rate distribution. Each statistic is calculated for every year in the sample, then averaged across years.¹⁸

A few key patterns are evident from Figure 20 (see Decker et al. (2014) for more discussion). First, median employment growth is only positive among young firms; the typical mature firm has zero employment growth, consistent with the age profiles described above. Second, growth outcomes are highly dispersed among young firms, with dispersion declin-

¹⁸The population of firms included in Figure 20 differs from the population included in Figure 17 in that 17 potentially includes all firms or all incumbents (including firms that exit, with a growth rate of -200 percent), but 20 includes only surviving firms. That is, in Figure 20, the bars corresponding with firm age 1 include firms that survived to reach age 1, omitting those that exited between ages 0 and 1.

ing as firm cohorts age. This fact illustrates the high pace of churn among young firms, with many outcomes of both extreme growth and extreme decline. Third, the growth rate distribution of young firms is characterized by skewness, shown as the distance from the 90th percentile to the median compared with the distance from the 10th percentile to the median; this skewness illustrates that the substantial job growth contribution of young firms includes not widespread growth but in fact a few firms with extremely high growth. Skewness disappears entirely by age five, a reason that much of the literature studies young firms with an age cutoff around five. High growth is a characteristic of (some) young firms.

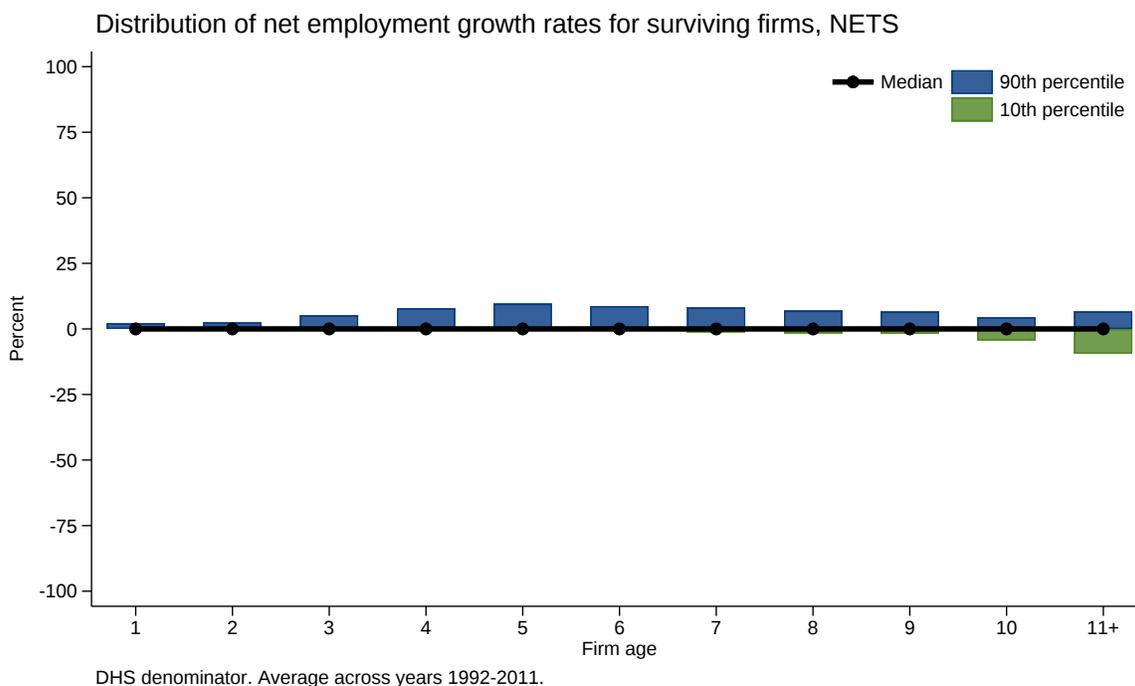


Figure 21

As with the data on firm growth by size and age, the patterns of dispersion and skewness over the (surviving) firm lifecycle evident in Figure 20 are critical stylized facts about the behavior of young firms and the sources of aggregate employment growth. We evaluate the ability of NETS to exhibit these patterns on Figure 21, which mimics Figure 20.

The difference between the figures is very concerning: While LBD data in Figure 20 exhibit significant growth rate dispersion among firms of all ages (and particularly young firms), very little dispersion is evident in the NETS data shown on Figure 21. Since these are employment-weighted distributions, the latter figure indicates that 90 percent of surviving-firm employment is at firms with a growth rate around zero percent or higher for almost all age groups, while in the LBD we observe very young firms with growth approaching negative 50 percent and even many mature firms with growth around negative 25 percent. And while negative growth is nearly absent from NETS data, positive growth is almost as rare. For example, in the LBD we observe *young* firms that account for around 10 percent of employment growing at a rate of 50 percent or more, but no firm age group in NETS reports a 90th percentile growth rate beyond 20 percent. The median growth rate in NETS, shown by the black line with dots, is close to zero for firms of all ages, in contrast to the positive growth rates seen by young firms in the LBD. NETS data therefore miss virtually the entire distribution of firm growth outcomes, whatever their performance tracking average growth patterns. This is a significant limitation of NETS generally and is particularly problematic for the study of young firms, which (as shown on Figure 20) are especially characterized by wide dispersion and high skewness of firm growth rates. In unreported results we find that omitting imputed observations from NETS does not materially alter Figure 21.

Decker et al. (2014) also document the “up-or-out” nature of the young firm environment by contrasting exit and survival. Figure 22, which uses BDS data to replicate Decker et al. (2014), reports the experiences of firm cohorts as follows. The light blue bars report jobs destroyed (over one year) by firms that exit just before reaching a given age; that is, the blue bar for age 1 reflects exits of firms between age 0 and age 1, the blue bar for age 2 reflects exits of firms between age 1 and age 2, and so on. The dark green bars report net job creation (over one year) among firms that survive to a given age; that is, the green bar for age 1 reports jobs created by firms continuing from age 0 to 1, the green bar for age 2 reports jobs created by firms continuing from age 1 to 2, and so on. All figures are scaled by the

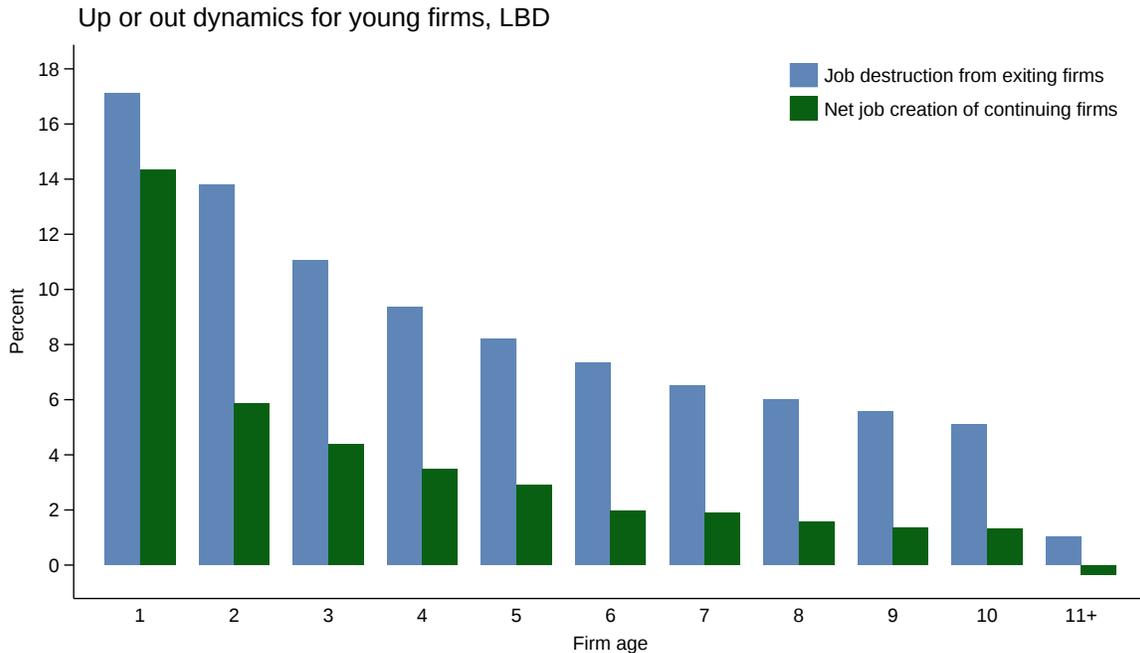


Figure 22

DHS employment denominator for the entire cohort, and rates are calculated by year then averaged over all years 1992-2011.¹⁹

Figure 22 illustrates three key points about the firm lifecycle. First, both job destruction from exits and job creation from entrants are high among young firms and decline monotonically with firm age, consistent with evidence above that young firm outcomes are volatile and highly dispersed. Second, an “up-or-out” pattern is evident in the sense that, while many jobs are destroyed by exiting firms, surviving firms have high average growth rates. Third, job creation from survivors is more than offset by job destruction from exiting firms for all age groups. Note that, by construction, age zero firms (not shown on Figure 22) only create (i.e., do not destroy) jobs, so creation far offsets destruction upon entry. A reason-

¹⁹As with the previous set of figures, we collapse all age categories above 10 into a single “11+” category, which is simple in this exercise since we rely on BDS data. We do this for comparability with NETS data but make a note of it because it differs slightly from the setup in Decker et al. (2014).

able characterization of young firm dynamics, then, is that each new cohort creates a large number of jobs upon entry, but firms immediately begin failing, destroying many jobs as firms age but with continued growth among surviving firms that partially offsets the job destruction.²⁰

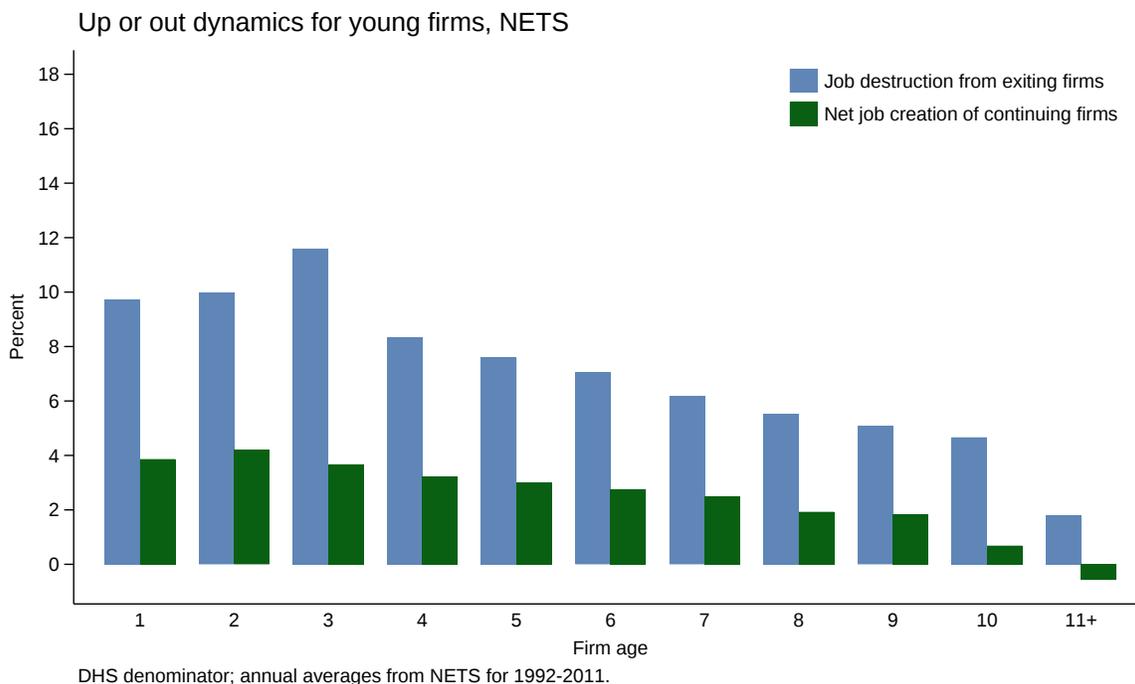


Figure 23

Figure 23 mimics Figure 22 using NETS data to assess the presence of “up-or-out” dynamics in NETS. The performance of NETS in this exercise is not as weak as in the previous skewness and dispersion exercise: the lifecycle pattern of exit-driven destruction and creation of continuers is not quite monotonic but is qualitatively similar to BDS data in that destruction from exit outpaces creation among continuers for all age groups. Moreover, among age groups above 5 the magnitudes of job destruction and creation appear reason-

²⁰Decker et al. (2014) note that the post-entry job destruction of exiting firms is still not enough to completely offset the jobs created upon entry: five years after entry, the employment of the typical cohort is still equal to 80 percent of the cohorts entry employment, such that new cohorts of firms make permanent contributions to aggregate employment despite high failure rates in early years.

ably accurate. However, the magnitudes illustrated by the figure indicate particularly poor measurement of *young* firm dynamics. The differences between young and mature firms, in terms of both job destruction and creation, are much smaller in NETS than in the BDS, and the monotonicity-by-age is wrong for the youngest age groups.

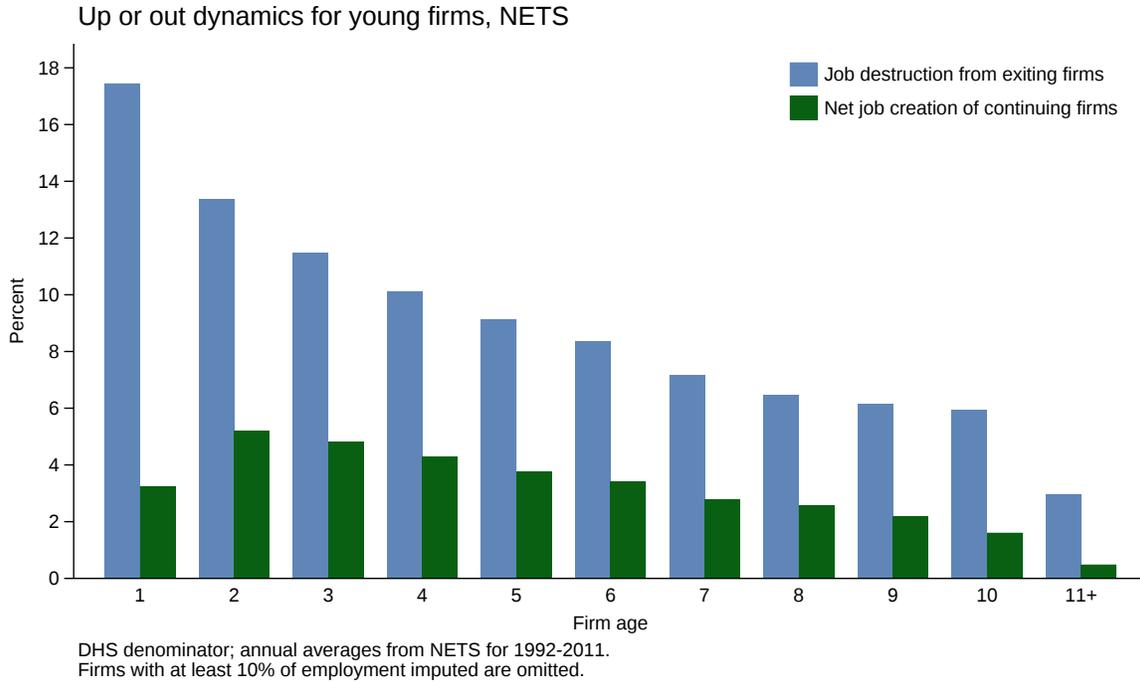


Figure 24

Figure 24 documents the same exercise in NETS omitting firms that have longitudinally imputed data comprising at least 10 percent of their employment; interestingly, NETS' inability to track the pattern of exit-driven job destruction among young firms shown in Figure 23 appears to be due to imputed observations; job destruction rates across the lifecycle look reasonably accurate among non-imputed firms. However, job creation rates among young firms appear little better among the non-imputed observations than among NETS firms generally. Again we observe that NETS is particularly limited in its measurement of young firm dynamics, and young firm dynamics are a critical component of the overall firm dynamics

literature.

5 Discussion

The foregoing comparisons reveal serious discrepancies between NETS and official administrative data. NETS displays markedly different patterns of young firm activity in terms of both aggregate activity shares and the micro behavior of young firms. NETS businesses generally exhibit patterns of business dynamics that are far less volatile than those seen in official sources. A key driver of these discrepancies is the high rate of imputation in NETS, particularly among young firms, most of whom lack fresh data observations. But restricting the sample to omit imputed data is no panacea, as imputation is extremely prevalent among young businesses and restricting the sample to non-imputed observations creates composition effects that do not usually mitigate the discrepancies. These limitations of NETS are serious and oblige researchers using NETS to use caution. Topics including reallocation, entrepreneurship, firm growth and exit, and inaction are highly vulnerable to the limitations of NETS.

One potential response to the NETS/LBD discrepancies we document here is that government data sources are also imperfect and may not have a claim on being the benchmark against which private data sources should be judged. We readily acknowledge that official sources have many limitations. For example, users of the LBD encounter problems with firm identifier longitudinal linkages, staleness of industry codes and firm organization details between census years, and lack of easily integrated coverage of the nonemployer universe. Indeed, methods of defining firm age and organic employment growth, now widely used in empirical literature but pioneered by [Haltiwanger et al. \(2013\)](#) and [Davis et al. \(2007\)](#), are designed to minimize the errors introduced by these limitations. [Barnatchez et al. \(2017\)](#) discuss limitations of official data more broadly, including the Census Bureau's County Business Patterns (which uses the same source data as the LBD), and show discrepancies between Census and Bureau of Labor Statistics (BLS) sources even when restricted to com-

mon industry scope (though these discrepancies are smaller than those between NETS and either official source).

However, there are at least two reasons to treat the official sources as authoritative. First, official data collection efforts are characterized by intense focus on consistency and measurement best practices. For example, the employment data on which the LBD is based always report employment as of March 12 of a given year; in contrast, the Dun & Bradstreet employment figures could be recorded at any point during the year, rendering them vulnerable to seasonal fluctuations. Other LBD variables are continually updated with information from administrative and survey sources, such as the annual Report of Organization survey, and Census surveys are conducted scientifically.²¹ More broadly, the U.S. statistical agencies employ large staffs of experts tasked with ensuring data quality as well as active researchers exploring and performing research and development on data products.²² These efforts are supplemented by robust exchanges between statistical agency staff and outside experts, such as those facilitated by the Federal Economic Statistics Advisory Committee (FESAC). For the purposes of D&B, scientific best practice is likely to be both excessively costly and unnecessary; for example, an estimated or imputed employment observation is often good enough for the needs of D&B clients while being much less useful for researchers of business dynamics.

Second, official sources are based in part on administrative data that are accurate by construction. The LBD source data are ultimately tax records, so the LBD represents the universe of in-scope employer businesses that are known to U.S. tax authorities—a clear and reasonable definition of business activity that contrasts with D&B’s looser goal of covering a less-defined employer and nonemployer business universe with large potential for undercoverage relative to its goal (as appears to have been the case prior to the likely scope

²¹See <https://www.census.gov/programs-surveys/cos/about.html> for details about the Report of Organization survey, also known as the Company Organization Survey.

²²For example, the Census Bureau’s Center for Economic Studies employs many social scientists who actively evaluate the research uses and limitations of the LBD and other Census data products. Other Census offices have similar features, and additional quality control is performed by authorized outside researchers using the Federal Statistical Research Data Centers.

improvement of the 2000s documented by [Barnatchez et al. \(2017\)](#)). To the extent that NETS differs from the combined Census employer and nonemployer universe (i.e., CBP and NES) in terms of establishment coverage, it must be that NETS is either including businesses defined in some other way than taxable entities or omitting taxable businesses. Likewise, annual employment snapshots in the LBD represent data that are routinely used for administrative purposes by the IRS and the Social Security Administration, limiting the scope for inaccuracy and imputation. While some LBD establishments only receive industry code and company organization updates after the semidecennial Economic Censuses, employment data come from administrative sources. NETS, by contrast, exhibits high rates of imputation of employment data, which is particularly problematic for the study of business dynamics.

Weaknesses and limitations of the official sources notwithstanding, then, the LBD and corresponding BDS are, in our view, best treated as more authoritative than NETS. The discrepancies between the sources are therefore cause for concern about the usefulness of NETS for business dynamics research.

For the study of firm dynamics, we conclude that the most promising use of NETS is not for broad studies of entrepreneurship or firm dynamics but instead for more detailed, narrower investigations of specific case studies in which the microdata can be evaluated against outside sources prior to analysis. [Echeverri-Carroll and Feldman \(2017\)](#), discussed above, is one such example, though it is limited to two cities (albeit ones with important recent entrepreneurship patterns, Austin, TX and the North Carolina “Research Triangle”). Those authors carefully describe ways to make the NETS data most reliable, including appropriate sample restrictions and relaxation of firm entry timing to windows broader than one year. Researchers with questions for which such restrictions and timing conventions are appropriate may find similar success, though we argue that validation against independent data sources will continue to be necessary before proceeding with case studies of other cities.

6 Conclusion

[Barnatchez et al. \(2017\)](#) argue that NETS can be useful for the study of static business distributions, provided that researchers exercise appropriate caution and pay careful attention to problems with coverage of certain kinds of establishments. In this paper, we obtain less optimistic results when using NETS to study business dynamics in a microdata approach. We show that one particular concept of high interest to firm dynamics researchers—the lifecycle dynamics of young firms—appears poorly measured in NETS data, as are broader concepts like gross and net job flows among firms generally. This is a considerable limitation given the importance of these concepts for studies of business and labor market dynamics. Popular topics including entrepreneurship, job reallocation, high-growth firms, and firm failure may be difficult to study with high confidence in NETS, a finding consistent with the more limited investigation performed by [Haltiwanger et al. \(2013\)](#). Through painstaking firm-level comparison work researchers may find specific settings in which NETS can be reliable, such as [Echeverri-Carroll and Feldman \(2017\)](#), but in general we urge caution.

While we view our results as compelling, there are many aspects of NETS that we do not investigate. NETS includes a wealth of information on variables other than employment, industry, and location, such as sales, credit information, and legal form of organization. We leave investigation of these and other variables for future research.

References

Alon, Titan, David Berger, Robert Dent, and Benjamin Pugsley, “Older and Slower: The Startup Deficit’s Lasting Effects on Aggregate Productivity Growth,” *Journal of Monetary Economics*, 2018, 93, 68–85.

Barnatchez, Keith, Leland D. Crane, and Ryan A. Decker, “An Assessment of the National Establishment Time Series (NETS) Database,” Finance and Economics Discussion Series 2017-110, Federal Reserve Board 2017.

Chava, Sudheer, Alex Oettl, Manpreet Singh, and Linghang Zeng, “The Dark Side of Technological Progress? Impact of E-Commerce on Employees at Brick-and-Mortar Retailers,” 01 2018.

Cho, Clare, Patrick Mclaughlin, Eliana Zeballos, Jessica Kent, and Chris Decken, “Capturing the Complete Food Environment With Commercial Data: A Comparison of TDLinx, ReCount, and NETS Databases,” Technical Bulletin 1953, USDA Economic Research Service March 2019.

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

—, **John Haltiwanger, Ron S. Jarmin, and Javier Miranda**, in Daron Acemoglu, Kenneth Rogoff, and Michael W, eds., *NBER Macroeconomics Annual 2006, Volume 21*, The MIT, 2007, chapter Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms, pp. 107–180.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda, “Changing Business Dynamism and Productivity: Shocks vs. Responsiveness,” *NBER working paper no. 24236*, 2018.

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

DeSalvo, Bethany, Frank Limehouse, and Shawn D. Klimek, “Documenting the Business Register and Related Economic Business Data,” Technical Report, Center for Economic Studies working paper CES-WP-16-17 2016.

Echeverri-Carroll, Elsie and Maryann Feldman, “Chasing Entrepreneurial Firms,” Technical Report 2017.

- Faccio, Mara and Hung-Chia Hsu**, “Politically Connected Private Equity and Employment,” *The Journal of Finance*, 2017, 72 (2), 539–574.
- Guzman, Jorge and Scott Stern**, “The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 15 US States, 1988-2014,” Working Paper 22095, National Bureau of Economic Research March 2016.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda**, “Who Creates Jobs? Small versus Large versus Young,” *The Review of Economics and Statistics*, 2013, 95 (2), 347–361.
- , –, **Robert Kulick, and Javier Miranda**, “High growth young firms: Contribution to job growth, revenue growth, and productivity,” in “Measuring entrepreneurial businesses: Current knowledge and challenges,” University of Chicago Press, 2017.
- Heider, Florian and Alexander Ljungqvist**, “As certain as debt and taxes: Estimating the tax sensitivity of leverage from state tax changes,” *Journal of Financial Economics*, 2015, 118 (3), 684 – 712. NBER Symposium on New perspectives on corporate capital structures.
- Jarmin, Ron S. and Javier Miranda**, “The Longitudinal Business Database,” Technical Report, Center for Economic Studies 2002.
- Kimbrough, Gray**, “Uncluttered Stata Graphs,” 2018. Github repository. At <https://github.com/graykimbrough/uncluttered-stata-graphs>. Accessed April 29, 2019.
- 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,” *The Review of Economics and Statistics*, 2011, 93 (1), 16–29.
- , **Junfu Zhang, and Brandon Wall**, “Employment Dynamics and Business Relocation: New Evidence from the National Establishment Time Series,” Working Paper 11647, National Bureau of Economic Research October 2005.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter, “Diverging Trends in National and Local Concentration,” Working Paper 25066, National Bureau of Economic Research September 2018.

Understanding Data in the NETS Database

Understanding Data in the NETS Database, Technical Report, Walls & Associates NETS marketing materials 2009.

Walls, Donald, “*Understanding Data in the NETS Database,*” *NETS Documentation*, 2008.