Do startups pay less?

M. Diane Burton
Cornell University
ILR School
170 Ives Hall
Ithaca, NY 14853
burton@cornell.edu

Michael S. Dahl
Aarhus University
Department of Management
Bartholins Allé 10
DK-8000 Aarhus
msd@mgmt.au.dk

Olav Sorenson
Yale University
165 Whitney Ave.
P.O. Box 208200
New Haven, CT 06520
olav.sorenson@yale.edu

ACKNOWLEDGEMENTS

We thank Yale University for generous financial support. We received helpful comments from seminar participants at CBS, Cornell, MIT, Oxford, and the University of Toronto. We are particularly indebted to Paige Ouimet, our discussant at the NBER 2015 Summer Institute. We also benefitted from discussions with Lisa Kahn, Toke Reichstein, Scott Stern, and Niels Westergård-Nielsen.

Do startups pay less?

June 19, 2017

Abstract: We analyzed Danish registry data from 1991 to 2006 to determine how firm age and size influence wages. Unadjusted statistics suggest that smaller firms paid less than larger ones and that firm age had little or no bearing on wages. After adjusting for differences in the characteristics of employees hired by these firms, however, we observed both firm age and firm size effects. We found that larger firms paid more than smaller firms for observationally-equivalent individuals but, contrary to conventional wisdom, that younger firms paid *more* than older firms. The size effect, however, dominates the age effect. Thus, while the typical startup – being both young and small – paid less than a more established employer, the largest ones paid a wage premium.

Introduction

Policymakers around the globe have become interested in promoting entrepreneurship as a means of creating jobs and stimulating economic growth. But there has been relatively little discussion about the quality of the jobs that startups create, particularly in terms of the salaries that they pay and the benefits that they offer. If the process of creative destruction involves replacing higher paying jobs at incumbent firms with lower paying ones at startups, then a simple examination of the number of jobs created, even net of jobs lost, may overstate the value of entrepreneurial activity to the economy and society.

Although some research has examined the wages of startups, the findings have been mixed and we cannot yet say with certainty whether startups pay more or less than established firms. Some of the inconsistencies across studies arise from different methodological choices. But some of the uncertainty arises from flawed assumptions about the ways that firm age and firm size relate to wages and from a failure to account for employee characteristics that influence both pay and the probability of joining a startup.

We study the extent to which pay varies as a function of firm age and firm size using comprehensive registry data on the population of Danish workers, from 1991 to 2006. We expand on the existing literature in at least four important ways: (1) by estimating the age and size effects as categories, we allow firm age and firm size to have flexible relationships with pay; (2) by focusing on new hires to the firm, we eliminate the extent to which variation in average job tenure might account for firm age and size effects; (3) by matching employees of smaller and younger firms with observationally-equivalent counterparts at large, established firms, we account for human capital differences that might explain earnings differentials; and (4) by including fixed effects for fine-grained industry categories, we adjust for the fact that firm age and size might vary with industry competitiveness and employment growth.

Across our analyses, younger firms paid *more* than older ones, and smaller firms paid less than larger ones. We document how methodological choices have substantial implications for the estimated magnitudes of annual earnings differentials across firms. First, our approach to treating firm age and size as categories (bins) yielded relationships as much as twice that of the standard approach of using logged firm age and size as covariates. Second, although focusing on new hires had little influence on the relationship between firm size and pay, the differences in pay with respect to firm age more than doubled once the confounding effects of employee tenure had been eliminated. Third, taking into account differences in the human capital characteristics of employees accounted for roughly one-third of the firm size wage effect and almost two-thirds of the firm age wage effect. Industry adjustments and other methodological choices, by contrast, made little difference to our wage estimates.

Through a variety of methodological adjustments, and consistent with prior research, we found a firm size wage penalty. The smallest firms paid less than the largest ones, by a factor of 10%-15%. Somewhat unexpectedly, however, younger firms paid more than older firms (though by less than 5%). In fact, large young firms appeared to pay a wage premium over established employers. In most cases, however, the size effect dominates the age effect, meaning that the typical startup – being both small and young – pays less than the average incumbent employer.

Firm age, firm size, and wages

Young firms account for an outsized share of all net job creation in the United States (Haltiwanger et al. 2013), in Denmark (Ibsen and Westergård-Nielsen 2011; Malchow-Moller et al. 2011) and in most other countries (Ayyagari et al. 2014; de Wit and de Kok 2014; Lawless 2014; Anyadike-Danes et al. 2015). Yet, despite enthusiasm for entrepreneurship on the part of policymakers and the evidence that startups account for the majority of net job creation, there are reasons to be pessimistic about entrepreneurship as an engine for creating good jobs and generating broad-based economic benefits. For starters, firms typically begin small. But large firms enjoy economies of scale and scope and therefore they can increase employee productivity. Indeed, an extensive empirical literature has examined the relationship between firm size and average compensation, finding that larger firms pay more and offer better benefits than smaller ones (for a review, see Oi and Idson 1999).

Startups may pay less than older firms even independent of these size effects. Fledgling firms have not had the opportunity to improve their operations through learning-by-doing (Arrow 1962), or by investing in equipment (Thompson 2001). Nor have they had time to build social capital (Sorenson and Rogan 2014). To the extent that these factors represent complements in production (Griliches 1969), startups should operate at lower levels of productivity than more established firms and consequently pay their employees less. Only a handful of studies to date, however, have examined the relationship between firm age and wages, net of firm size effects. Troske (1998), for example, reported that the youngest manufacturing plants in the United States paid nearly 20% less than the oldest ones in the late-1980s, even after adjusting for differences due to firm size. Similarly, Brixy et al. (2007), examining evidence from Germany, found that newly-founded firms paid roughly 8% lower wages on average than their older counterparts in the late-1990s.

Although studies typically find that older firms pay more, on average, than younger ones (Audretsch et al. 2001; Brixy et al. 2007), most of the studies informing our understanding of firm age and wages have only had information on the average wages paid by firms and therefore have been unable to adjust for differences in the characteristics of the employees of startups relative to other firms. However, startups may prove less appealing to employees because of their lower levels of capital investment and uncertain prospects. Indeed, recent research has documented that smaller and younger firms hire younger, less educated, and less experienced individuals (Nystrom and Elvung 2014; Ouimet and Zarutskie 2014). These employees have less human capital and would earn less at any employer, young or old, large or small. The observed firm age effects may therefore reflect differences in who firms hire rather than in what they pay.

Another complication to understanding firm age effects on wages stems from the fact that employees gain firm-specific experience over time, increasing their value to their employers. Older firms have employees with longer tenure. Accounting for this fact, moreover, can prove vexing. The typical approach has been to assume that wages adjust linearly with firm tenure (e.g., Brown and Medoff 2003; Heyman 2007). But, given that younger firms and older firms do not even overlap over most of the range of the tenure variable, that assumption could prove problematic. A two-year-old firm, for example, cannot have any employees with more than two years of experience at the firm, but an established firm might have few employees with less than two years of tenure at the firm. Comparing average pay across firms of different ages conflates the effects of firm age with those of employee tenure.

These methodological challenges imply that the apparent effects of firm age and size on wages could stem more from employee selection and retention than from differences across firms in their productivity or ability to pay (e.g., Abowd et al. 1999). At least four studies have addressed these issues. But those studies have used different approaches and have found conflicting results, particularly with respect to the effects of firm age, a fact that may reflect their differing methodological choices. Brown and Medoff (2003), in the first study of this type, included the characteristics of employees, firm age, and logged firm size in estimates of the wages of 1,410 American workers. In initial specifications that included only firm characteristics, they found average pay rising with firm age. But this relationship flipped after they accounted for employee characteristics.

Heyman (2007) and Nystrom and Elvung (2014), however, using registry data from Sweden, both found that older firms paid somewhat higher wages. Heyman (2007) took an approach similar to Brown and Medoff (2003), including a variety of individual-level characteristics in a wage equation for roughly 500,000 employee-years observed in three cross sections (1987, 1991, and 1995). Although he could only adjust for firm-specific experience in the latter two cross-sections (by including a linear term for tenure), in models adjusting for tenure, he found non-negative relationships between firm age and pay. In the 1995 cross section, moreover, older firms consistently paid more than younger ones.

Nystrom and Elvung (2014), by comparison, adopted a propensity score matching approach to adjust for an even larger set of employee characteristics among a sample of more than 150,000 entrants to the labor market. By definition, these individuals had no experience at the firm, the research design therefore eliminated firm tenure as a potential confound. But, in contrast to other studies that have accounted for employee characteristics, their adjustments for employee sorting had little effect on the estimated wage penalty of joining a young firm, perhaps because the models that they used to create propensity scores explained less than 2% of the variation in who joined a startup.

Schmieder (2013), meanwhile, found the opposite relationship between firm age and wages using registry data from Germany, with young firms offering as much as a 6% *premium* over older employers. To address the tenure problem, Schmieder focused only on job changers (movers). To adjust for heterogeneity in employees and employers, he used individual-level and firm-level fixed effects. His estimates therefore focus on the changes in pay associated with moving across firms as those firms mature and grow in size. Although his empirical approach addresses both the tenure issue and the sorting of employees into firms, the German data that he uses primarily capture older and larger employers. To the extent that firm age and size effects have decreasing marginal returns, his results may therefore underestimate the overall importance of these firm characteristics. He also does not allow firm age and size to interact in their determination of wages.

As this brief review reveals, although scholars have begun to address the question of how wages vary with firm age and size, many questions remain: Do firm age and size effects follow the functional forms typically used? Do they interact in their determination of wages? To what extent do the apparent age effects confound firm age with employee tenure? To what extent might unobserved heterogeneity either on the side of the employee or on the side of the firm account for differences after adjusting for the observed characteristics of employees and employers? Many studies, moreover, have treated firms and establishments interchangeably without careful attention to distinguishing newly created firms from either foreign subsidiaries or expansions of an existing enterprise. We tackle these questions below.

Empirical Strategy

To advance our understanding of the relationships between firm age, firm size, and wages, and to explore whether startups pay more or less than established firms, we examined Danish registry data, the Integrated Database for Labor Market Research, commonly referred to by its Danish acronym, IDA (for a useful English-language overview of this data, see Timmermans (2010)). Although these data cover every employee in the country, we began by restricting the sample to full-time employees between the ages of 18 and 60 to focus on adults and on those who had not yet begun to shift their employment choices in anticipation of retirement.

Although the IDA database begins in 1980, we considered only the post-1991 period. In the late-1970s and 1980s, a series of regulatory reforms dismantled most of the centralized wage-setting system, allowing firms much more flexibility in their compensation practices (Madsen et al. 2001). Denmark now has some of the least restrictive labor market policies in Europe (Bingley and Westergård-Nielsen 2003; Sørensen and Sorenson 2007), usefully allowing comparison to larger economies such as Canada, the United Kingdom, and the United States. Although Denmark has made it easy to hire and fire employees, and has a flexible wage-setting regime, the country nevertheless retains a strong social support net, a combination sometimes referred to as *flexicurity* (Madsen 2004).

Unlike the United States, for example, benefits, such as health insurance and retirement plans, come from the central state rather than from employers. For our purposes, it has the advantage of ensuring that most of the differences between employers in the quality of jobs stems from what they pay, rather than from a combination of wages and benefits.

Two additional features of the Danish context also help to simplify our analysis. First, the Danish tax system does not encourage the use of equity as a form of compensation. Denmark taxes equity awards as income rather than as capital gains. As a consequence, most companies use bonuses rather than stock or stock options as a means of paying for performance (Eriksson 2001). Eriksson (2001), moreover, finds that small firms in Denmark use equity compensation – as well as all other forms of variable compensation – less than large ones. Because bonuses appear as income, our analyses should include all forms of compensation for the vast majority of employees.

Second, Denmark has strong norms against long work hours. In 2015, for instance, only 2% of Danish full-time workers reported working more than 50 hours per week compared to an OECD average of 13% (OECD 2015). Any observed annual pay differences therefore likely do not reflect differences in the number of hours worked.

Average wages by age and size

We first divided and classified each employer into one of four size categories: 1-10 full-time employees, 11-49 full-time employees, 50-249 full-time employees, and more than 250 full-time employees. We also divided and classified each employer into one of four age categories: 1-2 years, 3-4 years, 5-8 years, and 9 or more years.[1] Although we chose these categories for their consistency with and comparability to the categories routinely applied to employers in the United States, we should note that our age and size characteristics refer to the firm (organization), not to the establishment or plant (subunit).

We excluded foreign subsidiaries and other cases of incumbent firms creating new establishments. To err on the conservative side, we also removed from the analysis startups where a large proportion of the employees had worked together in the prior year for an employer in the same industry and region, but with a different firm identification code. Although this procedure probably resulted in the elimination of several startups that had been founded by groups of individuals from the same employer, treating some well-established entities as start-ups seemed more problematic than the exclusion of a few employee spin-offs.

Table 1 reports the median, mean, and standard deviation of the annual earnings for all full-time employees, between the ages of 18 and 60, in each of these size and age

[1] Occasionally, a firm had no employees associated with it for one or more years. In cases where a firm had no employees for a single year, we treated it as though it had been operating continuously. In cases where a firm had no employees for multiple consecutive years, we reset its age to one upon reentry.

categories across the entire period, 1991-2006, for our population of over 12 million employee-years. We also report the number of employee observations and the number of firms for each category of firm age and size. Looking down the columns, one can see a clear size gradient. Within each age range, larger firms paid more than smaller ones. The smallest employers – those with 1-10 employees – paid their employees 18-23% less than those with 250 or more employees. Looking across the rows, patterns become more difficult to detect.

Table 1: Mean and median annual earnings for all employees by firm size and firm age

	1-2 years	3-4 years	5-8 years	9+ years	Total
1-10 employees					
Mean	239,938	239,224	235,315	225,448	231,334
Median	216,996	216,217	214,478	210,107	212,791
Standard Deviation	160,620	325,958	153,750	133,422	217,292
Observations	414,349	325,888	474,790	1,363,716	2,578,743
Number of firms	110,634	77,100	66,470	72,007	326,211
11-49 employees					
Mean	278,311	282,161	281,200	265,256	270,050
Median	246,560	249,440	249,078	239,992	242,641
Standard Deviation	185,361	188,510	184,084	160,083	179,867
Observations	208,350	221,110	405,741	1,861,437	2,696,638
Number of firms	8,720	8,445	9,930	16,293	43,388
50-249 employees					
Mean	291,341	292,146	297,941	284,591	286,493
Median	253,899	252,599	257,594	251,433	252,176
Standard Deviation	191,051	194,368	196,895	182,429	191,264
Observations	122,053	127,559	261,668	2,265,351	2,776,631
Number of firms	1,255	1,268	1,658	4,061	8,242
250+ employees					
Mean	311,357	309,104	305,273	289,589	292,215
Median	283,171	279,605	270,175	255,302	258,268
Standard Deviation	167,542	172,877	191,917	174,679	176,990
Observations	167,406	167,064	326,370	3,917,366	4,578,206
Number of firms	302	289	307	940	1,838
Total					
Mean	268,688	272,397	274,698	274,273	273,794
Median	240,832	243,042	244,098	244,790	244,307
Standard Deviation	176,584	228,829	182,432	163,723	189,494
Observations	912,158	841,621	1,468,569	9,407,870	12,630,218
Number of firms	120,911	87,102	78,365	93,301	379,679

Movers: One of the most consistent complications noted in the prior literature on the relationship between firm age and wages has been that older firms also tend to employ individuals who have longer tenure with the firm (Brown and Medoff 2003; Heyman 2007; Ouimet and Zarutskie 2014). Prior research has typically dealt with this issue by including firm tenure as a covariate. But if the returns to firm tenure decline over time or vary across individuals, any linear adjustment for firm tenure may underestimate the "true" firm tenure effect and would then probably attribute a portion of these tenure differences to firm age.

We adopted a quite conservative approach to addressing this issue by examining only new hires and the wages that they earned. By definition, these individuals have no prior experience in the firm; therefore, our estimates compare similar individuals – at least in terms of firm tenure – across both young and old firms. In particular, we restricted the sample to those who had worked for a firm for at least 30 days but no more than one year. We excluded all individuals listed as founders, employers, or entrepreneurs, as their compensation may involve equity as well as wages. These restrictions reduced our sample size, but still left us with over 3.1 million observations across more than 260,000 firms.

Table 2 reports the median, mean, and standard deviations of the wages for these recent hires as well as the number of observations and firms for each age and size category. Note that Table 2 reports consistently lower wages than Table 1, reflecting the fact that pay rises with firm tenure, as one would expect. Reading down the columns, one continues to see the strong relationship between firm size and wages, with the largest firms paying new hires 18% to 21% more than the smallest ones. But the pattern for firm age changes noticeably. Looking across the rows, one now notices a *negative* relationship between firm age and the average wages paid to recent hires. Within each of the size categories, the oldest firms paid the lowest wages. The youngest firms paid 9% to 13% more to new hires than the oldest ones, and both mean and median wages at firms in the youngest two age categories (4 or fewer years old) consistently exceeded the overall mean and median for each size category.

Table 2: Mean and median annual earnings for new hires by firm size and firm age

	1-2 years	3-4 years	5-8 years	9+ years	Total
1-10 employees					
Mean	221,932	212,674	207,906	199,925	209,133
Median	209,496	202,911	199,888	193,638	200,382
Standard Deviation	140,783	128,440	123,108	112,021	126,512
Observations	142,098	84,950	105,506	216,144	548,698
Number of firms	63,030	40,195	38,360	46,720	188,305
11-49 employees					
Mean	259,284	252,437	249,179	235,099	242,785
Median	237,420	232,509	230,577	220,225	225,746
Standard Deviation	161,576	147,481	140,816	136,324	146,859
Observations	70,738	59,604	95,501	306,166	532,009
Number of firms	7,785	7,480	8,876	14,556	38,697
50-249 employees					
Mean	274,507	264,220	266,412	253,423	257,130
Median	242,250	234,108	238,448	230,167	232,253
Standard Deviation	190,235	172,847	155,758	146,250	167,119
Observations	33,856	30,255	58,874	363,947	486,032
Number of firms	1,118	1,155	1,530	3,818	7,621
250+ employees					
Mean	298,975	276,950	267,888	258,301	262,858
Median	263,854	250,740	243,772	232,402	236,495
Standard Deviation	187,801	141,998	148,116	152,463	158,600
Observations	42,794	28,501	62,318	496,044	629,657
Number of firms	275	269	298	889	1,731
Total					
Mean	248,597	241,012	242,431	242,750	243,313
Median	228,186	222,936	224,518	223,062	223,939
Standard Deviation	171,308	148,566	142,464	137,629	150,545
Observations	289,486	203,310	322,199	1,382,301	2,197,296
Number of firms	72,208	49,099	49,064	65,983	236,354

Adjusted wages by age and size

Employee sorting: Although restricting the sample to recent hires accounts for differences on the most obvious dimension on which young and old firms differ – firm tenure – employees might nonetheless sort into firms on a host of other characteristics related to productivity and therefore also to expected wages. We first explored the extent to which young and small firms differed from the overall population of firms in terms of the individuals they hired. Although past studies have reported differences between younger and older firms in the characteristics of their employees (Nystrom and Elvung 2014; Ouimet and Zarutskie 2014), the cross-sectional information on which those studies have relied depends on the joint combination of differential hiring, maturation, and differential retention. Whether young firms, in fact, hire different kinds of individuals therefore remains an open question.

Table 3 reports the demographic characteristics for the full population of new hires and for two relevant subsets. The first column presents averages across all firms for any new hire. The second column, meanwhile, restricts this set to those hires coming from a business establishment that closed. We include this group because one might worry that movers self-select into moving and therefore differ in important ways from non-movers. By focusing on those who had to find a new job due to the closing of their prior establishment, examining this subsample should eliminate most – if not all – self-selection (Gibbons and Katz 1991; Gruetter and Lalive 2009). Although our sample includes the service sector, we refer to this subsample as the "plant closings" group. Note that our sample size falls dramatically from more than 2 million observations to fewer than 215,000. Interestingly, within this restricted set of involuntary movers, the means and medians rise for all age and size bins, suggesting adverse selection on average among movers. One can also see that by comparing the average prior wages across columns one and two in Table 3.

The next two columns report means for the subsets of all hires and of plant closing movers moving to young firms. The final two columns do the same for those moving to small firms.

Interestingly, across all employee characteristics, few differences exist between the populations of new hires in the first and second columns and those going to young firms in the third and fourth columns. Somewhat larger differences appear between smaller and larger firms, comparing the first and second columns to the fifth and sixth columns. Smaller employers, particularly in the plant closing group, hired less educated individuals, who had less labor market experience, and had longer spells of unemployment. The compositional differences in who startups hire, therefore, appears to be more of a size effect than an age effect.

Table 3: Sample demographics by move type and destination

	All		Young (<5 yrs)		Small (<50 empls)	
	All hires	Plant closing	All hires	Plant closing	All hires	Plant closing
Age	34.02	37.15	34.42	37.09	33.88	36.15
	(10.17)	(10.68)	(10.12)	(10.61)	(10.31)	(10.86)
Female	0.34	0.33	0.34	0.34	0.33	0.34
	(0.47)	(0.47)	(0.48)	(0.47)	(0.47)	(0.47)
Months of education	148.31	147.79	148.51	147.65	146.71	144.98
	(28.28)	(29.24)	(28.25)	(28.97)	(27.46)	(28.06)
Type of education						
Primary school	0.27	0.28	0.27	0.27	0.29	0.31
	(0.44)	(0.45)	(0.44)	(0.45)	(0.45)	(0.46)
High-school/gymnasium	0.10	0.08	0.09	0.08	0.09	0.08
	(0.30)	(0.28)	(0.29)	(0.28)	(0.28)	(0.28)
Vocational training	0.44	0.45	0.45	0.46	0.47	0.46
	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)
College	0.13	0.13	0.12	0.12	0.11	0.10
	(0.33)	(0.34)	(0.33)	(0.32)	(0.31)	(0.30)
University	0.06	0.06	0.07	0.06	0.05	0.04
	(0.25)	(0.25)	(0.25)	(0.24)	(0.22)	(0.20)
Labor market experience	13.80	16.04	13.65	15.55	13.53	14.62
	(8.88)	(9.26)	(8.83)	(9.19)	(8.73)	(9.09)
Unemployment history	1.27	1.13	1.34	1.18	1.37	1.42
	(1.93)	(1.84)	(1.99)	(1.89)	(2.01)	(2.06)
Ln(Prior wage)	9.22	10.10	9.11	9.99	9.00	9.17
	(5.20)	(4.67)	(5.28)	(4.74)	(5.28)	(5.19)
Observations	2,197,296	213,692	492,796	57,306	1,080,707	84,011

Note: Standard deviation in parentheses.

Employee matching: Although few prior studies have adjusted for these differences, those that have generally relied on adjustments through linear regression. In other words, researchers estimated a wage equation, effectively assuming each of the relevant human capital dimensions had additive effects on the expected wage, or its logged value (e.g., Brown and Medoff 2003). Having data on the entire population allowed us to adopt a more flexible and non-parametric approach to adjusting for individual differences. Rather than estimating a wage equation with linear adjustments for the effects of age, gender, education, and other factors, we instead matched on these characteristics and included a fixed effect for each matched group. Because the fixed effect adjusts for a specific combination of attributes, it effectively allows these attributes, such as education and experience, to have completely flexible relationships to earnings and to interact in their determination of wages (i.e. allowing the returns to one dimension of human capital to depend on the others).

To minimize the possibility that some confounding factor accounts for the results, one would ideally match cases and controls *exactly* on all relevant observed dimensions. Of course, with continuous variables, that proves impractical if not impossible as no two individuals may have, for example, been born at precisely the same instant or earn exactly the same amount. We therefore adopted a modified version of this approach, combining coarsened exact matching (CEM) on several dimensions with nearest-neighbor matching on income in the previous year. One can find extended discussions of the advantages of this approach in Iacus et al. (2012) and King and Nielsen (2015).

Our matching procedure operated as follows. We treated each cell in the firm age-size matrix as a subsample. For each employee within a subsample, such as those beginning jobs at companies with 1-10 employees that have been operating for 1-2 years, we found all observationally-equivalent individuals in our baseline category of large, established firms (those beginning jobs at employers with at least 250 employees and that have been operating for at least nine years). We considered two individuals to be observationally equivalent if they had the same gender (male/female), the same age (coarsened to the year of birth), the same level of education (coarsened to the highest degree: primary school only, high school or gymnasium, a vocational training certification, undergraduate college, or graduate level), and the same prior occupation (using the one-digit version of the occupational codes for Denmark which delineate 10 major occupational categories that distinguish between skilled and unskilled jobs and between white and blue collar occupations).[2]

Although this matching accounts for differences across employees on some of the most important factors influencing wages, workers differ on a host of difficult-to-observe dimensions that also affect productivity and pay. Most of these factors should, however, remain relatively stable for a given individual over relatively short intervals of time. We therefore used information on the prior wages of individuals to account for these differences. From the set of available individuals who matched exactly on gender, age, education, and occupation, we only included the two nearest neighbors on the prior year wage distribution – the closest observation above and the closest below what an employee earned in that previous year – in the comparison set. To achieve high quality matches and ensure that the matching estimator replicates the population average treatment effect (Abadie and Imbens 2006), we matched with replacement, meaning a control could serve as a match for multiple focal individuals. This procedure yielded statistically-identical average wages across all sets of cases and controls.

Consider an example. Beginning first with the individuals who joined small (1-10 employees), young (1-2 years) firms (the top left cell in our tables). For the 142,098 "focal" individuals who joined these firms (see Table 2), we found control individuals who joined large (250+ employees), established (9+ years) firms (the baseline category) in the same year, who matched the focal individuals on age, gender, education, and prior occupation. For each focal individual, we selected the exact match closest but just above the person in earnings

[2] The number of available matches limited our ability to employ more fine-grained occupational categories at this stage. Including more detailed two-digit occupational categories as fixed effects, however, did not absorb much additional variance, consistent with the results reported by Schmieder (2013).

$(t-1)$ and the exact match closest but just below in earnings $(t-1)$ forming an observation triad. We successfully identified matches for 135,530 (98%) of the focal individuals.[3]

Adjusting for human capital: For each of our 15 matched samples, we estimated the effect of being in the "treated" group (that is, not being employed by a firm in the oldest and largest category). Specifically, we estimated the following equation:

$$Ln(W_i) = \beta_{as} AS_i + \gamma_j + E_i, \tag{1}$$

where W_i represents the starting wage for individual i, AS_i denotes a dummy variable that takes the value one when the individual in question works for a firm in the younger age and/or smaller size category, γ represents a vector of fixed effects specific to each triad j (i.e. a focal individual plus two matched controls), and E_i denotes an individual-specific error term. By adjusting for individual characteristics through a series of fixed effects, this model controls flexibly for any shape that the relationship between each of these factors and wages might take, as well as for any interactions between these characteristics in the determination of wages. Because both the cases and controls in each set changed jobs in the same year, these fixed effects also absorb any period effects, such as the business cycle. We repeated this procedure for each of the 15 matched samples.

Table 4 reports the β_{as} values from these 15 regressions. In the interest of saving space, each cell in the tables below simply reports the β_{as} coefficient and standard errors for the regression using the relevant matched sample. We also report the number of case-control triads used in each regression.

One can read the value in each cell as estimating the pay for observationally-equivalent new hires in a firm in that particular age and size bin relative to the pay offered in an established (9+ years), large (250+ employees) firm. Thus, for example, the top left cell indicates that an individual hired by a firm in the smallest, youngest group would receive an

---

[3] In total, we have 15 sets of matched samples (one for each cell in the age-size matrix, except for the baseline category) and obtained a match rate of 97% or better across 11 of them. Our lowest match rate, 78%, occurred in the smallest (1-10 employees) and oldest (9+ years) category.

annual wage only 84% (exp(-0.179)=0.836) as large as a similar individual hired by a large, established firm. Given that the average new hire in a large, established firm earned about 258,300 Danish kroner (DKK), the comparable person hired by the small start-up would earn about 216,000 DKK, an annual difference of approximately $6,500

Table 4: Regression matrix: New hires matched on age, gender, education, prior job, and prior earnings

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees	-0.179*	-0.204*	-0.206*	-0.192*
	(0.003)	(0.003)	(0.003)	(0.002)
Triads	135,520	81,548	101,582	208,930
11-49 employees	-0.046*	-0.061*	-0.062*	-0.071*
	(0.003)	(0.003)	(0.002)	(0.002)
Triads	68,300	57,598	92,550	297,818
50-249 employees	-0.006	-0.013*	-0.014*	-0.019*
	(0.004)	(0.004)	(0.003)	(0.001)
Triads	32,744	29,322	57,147	354,571
250+ employees	0.094*	0.041*	0.021*	
	(0.003)	(0.004)	(0.003)	
Triads	41,319	27,603	60,507	

* $p < 0.01$. Robust standard errors in parentheses.
Note: All regressions include fixed effects for each matched triad of one individual from the treatment cell and two matched individuals from the baseline cell of firms with 250 or more employees that are more than 9 years old.

Reading across the first row, all four firm age categories have negative wage coefficients, revealing a sizable wage penalty for employees hired by the smallest firms. The steepest gradient, moreover, appears as one moves across the first two ages groups. Reading these coefficient values down the columns, comparing wages in similarly-aged firms of different sizes, one can see a strong positive size effect. Wages penalties decrease as firm size increases. The size effects nevertheless appear larger for younger firms than for older ones, suggesting that firm age and size interact to some extent in determining wages.

Interestingly, some startups paid wages equivalent to or higher than the largest, most established firms. Notably, all three columns of the largest size category (250+ employees) have positive wage coefficients. Start-ups, particularly those that begin large or become large rapidly, would appear to create the highest-paying jobs. But how common are such firms and how prevalent are these jobs? In terms of firms, recall that nearly 90% of firms in the youngest column occupy the top-left cell, being both young and very small. Low paying startups therefore dominate the mix. But in terms of the typical job offering, because the larger firms account for more jobs, the numbers are more encouraging. Roughly one-quarter of jobs in startups pay a premium over those of large, established firms.

Methodological choices and their implications

Although our approach accounts for a variety of issues not addressed in previous studies, to what extent do the differences in our estimates versus earlier ones stem from our modeling choices versus from our setting? On the one hand, better accounting for firm tenure and employee characteristics may address issues in the estimates found in prior studies. On the other hand, Denmark may simply exhibit a different pattern of pay with respect to firm age and size. Even the prior inconsistencies in the estimates may stem not from the approaches used, but from the fact that the results have variously been from samples in the United States, Sweden, and Germany.

Table 5: Firm age and firm size effects for different samples and modeling choices

	Firm age 1-10	Firm size 5-250
(1) Population (Brown/Heyman)	-2.52%	12.69%
(2) Movers (Brown/Heyman)	-3.68%	16.98%
(3) Population (bins)	-3.96%	33.14%
(4) Movers (bins)	-9.31%	35.84%
(5) + matched individuals (Table 4)	-3.99%	23.70%
(6) + industry FEs	-5.25%	20.30%
(7) Plant closings (bins, matched, industry FE)	-4.79%	23.78%
(8) Movers (bins, matched, industry FE, future growth)	-3.87%	18.68%
(9) Entrants (bins, matched, industry FE)	-3.21%	14.95%

To address these issues, we estimated the magnitudes of the firm age and firm size effects using a variety of methodological techniques in order to mimic prior research. We summarize these results in Table 5. Table 5 has two columns: the first column reports the estimated wage differential associated with moving from a startup to a firm that has been operating for ten years; the second column reports the estimated wage differential associated with moving from a firm with five employees to one with 250. Each row corresponds to a set of estimation choices.

We first replicated the approach employed by both Brown and Medoff (2003) and Heyman (2007). This model used the full population as a sample, adjusting for both a variety of employee and firm characteristics as covariates. It estimated the effects of firm age and of firm size as log functions. As noted above, accounting for employee tenure may represent one of the most important adjustments, particularly when trying to understand the effects of firm age. The second row, therefore, implements the same method but using the sample of job changers instead of including firm tenure as a covariate (Full results and more detailed explanation these analyses are available in an online appendix).

Beginning with the third row, the table explores the various elements of our modeling approach. It reports the results of the estimated effect sizes using our approach to binning age and size. Although the splining used by Heyman (2007), in his supplementary analysis, and by Schmeider (2013) also allows for non-linearity, the binning further allows for firm age and firm size to interact in their determination of wages. Note that when calculating the average effects of firm age and firm size in our binned models, the reported values weight the effect in each row (column) according to the number of jobs represented by that row (column). The third row reports these estimates for the full population, the fourth row for only job changers (movers). The fifth row reports estimates for the models that we report in Table 4, on the sets of matched observationally-equivalent individuals.

Comparing across these five rows, one can see three effects. First, the binned approach, with its flexible functional form, generally produced much larger firm age and firm size effects. Two factors combine to produce this result: (i) the interactive effects of firm age and firm size, and (ii) the rate of change in the firm size-wage relationship. As one can see in Table 4, the youngest firms differ more in their wages as a function of size than the oldest ones. The largest firms also differ most in their wages as a function of age. Those effects remain consistent across all variations of the models discussed below, suggesting a strong interaction between firm age and firm size. In addition, the effects of firm size reach an asymptote faster than a linear or log function would predict.

Second, focusing on job changers yields much larger effects, particularly for firm age. The larger age effect probably stems from the linear effect for tenure under-correcting for the relationship between firm tenure and pay. Such an under-correction can lead to a positive bias in the estimates of the relationship between firm age and wages. Note that, if large enough, failure to account adequately for tenure effects could even flip the sign of the firm age effect positive, such that older firms would appear to pay more. In other words, using a linear adjustment for tenure may account for why Heyman (2007) found a positive relationship between firm age and wages.

Third, adjusting for compositional differences in the characteristics of the employees of firms of different ages and sizes accounts for a large share of the average observed differences at the firm level, roughly one-third of the size effect and more than one-half of the age effect, similar in magnitude to the results reported by Brown and Medoff (2003) and Heyman (2007). Failure to account for these differences will therefore generally lead to an overestimation of the firm age and size effects.

Adjusting for industry: Although the adjustments made up until this point address most of the factors that might confound the relationship between firm age and wages, they do not account for the fact that the firm age and size distributions might vary systematically across industries. New, rapidly growing, industries, for example, might have an unusual number of small, young firms. They may also face a thin labor market in which talent commands a wage premium. What appears to be a firm age or firm size effect might then actually reflect an industry effect on wages.

To account for the differences across industries, we re-estimated the models including fixed effects for four-digit industries. We found that the firm age penalty increases slightly and the firm size premium declines somewhat. But the addition of more than 500 industry intercepts notably had relatively little effect on either the magnitudes or the patterns of the wage differentials.

Accounting for selection into mobility: Although our approach – considering only the wages of recent hires – has the advantage of holding constant firm tenure, one might worry that these job changers differed systematically on other factors from those who remained with their employers. But the direction of this bias remains uncertain. On the one hand, the least productive employees might get fired and need to find new jobs. On the other hand, the most productive ones might move in search of more attractive job opportunities.

To prevent such selection from influencing our estimates, we further restricted the sample to include only those individuals who had left their prior employers because the plant or business location at which they had been working closed ("plant closings"). Although this restriction reduces our sample size dramatically, it allows us to examine the subset of individuals who plausibly sought employment for reasons exogenous to their individual ability or productivity (Gibbons and Katz 1991; Gruetter and Lalive 2009). Again, the general patterns of wage penalties and premia with respect to firm age and size appear largely the same for these involuntary movers.

Exploring alternative explanations

Although our results suggest that the firm age and firm size effects remain robust to a variety of modeling choices, one might nonetheless worry that they stem from some form of unobserved heterogeneity. For example, if firms can accurately assess the quality of potential employees not captured in their prior wages, and if individuals accurately evaluate the prospects of their potential employers, then one might see assortative matching—the most productive employees joining the startups with the greatest potential. Perhaps the "sure bets" can pay higher wages and therefore attract the best employees? Dahl and Klepper (2015), for example, found that those firms with the best survival prospects, based on the attributes of the founders and of the firm at the time of its founding, paid somewhat higher wages than firms with worse survival prospects. Note that differences in the future prospects of firms might stem from a wide variety of factors. They might result from quality differences in the founders or their ideas. They could stem from externalities, such as being located in an industrial cluster. Or they might reflect the underlying ambitions of the founders of the firm. Although one might expect the industry intercepts to capture some of these differences, substantial variation probably exists even within industries.

Unobserved firm characteristics: To address this possibility we took advantage of the longitudinal nature of our data. For each year of the sample and for each firm in the sample, we considered its *future* growth for the next five years (defined in terms of the number of employees in year $t+5$ divided by the number in year t). Schmieder (2013) also argues that growth rates should account for the wage differential between young and old firms to the extent that those differences emerge from older firms effectively having more bargaining power vis-à-vis employees. Although distinguishing monopsony from assortative matching would prove difficult, for our purposes, we care simply whether accounting for differential growth rates attenuates the observed negative relationship between firm age and wages. Note that we do not have five-year forward projections for all firms, because many firms fail. We excluded any firms without $t+5$ data from this analysis. To allow for a flexible relationship between firm growth and wages, we used future growth to assign each firm to a growth decile (across all firms in the sample for that year) and included a vector of indicator variables to capture any wage differentials associated with firm growth rates.

Although the general patterns with respect to firm age and firm size remain the same, they do contract somewhat in magnitude—on the order of 25% for firm age and 8% for firm size (see the second to the last row in Table 5). We would note that these results seem quite consistent with those of Gibson and Stillman (2009). Using rich and detailed measures of worker skills, they found little evidence for the idea that the sorting of better employees into larger firms could account for the firm-size wage effect. Our results suggest that sorting also probably does not account for the firm-age wage effect.

Unobserved employee characteristics: Because the analyses above match individuals according to their wages in their prior jobs, they effectively restrict the sample to those already in the labor market. Although this approach has some advantages, in terms of more tightly accounting for difficult-to-observe differences in human capital and productivity, it potentially also raises some issues. Movers within this sample, for example, may sort into larger versus smaller and younger versus older firms based on their prior labor market experience. Focusing on involuntarily movers – those employed at plants that closed – partially accounts for these differences, but prior experience may still allow workers to signal their quality.

We therefore estimated the wage premia and penalties again using only new entrants to the labor market, following Nystrom and Elvung (2014). Firms hiring labor market entrants necessarily have much weaker signals of worker quality. This subpopulation should therefore have much less potential for sorting employees to employers on the basis of quality or productivity. Although these estimates parallel those for Table 4, plus industry fixed effects, because these individuals have not had jobs, we could not include nearest-neighbor matching on prior wages. Though somewhat smaller in magnitude, the same pattern of results appears in this subpopulation (see the last row of Table 5). Some form of assortative matching therefore seems unlikely to account fully for either the firm age of the firm size effect.

Discussion

Do startups pay well? Our answer seems mixed: Most do not, but a few do. We explored the relationship between the amount that firms paid and firm age and size among the population of Danish employers and employees and found both a firm size effect and a firm age effect on the wages of new hires. Larger firms paid recent hires more than smaller ones, even for observationally-equivalent individuals who had earned roughly the same amount in their previous jobs. On the other hand, young firms actually paid recent hires *more* than older firms of similar size. But firm size had larger effects than firm age. Hence, to the extent that startups begin both young and small – nearly 90% of firms in our population do – they do tend to pay less than large, established firms.

But those rare startups that begin large, or become large very quickly, actually pay a premium relative to more established employers. Firms four years of age or less with at least

250 employees paid substantial premia over more established firms. Although these firms amounted to a small minority of employers, because each of them hired hundreds of individuals, they accounted for roughly one-fifth of the jobs created by firms in their first four years.

Our most novel finding, however, is that young firms paid more than older ones. Why might they have done so? One possibility is that startups need to compensate for the greater instability of the jobs that they offer. Because the average startup has a half-life of only four years, employees face a substantial risk of losing their jobs as a consequence of the firm itself failing. Higher wages, therefore, may provide something of a compensating differential for this instability. If they do, one might reasonably ask whether startups should pay even larger premia. Not only do these firms fail at high rates, meaning that employees may find themselves involuntarily unemployed or looking for a job, but also the probability of failure probably rises during periods of economic contraction, precisely the times during which laid-off employees would find it most difficult to secure another job.

A second possibility is that startups, particularly those that start at a larger scale or that grow rapidly, face different labor market challenges than established firms. In particular, startups may want to hire numerous employees in a short period of time. They may therefore need to offer something akin to signing bonuses to entice would-be employees (Schmieder, 2013). Although this explanation might account for a portion of the differential, the fact that controlling for future growth had only a modest effect on the relationship between firm age and wages suggests that it cannot be the full story.

In addition to establishing a set of empirical facts about wages for the jobs being created by startups, we believe that our research also offers multiple methodological contributions. First, we disentangle the effects of firm age and firm tenure on wages by comparing new hires to all employees. In eliminating the effects of firm-specific human capital accumulation, we reveal that prior research that fails to account for employee tenure likely biases the estimates toward finding a more positive relationship between firm age and wages. Although this bias simply attenuates the negative relationship in our sample, this bias might account for why Heyman (2007) concluded that older firms pay more.

Second, we demonstrate that firm age and size are neither linearly (nor log linearly) related to wages nor are they independent in their effects. Instead, firm age and firm size jointly affect wages. Failure to accommodate these interactions appears to lead to underestimation of the effects of firm age and firm size. Although strategy of binning firms into age and size categories challenges common empirical practice, it accords well with the idea that firms fall into conceptual categories (e.g. micro-enterprises, mittlestand or SMEs, high-growth firms or gazelles). Though not a concern in the Danish context, this approach can also account for the fact that laws and regulations often come into play only above specific size thresholds (for a thorough discussion of these thresholds in the United States, see Eyal-Cohen 2013).

Third, we provide additional evidence that human capital characteristics can become a confound to the effects of firm characteristics given that firms of different ages and sizes draw from somewhat different labor pools. One of the difficulties in assessing job quality is that one cannot say whether one job is better than another without understanding the characteristics of the would-be occupants of those jobs. Being a truck driver, for example, might pay well relative to the alternatives for someone lacking a high school degree. Although the extant research has been aware of this issue, the typical approach to adjusting for job holder characteristics has been to include the observed characteristics of job holders as covariates in a wage equation (or in regressions on some other measure of job quality). That approach however, has the limitation of essentially requiring one to assume that these characteristics have additive (and usually linear or log-linear) relationships to productivity and wages.

The increasing availability of longitudinal registry data, however, opens the door for alternative approaches. The Danish data, for example, include more than 20 million person-years of information. Instead of adjusting for observed characteristics through regression, we used matching to create sets of cases and controls nearly identical on the observed dimensions and allowed each group – with its potentially unique combination of characteristics – to have its own intercept. Doing so allows us to adjust for the characteristics of the employees without requiring any assumptions about the functional forms of the relationships between these characteristics and wages, or about the ways in which these attributes may interact in determining wages.

Overall, these methodological issues appear sufficient to account for apparent inconsistencies in the prior literature. But without estimating similar sets of models on data from each of the countries, one cannot say for certain whether they do. Notably, Heyman (2007) finds somewhat inconsistent results even using the same method on different cross sections from the same country. The relationships between firm age, firm size, and wages likely also depends to some extent on the context.

Although our results provide some initial insight into the question of whether startups create good jobs, they represent more of a first step in a research agenda than a definitive answer. Consider some of the closely-related questions that remain open: Although young firms pay their employees a premium in the first year, how do these effects evolve? Do the employees of younger and older firms experience similar wage trajectories or do their wages change at different rates? It would seem that these effects might go either way: On the one hand, rapidly growing firms might promote employees faster and give them larger raises. On the other hand, the managers of young firms with higher probabilities of failure may invest less in firm-specific human capital that would enhance their productivity over time.

Entrepreneurship has been and will continue to be an important driver of economic vitality. As such, understanding better how the jobs created by entrepreneurs affect the earnings and lives of the people who occupy them will importantly inform both policy and practice.

References

Abadie, Alberto, Guido W. Imbens. 2006. Large sample properties of matching estimators for average treatment effects. *Econometrica* **74**(1) 235–267.

Abowd, John M., Francis Kramarz, David N. Margolis. 1999. High wage workers and high wage firms. *Econometrica* **67**(2) 251–333.

Anyadike-Danes, Michael, Carl-Magnus Bjuggren, Sandra Gottschalk, Werner Holzl, Dan Johansson, Mika Maliranta, Anja Myrann. 2015. An international cohort comparison of size effects on job growth. *Small Business Economics* **44**(821-844).

Arrow, Kenneth J. 1962. The economic implications of learning by doing. *Review of Economic Studies* **29**(3) 155–173.

Audretsch, David B., George van Leeuwen, Bert Menkveld, Roy Thurik. 2001. Market dynamics in the Netherlands: Competition policy and the role of small firms. *International Journal of Industrial Organization* **19** 795–821.

Ayyagari, Mechana, Asli Demirguc-Kunt, Vojislav Maksimovic. 2014. Who creates jobs in developing countries? *Small Business Economics* **43** 75–99.

Bingley, Paul, Niels Christian Westergård-Nielsen. 2003. Returns to tenure, firm-specific human capital and worker heterogeneity. *International Journal of Manpower* **24** 774–788.

Brixy, Udo, Susanne Kohaut, Claus Schnabel. 2007. Do newly founded firms pay lower wages? First evidence from Germany. *Small Business Economics* **29**(1-2) 161–171.

Brown, Charles, James L. Medoff. 2003. Firm age and wages. *Journal of Labor Economics* **21** 677–696.

Dahl, Michael S., Steven Klepper. 2015. Whom do new firms hire? *Industrial and Corporate Change* **24**(4) 819–836.

de Wit, Gerrit, Jan de Kok. 2014. Do small businesses create more jobs? New evidence for Europe. *Small Business Economics* **42** 283–295.

Eriksson, Tor. 2001. How common are the new compensation and work organization practices and who adopts them? Working paper, Aarhus University.

Eyal-Cohen, Mirit. 2013. Down-sizing the little guy myth in legal definitions. *Iowa Law Review* **98**(3) 1041–1100.

Gibbons, Robert, Lawrence F. Katz. 1991. Layoffs and lemons. *Journal of Labor Economics*

9(4) 351–380.

Gibson, John, Steven Stillman. 2009. Why do big firms pay higher wages? Evidence from an international database. *Review of Economics and Statistics* **91**(1) 213–218.

Griliches, Zvi. 1969. Capital-skill complementarity. *Review of Economics and Statistics* **51** 465–468.

Gruetter, Max, Rafael Lalive. 2009. The importance of firms in wage determination. *Labour Economics* **16**(2) 149–160.

Haltiwanger, John, Ron S. Jarmin, Javier Miranda. 2013. Who creates jobs? Small versus large versus young. *Review of Economics and Statistics* **95** 347–361.

Heyman, Fredrik. 2007. Firm size or firm age? The effect on wages using matched employer- employee data. *Labour* **21** 237–263.

Iacus, Stefano, Gary King, Giuseppe Porro. 2012. Causal inference without balance checking: Coarsened exact matching. *Political Analysis* **20**(1) 1–24.

Ibsen, Rikke, Niels Christian Westergård-Nielsen. 2011. Job creation by firms in Denmark. Tech. Rep. DP No. 5458, IZA.

King, Gary, Richard Nielsen. 2015. Why propensity scores should not be used for matching. Working paper, Harvard University.

Lawless, Martina. 2014. Age or size? Contributions to job creation. *Small Business Economics* **42** 815–830.

Madsen, Jørgen Steen, Søren Kaj Andersen, Jesper Jørgen Due. 2001. From centralised decentralisation towards multi-level regulation. *Proceedings of the 6th European IIRA Congress*.

Madsen, Per Kongshøj. 2004. The Danish model of ‘flexicurity’: Experiences and lessons. *Transfer: European Review of Labour and Research* **10**(2) 187–207.

Malchow-Moller, Nikolaj, Bertel Schjerning, Anders Sorensen. 2011. Entrepreneurship, job creation and wage growth. *Small Business Economics* **36**(1) 15–32.

Nystrom, Kristina, Gulzat Zhetibaeva Elvung. 2014. New firms and labor market entrants: Is there a wage penalty for employment in new firms? *Small Business Economics* **43** 399–410.

OECD.2015. *OECD Better Life Index*.URL: `http://www.oecdbetterlifeindex.org/countries/denmark`.

Oi, Walter Y., Todd L. Idson. 1999. *Handbook of Labor Economics*, vol. 3, chap. Firm size and wages. Elsevier Science, 2165–2214.

Ouimet, Paige, Rebecca Zarutskie. 2014. Who works for startups? The relation between firm age, employee age, and growth. *Journal of Financial Economics* **112** 386–407.

Schmieder, Johannes. 2013. What causes wage dispersion? Evidence from new firms. Working paper, Boston University.

Sørensen, Jesper B., Olav Sorenson. 2007. Corporate demography and income inequality. *American Sociological Review* **72** 766–783.

Sorenson, Olav, Michelle Rogan. 2014. (When) do firms have social capital? *Annual Review of Sociology* **40** 261–280.

Thompson, Peter. 2001. How much did the Liberty shipbuilders learn? New evidence for an old case study. *Journal of Political Economy* **109**(1) 103–137.

Timmermans, Bram. 2010. The Danish Integrated Database for Labor Market Research: Towards demystification for the English-speaking audience. DRUID Working Paper No. 10-16.

Troske, Kenneth R. 1998. *Labor Statistics Measurement Issues*, chap. The worker establishment characteristics database. University of Chicago Press, Chicago, 371–403.