Do startups create good jobs?*

Olav Sorenson[†] Michael S. Dahl[‡] M. Diane Burton[§]
Yale University Aalborg University Cornell University

June 25, 2015

Abstract: We analyze Danish registry data from 1991 to 2006 to determine how firm age and size influence wages. Unadjusted statistics suggest that smaller firms pay less than larger ones and that firm age has no bearing on wages. After adjusting for differences in the characteristics of employees hired by these firms, however, we observe both firm age and firm size effects. We find that larger firms pay more than smaller firms for observationally-equivalent individuals but, contrary to conventional wisdom, that younger firms pay *more* than older firms. Moreover, we find that the size effect dominates the age effect. Thus, while the typical startup – being both young and small – pays less than a more established employer, those that grow rapidly often pay a wage premium.

*We thank Yale University for generous financial support. The usual disclaimer applies.

†165 Whitney Ave., P.O. Box 208200, New Haven, CT 06520, olav.sorenson@yale.edu

‡Entrepreneurship and Organizational Behavior, Fibigerstræde 11, DK-9220 Aalborg Ø, Denmark, md@business.aau.dk

§ILR School, 170 Ives, Ithaca, NY 14853-7601, burton@cornell.edu

Introduction

Entrepreneurship and the idea that entrepreneurs would create jobs, reduce unemployment, and stimulate economies has captured the imaginations of policymakers around the globe. A great deal of research in economics, moreover, suggests that these hopes have some basis in reality (Audretsch 2007). Early studies, by Birch (1987), for example, pointed to small establishments as the engines of job creation in the United States. According to his analysis of data from the early 1980s, firms with fewer than 20 employees accounted for more than 80% of gross job creation. Early studies, however, had a number of limitations. They could not, for example, distinguish startups from existing firms that had moved or had opened additional plants or offices, and they focused on gross, as opposed to net, job creation and therefore did not account for the fact that some newly-created jobs might simply displace existing ones (Davis et al. 1996). More recent research, using detailed microdata to address these limitations, has, if anything, shined even brighter light on the employment growth benefits of entrepreneurship. Haltiwanger et al. (2013), for example, found that it is not small firms, but rather startups – firms in their first few years of operation – that account for an outsized share of all net job creation in the United States (see also Audretsch 2002). Most other countries, moreover, appear to exhibit quite similar patterns (Ayyagari et al. 2014; de Wit and de Kok 2014; Lawless 2014; Anyadike-Danes et al. 2015).[1]

Largely absent in this literature on startup job creation, however, has been a consideration of the quality of the jobs that are being created, in terms of the salaries that they pay and the benefits that they offer. If the process of creative destruction largely involves the replacement of higher paying jobs at incumbent firms with lower paying ones at startups, then a simple

[1]For evidence on job creation specific to Denmark, see Malchow-Moller et al. (2011) and Ibsen and Westergård-Nielsen (2011).

examination of the numbers of jobs created, even net of jobs lost, may overstate the value of entrepreneurial activity to the economy. We would therefore like to know more about whether the jobs being created by startups are better or worse than existing ones.

Prior research does shed some light on this question. Studies have, for example, found that larger firms pay more on average than smaller firms (Davis and Haltiwanger 1991; Oi and Idson 1999), perhaps because economies of scale allow their employees to be more productive. Research has also demonstrated that older firms pay more on average than younger ones (Audretsch et al. 2001; Brixy et al. 2007), even after adjusting for differences in firm size. Older firms may enjoy higher productivity because competition weeds out the less productive firms over time or because firms become more productive as they gain experience and invest in equipment and infrastructure. Given these patterns, one might therefore expect that startups, being both young and small, would pay substantially less than more established firms.

But comparisons of average wages do not account for differences in the characteristics of the employees working at these firms. Recent research demonstrates that startups tend to attract a somewhat different set of employees than established firms. Ouimet and Zarutskie (2014), for example, find that the average employee of an entrepreneurial firm is younger, less educated, and less experienced than one would find in the workforce as a whole. Given that these individual characteristics influence the amount that an employee could expect to earn at any job, whether a startup or a more established employer, failure to account for these differences means that the apparent effects of firm age and size may instead reflect differences in the compositions of the workforces of these different types of firms. Two other issues further complicate attempts to understand the firm age wage effect: the tenure problem – where it is difficult to disentangle the effect of employee tenure from that of firm age – and the mobility problem – where wages, employment prospects, and bargaining power likely differ across voluntary and involuntary job changers.

We address these issues by using comprehensive registry data on the population of Danish workers, from 1991 to 2006, to examine how wages vary with firm age and size, and to assess the extent to which those differences remain after adjusting for differences in the characteristics of their workforces. Our large population size – nearly 30 million employee-years – allows us to contribute to and extend existing research on firm age and wages in at least five ways: (i) To eliminate the confounding effects of differential firm tenure, our estimates focus only on the wages paid to those newly hired; (ii) to further eliminate the potential selection effects associated with who leaves their prior employer, we also estimate effects among a sub-sample of individuals changing jobs because their prior employer closed; (iii) to account for differences in the characteristics of the individuals employed by firms of varying age and size, we use coarsened-exact matching to focus on effects within sets of observationally-equivalent individuals; (iv) to account in part for unobserved differences in the productivity of individuals, we further match individuals to their nearest-neighbors in terms of wages in their previous jobs; and (v) we also adjust for fine-grained (4-digit) industry differences in wages.

We demonstrate in our analysis below that these adjustments for differences in employee characteristics and tenure have important consequences for the apparent relationship between wages and firm age and size. Whereas larger firms, on average, pay more than smaller firms, roughly half of this effect appears to stem from differences in who they employ. Moreover, although firm age, on average, has almost no relationship to wages in Denmark, this negligible net effect appears to stem from the fact that young firms pay a wage premium but systematically employ workers who would receive less in any job. Startups, particularly those growing fast enough to overcome the size effect quickly, therefore do appear to create better-paying jobs.

Firm age, firm size, and wages

Despite the enthusiasm for entrepreneurship on the part of policymakers and the evidence from economists that startups account for the majority of net job creation, we still have reasons to be pessimistic about entrepreneurship as an engine for creating good jobs and generating broad-based economic benefits. A fairly substantial empirical literature has examined the relationship between firm size and compensation (for a review, see Oi and Idson 1999). Researchers typically find that larger firms pay more and offer better benefits than smaller ones (e.g., Brown and Medoff 1989; Davis and Haltiwanger 1991). Large firms enjoy economies of scale and scope, and, because firms generally only become large over time, they also may benefit from economies of experience and the favorable selection of firms with better business strategies and operational routines (Moore 1911; Doeringer and Piore 1971; Syverson 2011). Note, however, that relatively few of the studies of firm size and wages have adjusted for differences in the characteristics of the employees of larger versus smaller firms. Yet, larger firms also systematically employ individuals with more education and experience. Studies adjusting for this fact generally find much smaller wage premiums associated with firm size (Abowd et al. 1999; Troske 1999; Winter-Ebmer and Zweimuller 1999). But even these studies find a firm size wage effect and, to the extent that startups begin small, one might then expect them to pay poorly.

Startups may even pay less than older firms independent of these size effects. Fledgling firms, for example, 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. Due to their lower levels of capital investment and to the uncertainty surrounding their future prospects, startups may also prove less appealing to employees and therefore find themselves relegated to employing less-productive individuals than older organizations (Moore 1911; Kremer 1993).

Only a handful of studies to date, however, have examined the relationship between firm age and wages, net of firm size effects.[2] Troske (1998), for example, reports 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. This differential appeared to dissipate over time, though slowly: Even five years after their founding, these young firms continued to pay roughly 5% less than more established employers.

While there is an emerging consensus that older firms pay higher wages than startups, most of the studies informing this view 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. But not only do employees gain experience during their tenure with a firm but also research suggests that smaller and younger firms hire younger, less educated, and less experienced individuals (Nystrom and Elvung 2014; Ouimet and Zarutskie 2014). The apparent effects of firm age and size on wages therefore may stem more from who these firms hire than from differences across firms in their productivity or

[2]Even less research has examined the relationship between firm age and fringe benefits, but it appears to find a similar effect, with younger firms offering less generous benefits (e.g., Litwin and Phan 2013).

ability to pay (e.g., Abowd et al. 1999). After accounting for employee characteristics, Brown and Medoff (2003), for example, in a sample of 1,410 American workers found no significant relationship between firm age and wages. Heyman (2007) and Nystrom and Elvung (2014), meanwhile, using data on roughly 170,000 and 150,000 (before matching) Swedish employees, both found that – even after adjusting for employee characteristics – older firms paid slightly higher wages.

Although studies have begun to address the question of how wages vary with firm age, questions nevertheless remain. First, the only large-scale studies both rely on data from Sweden. Do other countries show similar patterns? Second and more importantly, the few studies that have adjusted for differences in employee characteristics have made strong assumptions about how these characteristics are related to wages. Scholars have either included individual-level covariates in a standard wage equation or have relied on propensity-score matching. Both of these techniques essentially assume that individual wages have either linear or log-linear relationships to wages (in the absence of firm-level effects). But much research suggests that wages and productivity may have more complex relationships with individual characteristics. For example, individual characteristics such as gender or experience could have complementary or substitutive effects on productivity and thus shape wages. How might the patterns change if one adopted a more flexible approach to adjusting for individual characteristics? Third, does firm age really have independent effects from employee tenure. Nystrom and Elvung (2014) address the issue by focusing on those with no experience, those first entering the labor market. But these early job matches tend to involve a high degree of instability and experimentation and therefore they may not represent well the dynamics of the labor market as a whole (Topel and Ward 1992).

Empirical Strategy

To advance our understanding of the relationship between firm age, firm size, and wages, we examined Danish registry data covering every employee in the country from 1991 to 2006 (using the Integrated Database for Labor Market Research, commonly referred to as the IDA database). In total, the database covers roughly 30 million person-years of employment data. The use of such a large dataset allowed us to adjust for a variety of factors on which the employees of large and established firms might differ from small or young firms.

Analysts widely regard Denmark as having one of the most flexible labor markets in Europe (Bingley and Westergård-Nielsen 2003; Sørensen and Sorenson 2007). This flexibility is the result of a series of wages reforms, the last of which went into place in 1991 (Madsen et al. 2001). Because observed firm differences prior to this point might stem, in part, from the unusual wage setting regime that existed at the time and may, therefore, not apply to countries with more decentralized systems, we built our panel dataset using only employees and employers from 1991 to 2006, the period following the wage reforms.

Despite the fact that Denmark has made it easy to hire and fire employees, and despite the fact that it has a flexible wage-setting regime, the country retains a strong social support net. Unlike the United States, for example, most benefits, such as health insurance and retirement plans, come from the central state rather than from employers. This fact also has the advantage for our purposes of ensuring that most of the differences between employers in the quality of jobs stems from the wages that they offer, rather than from a combination of wages and fringe benefits.

Average wages by age and size

To begin the analysis, we 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.[3] Although we chose these categories for their consistency with and comparability to the categories routinely used to classify 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).

Table 1 reports the median, mean and standard deviation of the wages for all employees in each of these size categories across the entire period. Looking down the columns, one can immediately see a size gradient. Within any age range, larger firms pay more than smaller ones. This differential appears most pronounced for younger firms, those that have been operating for fewer than nine years. Within these age categories, employers with 1-10 employees pay their employees 21% to 30% less, on average, than similar-aged firms with 250 or more employees. By contrast, looking across the rows, firm age has almost no apparent relationship to wages, once one controls for firm size. In most cases, the means across age ranges differ by no more than 2% to 4%. The exception would be the largest firms, which seem to have a non-monotonic pattern, with wages first rising and then falling with firm age.

[3]Occasionally, a firm will have no employees associated with it for one or more years and will then reenter the data. In cases where a firm has no employees for a single year, we treat it as though it has been continuously in existence. In cases where a firm has no employees for multiple consecutive years, we reset its age to one when it reenters. This coding decision nevertheless affects a relatively small number of firms and therefore has no meaningful influence on our estimates of firm age and size effects.

Table 1: Mean and median wages for all employees by size and age

	Age of employer			
	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees				
Median	216,657	215,721	215,143	215,873
Mean	236,558	235,758	233,810	229,586
Standard Deviation	150,216	262,209	144,218	130,594
Observations	713,498	564,430	843,861	2,055,623
11-49 employees				
Median	241,491	243,509	245,020	244,593
Mean	265,833	269,149	271,008	266,317
Standard Deviation	158,659	161,626	164,728	151,765
Observations	562,938	534,541	933,372	3,006,214
50-249 employees				
Median	256,508	258,951	262,862	254,782
Mean	283,725	289,681	296,477	286,602
Standard Deviation	159,388	167,088	184,248	178,582
Observations	421,312	378,354	661,398	3,173,925
250+ employees				
Median	262,808	258,131	274,921	232,700
Mean	291,495	286,521	305,151	256,221
Standard Deviation	226,572	156,494	181,147	137,522
Observations	1,831,120	1,078,998	1,550,696	10,745,356

Adjusting for firm tenure

One of the most consistent complications noted in the prior literature on the relationship between firm age and wages has been the older firms also tend to employ individuals who have longer tenure with the firm, and may generally have more experience (Brown and Medoff 2003; Heyman 2007; Ouimet and Zarutskie 2014).

Although a couple of studies have attempted to address this issue, the typical approach has been to assume that wages adjust linearly with firm tenure (e.g., Brown and Medoff 2003; Heyman 2007). That assumption has probably been necessary in most of these prior analyses given the limited amount of data they have had available. 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 quite problematic. A two-year-old firm, for example, cannot have any employees with more than two years of experience at the firm but a firm of ten years of age might have few employees with less than two years of tenure at the firm. If the returns to firm tenure decline over time, any linear adjustment for firm tenure would then underestimate the "true" firm tenure effect and probably attribute a portion of these tenure differences to firm age.

We therefore adopted a quite conservative approach to addressing this issue by examining only new hires and the wages that they earn. 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 full-time employees between the ages of 18 and 60, who had worked for a firm for at least 30 days but no more than one year. We also excluded all individuals listed as founders, employers, or entrepreneurs to ensure that we only included employees in our analysis.

Table 2 reports the median, mean and standard deviations of the wages for these recent

Table 2: Mean and median wages for recent hires by size and age

	Age of employer			
	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees				
Median	206,259	200,397	198,131	192,416
Mean	217,428	208,845	204,966	195,929
Standard Deviation	134,628	125,059	119,792	112,403
Observations	202,555	127,550	166,168	324,627
11-49 employees				
Median	228,299	225,920	225,652	219,915
Mean	246,056	241,965	240,556	231,561
Standard Deviation	142,451	134,173	131,353	129,628
Observations	141,016	117,946	190,135	497,129
50-249 employees				
Median	231,733	230,406	235,263	230,603
Mean	255,811	254,220	258,716	252,187
Standard Deviation	160,116	149,909	149,265	146,017
Observations	83,125	69,471	125,776	507,947
250+ employees				
Median	235,735	226,825	235,562	203,032
Mean	258,059	245,510	254,309	222,214
Standard Deviation	184,106	154,422	148,084	131,665
Observations	285,063	131,277	216,330	1,264,676

hires. As one would expect, since firm tenure should not necessarily vary with firm size once firm age has been held constant, reading down the columns, one continues to see a strong relationship between firm size and wages, with the largest firms generally paying new hires 12% to 20% more than the smallest ones. But the pattern for firm age changes noticeably. Looking across the rows, one sees a *negative* relationship between firm age and the average wages paid to recent hires. Within each of the size categories, the youngest firms paid up to 16% more to these hires than the oldest ones.

Adjusting 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 differ systematically on other factors from those who remain 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 any such selection from influencing our estimates, we therefore 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 had been closed. Although our sample includes the service sector, we refer to this subsample as the "plant closings" group. These individuals presumably sought employment for reasons exogenous to their individual ability or productivity (Gibbons and Katz 1991).

Table 3 reports the median, mean and standard deviations of the wages for these individuals who sought employment due to the closing of their prior employer. Note first that nearly all of the medians and means increase within this subsample as compared to the full

Table 3: Mean and median wages by size and age for employees changing jobs due to plant closing

	Age of employer			
	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees				
Median	209,023	203,061	200,201	195,417
Mean	224,490	213,564	210,422	201,567
Standard Deviation	161,755	135,605	125,876	138,612
Observations	17,543	9,122	11,533	20,964
11-49 employees				
Median	240,968	234,583	231,697	230,121
Mean	262,944	256,855	252,387	246,634
Standard Deviation	158,278	147,576	145,436	155,385
Observations	14,107	8,924	13,704	34,351
50-249 employees				
Median	247,848	246,177	259,543	246,939
Mean	281,442	279,381	294,374	277,769
Standard Deviation	192,555	183,788	184,855	195,194
Observations	8,772	6,737	11,208	41,358
250+ employees				
Median	276,480	252,537	267,504	238,656
Mean	302,685	284,343	293,010	276,626
Standard Deviation	174,215	218,012	154,155	188,550
Observations	43,149	12,675	22,702	96,912

sample presented in Table 2 that included voluntary and involuntary movers. This suggests that, on average, recent hires exhibit some adverse selection relative to those who remain in their jobs. The general patterns seen in the previous table nevertheless hold in this more restrictive sample. Within each age range, the largest firms pay 22% to 34% more to new hires than the smallest firms. Meanwhile, within each size range, the youngest firms pay as much as 16% more than the oldest firms.

Adjusting for human capital

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 (Moore 1911; Kremer 1993). We first explored the extent to which these young and small firms differed from the population as a whole 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 hire different kinds of individuals therefore remains an open question.

Table 4 reports the averages across all firms for any new hire and for new hires coming from a plant closing in the first two columns. It compares these individuals to those moving to young firms and to those moving to small firms. Interestingly, we see almost no differences between the population of new hires and those going to young firms. By contrast, we see somewhat larger differences between smaller and larger firms, with smaller employers hiring more men, and less educated and younger individuals.

Table 4: Demographics for movers to different destinations, 1990-2007

	Any destination		<9 years		<50 employees	
	All hires	Plant closing	All hires	Plant closing	All hires	Plant closing
Age	35.12	38.87	35.63	39.63	34.66	36.94
	(10.41)	(10.99)	(10.59)	(11.02)	(10.46)	(10.93)
Female	0.44	0.44	0.42	0.47	0.36	0.36
	(0.50)	(0.50)	(0.49)	(0.50)	(0.48)	(0.48)
Months of education	151.89	150.90	152.39	152.12	148.27	146.19
	(29.43)	(29.78)	(29.62)	(29.70)	(28.05)	(28.60)
Type of education						
Primary school	0.24	0.24	0.23	0.23	0.27	0.30
	(0.43)	(0.43)	(0.42)	(0.42)	(0.44)	(0.46)
High-school/gymnasium	0.10	0.08	0.09	0.07	0.09	0.08
	(0.30)	(0.27)	(0.29)	(0.26)	(0.28)	(0.27)
Vocational training	0.40	0.42	0.41	0.42	0.46	0.46
	(0.49)	(0.49)	(0.49)	(0.49)	(0.50)	(0.50)
College	0.18	0.18	0.17	0.19	0.13	0.11
	(0.38)	(0.38)	(0.37)	(0.40)	(0.33)	(0.32)
University	0.09	0.08	0.10	0.08	0.06	0.05
	(0.28)	(0.27)	(0.29)	(0.27)	(0.23)	(0.22)
Labor market experience	13.99	17.11	14.51	17.75	13.74	14.97
	(9.08)	(9.74)	(9.36)	(9.85)	(8.83)	(9.19)
Unemployment history	1.36	1.25	1.32	1.24	1.41	1.43
	(2.04)	(1.99)	(2.03)	(2.00)	(2.06)	(2.07)
Observations	5,314,599	558,422	2,716,064	364,592	2,157,161	180,450

Note: Standard deviation in parentheses.

Although few prior studies have adjusted for these differences, those that have generally had to rely on adjustments through linear regression (for an exception using propensity score matching, see Nystrom and Elvung 2014).[4] In other words, the researchers estimate a wage equation and assume that each of the relevant human capital dimensions has additive effects to the expected wage, or its logged value (e.g., Brown and Medoff 2003). That approach assumes, for example, that the returns to education and experience do not vary across men and women.

Having data on the entire population allows us to adopt a more flexible and non-parametric approach to adjusting for these factors. Rather than estimating a wage equation with linear adjustments for the effects of age, gender, education, and other factors, we instead match on these factors and include a fixed effect for each 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).

Many forms of matching exist. Perhaps the most commonly used form of matching, propensity score matching, estimates a model that uses a set of observed variables to predict the probability that a particular individual would receive "treatment" – in this case, that an employee would join a firm of a particular age and size. One then compares those receiving the treatment to a set of controls, other individuals, that had identical probabilities of being treated as a means of assessing the effects of treatment (Rosenbaum and Rubin 1983). Propensity score matching has its advantages, particularly when one has a relatively small number of cases with which to work and therefore one wishes to retain as many of them

[4]Nystrom and Elvung (2014), however, used joining a startup versus an established firm as the treatment in creating the propensity score, they therefore essentially estimated the joint effects of firm age and size.

as possible to maximize the precision of the estimates. But it also has limitations. Most notably, researchers often find it difficult to achieve balance – statistically-indistinguishable distributions on observables – between cases and controls using propensity score matching (King and Nielsen 2015). In the absence of balance, one must worry that the apparent effects of the treatment arise instead from differences between the cases and controls on other dimensions.

In order to minimize the possibility that some confounding factor accounts for the results, one would ideally match cases and controls *exactly* on all of the 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 down to the dollar. We therefore adopt 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 an extended discussion of the advantages of this approach in Iacus et al. (2012), but one of the principle ones for our purpose is that the procedure guarantees balance between the cases (the individuals working for young or small firms) and the controls (employees of larger and older firms).

Our matching procedure operates as follows, 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 find all of the observationally-equivalent individuals beginning jobs at employers that have at least 250 employees and that have been operating for at least nine years (our baseline category). We consider two individuals observationally equivalent if they have 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.[5]

Although matching accounts for differences across employees in observed characteristics, workers likely differ on a host of unobserved dimensions that also affect productivity and pay. To account for these differences, from the set of available individuals who matched exactly on gender, age, and degree, we only included the two nearest neighbors on the prior year wage distribution (the closest above and the closest below) – what an employee earned in that previous year – in the comparison set. This procedure yielded statistically indistinguishable 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 each of the 202,555 "focal" individuals that 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. Of the 202,555 focal individuals in this category, we found at least two exact matches on these dimensions for 99,976 of them. For each focal individual, we selected the exact match closest but just above the person in earnings $(t-1)$ and the exact match closest but just below in earnings $(t-1)$.

In total, we have 15 sets of matched samples (one for each cell in the age-size matrix, except for the baseline category). For each of these matched samples, we estimated the effect of being in the "treated" group (that is, not being employed by a firm in the oldest

[5]We used the one-digit version of the occupation codes for Denmark. These codes distinguish between skilled and unskilled jobs and between white collar and blue collar occupations but they do not introduce a fine-grained classification that would distinguish between industries. Below, we adjust for industry differences by introducing a vector of indicator variables for 4-digit industry codes.

and largest categories). Specifically, we estimated the following equation:

$$W_i = \beta_{as} AS_i + \gamma_j + \epsilon_i, \quad (1)$$

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 ϵ_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. In the interest of saving space, each cell in the tables below simply reports the β_{as} coefficient for the model based on the relevant matched sample.

We repeated this procedure for each of the 15 matched samples. Table 5 reports the β_{as} values from these 15 regressions. 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 range relative to established (9+ years), large (250+ employees) firms. Thus, for example, the top left cell indicates that an individual hired by a firm in the smallest, youngest group would receive roughly 11,000 Danish kroner (about $2,000) less per year than a similar individual hired by a large, established firm.[6] Reading these values down the columns, one can see a strong size effect, with the largest firms typically paying about 20,000 kroner ($3,600) more than the smallest ones. Reading across the rows, one also sees a very consistent firm age effect, with the youngest firms paying about 10,000 ($1,800) kroner more than the oldest ones. Relative to the unadjusted figures in Table 2, the gradient appears roughly the same

[6] Currency conversions made using the average exchange rate for 2013: 5.6 DKK = 1 USD.

Table 5: 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	-10,913***	-15,147***	-16,570***	-21,813***
	(405.46)	(454.65)	(406.38)	(360.96)
N	299,928	230,658	268,674	369,258
11-49 employees	8,042***	4,646***	2,599***	-2,365***
	(467.18)	(463.98)	(408.42)	(336.75)
N	249,216	221,517	288,393	423,240
50-249 employees	10,418***	9,752***	9,175***	9,996***
	(584.53)	(587.82)	(491.39)	(328.18)
N	174,360	153,327	231,549	430,185
250+ employees	13,400***	9,141***	4,518***	
	(517.79)	(538.16)	(425.99)	
N	266,571	209,679	270,633	

* p <0.1, ** p <0.05, *** p <0.01. Robust standard errors in parentheses.

Note: All regressions are fixed effects OLS regressions, fixed on the matched trio of one individuals from the treatment cell exactly matched with two individuals from the baseline cell. Baseline cell is firms aged 9+ with 250+ employees. Matched on age, gender, education, prior occupational level, and prior wage.

Table 6: Regression matrix: Movers from plant closings matched on age, gender, education, prior job, and prior earnings

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees	-17,552***	-27,458***	-28,097***	-34,089***
	(2736.03)	(2074.68)	(2133.74)	(1575.51)
N	15,942	10,938	12,408	17,733
11-49 employees	3,795**	-9,978***	-9,483***	-14,559***
	(1910.11)	(2481.13)	(2029.39)	(1444.43)
N	15,405	11,247	14,955	23,499
50-249 employees	8,085***	5,831**	2,094	-1,885
	(2271.82)	(2389.24)	(2704.86)	(1577.23)
N	11,319	9,237	13,119	25,926
250+ employees	2,957*	126	-9,088***	
	(1789.77)	(2194.51)	(2027.31)	
N	18,117	11,850	14,313	

* p <0.1, ** p <0.05, *** p <0.01. Robust standard errors in parentheses.

Note: All regressions are fixed effects OLS regressions, fixed on the matched trio of one individuals from the treatment cell exactly matched with two individuals from the baseline cell. Baseline cell is firms aged 9+ with 250+ employees. Matched on age, gender, education, prior occupational level, and prior wage.

for the relationship between firm age and wages but somewhat less steep for the relationship between firm size and wages.

Further adjusting for selection into mobility. As noted above, one might nonetheless worry that job changers differ systematically from the population of employees. We therefore repeated our case-control construction and estimation, restricting the population to individuals who had left their prior employers because the plant or business location at which they had been working closed (see Table 6). Although the general patterns across the age-size cells appears consistent with the estimates based on the full population of movers, the effect sizes appear even larger within this subsample. Reading down the columns, the largest firms pay between 19,000 kroner ($3,400) and 34,000 kroner ($6,100) more than the smallest ones. Reading across the rows, the youngest firms pay about 17,000 kroner ($3,000) more than the oldest ones in the two smaller size categories but only 3,000 kroner ($540) more in the largest size category.

Adjusting for industry

Although these adjustments address a large number 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. What appears a firm age or firm size effect therefore might actually capture an industry effect on wages.

To address this issue, we reestimated the models above adjustments for four-digit industries:

$$W_i = \beta_{as} AS_i + \eta_i + \gamma_j + \epsilon_i, \tag{2}$$

where η_i represents a vector of four-digit industry dummies.[7] Table 7 reports the equivalent of Table 5, adjusting for industry effects. Even after adjusting for industry effects, larger firms still pay more than smaller firms and younger firms continue to compensate observationally-equivalent individuals better than older firms. Interestingly, while industry differences appear to account for as much as half of the relationship between firm age and wages, compositional differences across industries in the joint firm age-size distribution appear, if anything, to mask the relationship between firm size and wages.

[7]Alternatively, one could adjust for industry by matching the focal individuals with the controls on their industry of employment. That approach has the advantage of not only adjusting for average wages across industries but also for differences in the returns to human capital characteristics across industries. It nevertheless has the disadvantage of substantially increasing the difficulty of finding matches for the focal individuals. Despite the large number of cases from which we have to draw, such precise matching leaves us with only a few cases on which to estimate our β's and therefore little confidence in their precision or representativeness.

Table 7: Regression matrix: New hires matched on age, gender, education, prior job, and prior earnings (controlling for 4-digit industries)

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees	-20,760***	-23,811***	-26,302***	-25,738***
	(754.30)	(912.08)	(724.64)	(633.51)
N	299,928	230,658	268,674	369,258
11-49 employees	-3,830***	-7,957***	-10,465***	-10,398***
	(675.95)	(707.15)	(629.04)	(492.08)
N	249,216	221,517	288,393	423,240
50-249 employees	5,421***	2,272***	-1,292*	-4,494***
	(851.07)	(800.00)	(724.14)	(446.69)
N	174,356	153,327	231,549	430,185
250+ employees	8,055***	4,664***	421	
	(599.32)	(717.57)	(588.04)	
N	266,571	209,679	270,633	

* p <0.1, ** p <0.05, *** p <0.01. Robust standard errors in parentheses.

Note: All regressions are fixed effects OLS regressions, fixed on the matched trio of one individuals from the treatment cell exactly matched with two individuals from the baseline cell. Baseline cell is firms aged 9+ with 250+ employees. All regressions include 4-digit industry dummies. The sample is matched on age, gender, education, prior occupational level, and prior wage.

Table 8: Regression matrix: Movers from plant closings matched on age, gender, education, prior job, and prior earnings (controlling for 4-digit industries)

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees	-28,638*** (3792.83)	-35,125*** (4053.03)	-39,698*** (4164.93)	-39,211*** (2637.56)
N	15,984	11,022	12,312	17,700
11-49 employees	-4,034 (2738.52)	-11,059*** (2778.87)	-20,996*** (2888.23)	-16,728*** (2521.20)
N	15,363	11,307	14,877	23,286
50-249 employees	4,428 (3125.22)	-3,583 (3605.22)	-2,925 (4203.53)	-11,120*** (2087.57)
N	11,282	9,237	13,128	25,992
250+ employees	9,845*** (2159.96)	6,729** (3401.69)	-8,094*** (2827.04)	
N	18,150	11,769	14,385	

* p <0.1, ** p <0.05, *** p <0.01. Robust standard errors in parentheses.

Note: All regressions are fixed effects OLS regressions, fixed on the matched trio of one individuals from the treatment cell exactly matched with two individuals from the baseline cell. Baseline cell is firms aged 9+ with 250+ employees. All regressions include 4-digit industry dummies. The sample is matched on age, gender, education, prior occupational level, and prior wage.

Table 8 runs an equivalent set of models on the subsample of individuals who moved due to the closure of their prior workplace. These patterns largely hold when restricting the sample to those moving because their prior places of employment had closed.

Discussion

Do startups create good jobs? Our answer to this question seems somewhat mixed: We explored the relationship between the amount that firms paid and firm age and size among the population of Danish employers and employees and present evidence of both a firm size effect and a firm age effect on the wages of new hires. We found that 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, we found that young firms actually paid recent hires *more* than older firms of similar size. Our empirical strategy recognizes that employee characteristics affect wages and adjusts for the fact that young and small firms hire different kinds of people from larger and more established ones. Indeed, our pattern of results held after adjusting for differences in the characteristics of the people they hired and for industry differences in pay. But firm size had larger effects than firm age. Hence, to the extent that startups begin both young and small, they do tend to pay less than large, established firms. Those startups that grow rapidly, however, may overcome this size effect fast enough that they actually pay a premium relative to more established employers.

We believe that this research has two central contributions. The first has been more methodological. One of the difficulties in assessing job quality is that one cannot really 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 extant research has been aware of this issue, the typical approach to adjusting for these 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. This method therefore does account for the fact that certain characteristics may act as complements or substitutes in the determination of wages.

The increasing availability of longitudinal registry data, however, opens the door for alternative approaches. The Danish registry data, for example, include more than 30 million person-years of information. Our analytic strategy demonstrates two ways in which researchers can take advantage of such large datasets. One involves the creation of observationally-equivalent sets through matching. Instead of adjusting for observed characteristics through regression, we instead use matching to create sets of cases and controls nearly identical on the observed dimensions and allow 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.

The other involves the use of specific subsets to address other potential confounds. For example, we limit our analysis to new hires to eliminate the confounding effect of employee tenure on wages, allowing us to isolate better the relationship between firm age and wages. We also address potential concerns that people who change jobs are somehow different than those who do not by focusing our analysis on those who had no choice but to move—those employed at employers that closed. We are hardly the first to use this technique: Deaths, for example, have been used to identify peer effects and the importance of founders and managers (e.g., Azoulay et al. 2010; Gjerløv-Juel 2014); plant closings have been exploited as a source of exogenous job changes (e.g., Gibbons and Katz 1991); and the birth of daughters has been used as an exogenous shift in the gender beliefs of managers (e.g., Dahl et al. 2012). But it's value as a methodological approach rises rapidly in large datasets: These subsets often focus on a small fraction of the overall data, but in the very large numbers that characterize registry data one still has sufficient statistical power to generate relatively precise estimates.

The second contribution has been to call attention to the importance of the quality of the jobs created by startups. Although a substantial number of studies have examined the extent to which startups create more jobs than they destroy (e.g., Audretsch 2002; Haltiwanger et al. 2013; de Wit and de Kok 2014), relatively little attention has been given to whether those jobs are better than the ones that they replace.

Although our results provide some initial insight into this question, they represent more of a first step in a research agenda than a definitive answer to whether startups create good jobs. Consider some of the closely-related questions that remain open: Although young firms might pay their employees a premium in the first year, how do these effects evolve over time? 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 training and in the capital improvements that would enhance worker productivity over time.

How does job instability affect the employees of startups? One reason that startups might pay their employees more stems from the fact that these workers face a higher 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. But what if losing a job due to firm failure proves more problematic than expected, perhaps because former employees end up needing to find a job during a period of contraction or because would-be future employers view any spell of unemployment unfavorably (regardless of the reason). The negative effects of this job instability may then more than offset the higher wages paid by these smaller firms, meaning that startup jobs may not be as attractive as they would otherwise appear.

Do these patterns vary across space and time? One could easily imagine that the quality of jobs created by startups might vary across regions or over time. Better infrastructure and supporting services, for example, might allow small, young firms to operate closer to the same level of productivity as larger and more established rivals. Consistent with this idea, Hollister (2004) reports that the wage gap between large and small firms has been shrinking over time in the United States.

These unanswered questions point to a promising research agenda that brings together ideas and techniques from labor economics and organizational sociology to study the relationship between entrepreneurial activity, economic development, and inequality. Entrepreneurship has been and will continue to be an important driver of economic vitality, 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

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. 2002. The dynamic role of small firms: Evidence from the U.S. *Small Business Economics* **18** 13–40.

Audretsch, David B. 2007. *The Entrepreneurial Society*. Oxford University Press, New York.

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.

Azoulay, Pierre, Joshua S. Graff Zivin, Jialan Wang. 2010. Superstar extinction. *Quarterly Journal of Economics* **125**(2) 549–589.

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

Birch, David L. 1987. *Job Creation in America: How Our Smallest Companies Put the Most People to Work*. Free Press, New York.

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. 1989. The employer size wage effect. *Journal of Political Economy* **97** 1027–1059.

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

Dahl, Michael S., Cristian Dezso, David Gaddis Ross. 2012. Fatherhood and managerial style: How a male CEO's children affect the wages of his employees. *Administrative Science Quarterly* **57**(4) 669–693.

Davis, Steven J., John Haltiwanger. 1991. Wage dispersion between and within U.S. manufacturing plants, 1963-1986. *Brookings Papers on Economic Activity: Microeconomics* (1) 115–180.

Davis, Steven J., John Haltiwanger, Scott Schuh. 1996. Small business and job creation: Dissecting the myth and reassessing the facts. *Small Business Economics* **8** 297–315.

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

Doeringer, Peter, Michael Piore. 1971. *Internal Labor Markets and Manpower Analysis*. D.C. Heath, Lexington, MA.

Gibbons, Robert, Lawrence F. Katz. 1991. Layoffs and lemons. *Journal of Labor Economics* **9**(4) 351–380.

Gjerløv-Juel, Pernille. 2014. Who loses a leader without losing ground? unexpected deaths in top management teams and firm performance. Paper presented at DRUID Summer Conference.

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

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.

Hollister, Matissa N. 2004. Does firm size matter anymore? the new economy and firm size wage effects. *American Sociological Review* **69** 659–676.

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.

Kremer, Michael. 1993. The o ring theory of economic development. *Quarterly Journal of Economics* **108** 551–575.

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

Litwin, Adam Seth, Phillip H. Phan. 2013. Quality over quantity: Reexamining the link between entrepreneurship and job creation. *Industrial & Labor Relations Review* **66** 833–873.

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

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

Moore, Henry L. 1911. *Laws of Wages*. Macmillan, New York.

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.

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.

Rosenbaum, Paul R., Donald B. Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**(1) 41–55.

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.

Syverson, Chad. 2011. What determines productivity? *Journal of Economic Literature* **49**(2) 326–365.

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.

Topel, Robert H., Michael P. Ward. 1992. Job mobility and the careers of young men. *Quarterly Journal of Economics* **107**(2) 439–479.

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

Troske, Kenneth R. 1999. Evidence on the employer size-wage premium from worker-establishment matched data. *Review of Economics and Statistics* **81**(1) 15–26.

Winter-Ebmer, Rudolf, Josef Zweimuller. 1999. Firm-size wage differentials in switzerland: Evidence from job-changers. *American Economic Review* **89**(2) 89–93.