Do startups create good jobs?*

M. Diane Burton† Cornell University

Michael S. Dahl‡ Aarhus University

Olav Sorenson§ Yale University

March 15, 2016

DRAFT MANUSCRIPT
PLEASE DO NOT CITE OR QUOTE WITHOUT PERMISSION

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 pay a wage premium.

*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. The usual disclaimer applies.

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

‡Department of Management, Bartholins Allé 10, DK-8000 Aarhus, msd@mgmt.au.dk

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

Introduction

Entrepreneurship and the idea that entrepreneurs 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) 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 appear to exhibit quite similar patterns (Ayyagari et al. 2014; de Wit and de Kok 2014; Lawless 2014; Anyadike-Danes et al. 2015), lending credence to the idea that entrepreneurial firms are indeed an engine of job creation.[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

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

of higher paying jobs at incumbent firms with lower paying ones at startups, then a simple examination of the numbers of jobs created, even net of jobs lost, may overstate the value of entrepreneurial activity to the economy and society. 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 suggested 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 findings, one might expect that startups, being both young and small, would pay substantially less than the more established firms they replace.

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 attract a somewhat different set of employees than established firms. Ouimet and Zarutskie (2014), for example, reported that the average employee of an entrepreneurial firm is younger, less educated, and less experienced than 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 them means that any apparent effects of firm age and size may instead reflect differences in workforce composition as opposed to actual wage differences across types of firms. Two other individual-level 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.

Compositional issues also arise at the industry level. New firm entries may disproportionately appear in new industries that have different pay practices than established industries. These industries may also offer faster growth prospects with their own implications for wages.

We address all of 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 characteristics of their workforces, mobility patterns, and industry effects. Our large population size – over 20 million employee-years – allows us to contribute to and extend existing research on firm age and wages in at least six 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 address the potential selection effects associated with who leaves their prior employer, we estimated effects within 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 used (coarsened-exact) matching to focus on effects within sets of observationally-equivalent individuals; (iv) to account in part for unobserved differences in productivity, we further matched individuals to their nearest-neighbors in terms of wages in their previous jobs; (v) to address industry-level effects, we adjusted for fine-grained (4-digit) industry differentials in wages; and (vi) to account for additional firm-level heterogeneity and the potential sorting of high quality employees into faster growing firms on the basis of characteristics not observed by the researcher, we also accounted for the future growth rates of the employers. We further demonstrated the robustness of the patterns among a subsample of labor market entrants.

Some of these adjustments mattered more than others. Focusing on the newly hired roughly doubled the magnitude of the wage gradient associated with firm age, though it had almost no effect on the wage gradient with respect to firm size. By contrast, focusing on those moving because their prior employer closed, led to roughly 50% larger effects of firm size, though it did not change the effects of firm age appreciably. Across both the age and size gradients, differences in the characteristics of employees appeared to account for about 40% to 60% of the observed differences in average firm wages. Even after all of these adjustments, we found a small firm wage penalty, consistent with prior research. 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, rapidly growing young firms appeared to pay a substantial wage premium over large, 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 more established employers.

Firm age, firm size, and wages

Despite 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 (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, 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.

[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).

(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, 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 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, many questions remain. The only large-scale studies that consider both individual characteristics and industry effects both rely on data from Sweden. Do other countries show similar patterns? 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. The returns to experience, for example, might vary across men and women and with levels of education (e.g., Polachek 1981; Goldin 2006). How might the patterns change if one adopted a more flexible approach to adjusting for individual characteristics? Finally, does firm age really have effects independent from employee tenure. Nystrom and Elvung (2014) did address the issue by focusing on those with no experience at any job, those first entering the labor market. But these early job matches 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 relationships between firm age, firm size, and wages, we examined Danish registry data covering every employee in the country using the Integrated Database for Labor Market Research (commonly referred to as the IDA database). For our analyses 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 have not yet begun to shift their employment choices in anticipation of retirement.

The IDA database starts in 1980; however, we consider only the post-1991 period after a series of regulatory reforms had dismantled the centralized wage-setting process and left only a minimum wage (Madsen et al. 2001). This allowed firms much more flexibility in wage setting. Indeed, Denmark has the least restrictive labor market policies in Europe (Bingley and Westergård-Nielsen 2003; Sørensen and Sorenson 2007), usefully allowing comparison to market-oriented economies such as Canada, the United Kingdom, and the United States. 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 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 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 entering the Danish market as these firms are quite different from new startups.

We were also careful to exclude established firms that might appear to be new young firms by virtue of a change in ownership or legal form, such as spin-outs and privatizations. We identified these firms by flagging new firms where a large proportion of the employees worked

[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

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees				
Mean	239,938	239,224	235,315	225,448
Median	216,996	216,217	214,478	210,107
Standard Deviation	160,620	325,958	153,750	133,422
Observations	414,349	325,888	474,790	1,363,716
Number of firms	110,634	77,100	66,470	72,007
11-49 employees				
Mean	278,311	282,161	281,200	265,256
Median	246,560	249,440	249,078	239,992
Standard Deviation	185,361	188,510	184,084	160,083
Observations	208,350	221,110	405,741	1,861,437
Number of firms	8,720	8,445	9,930	16,293
50-249 employees				
Mean	291,341	292,146	297,941	284,591
Median	253,899	252,599	257,594	251,433
Standard Deviation	191,051	194,368	196,895	182,429
Observations	122,053	127,559	261,668	2,265,351
Number of firms	1,255	1,268	1,658	4,061
250+ employees				
Mean	311,357	309,104	305,273	289,589
Median	283,171	279,605	270,175	255,302
Standard Deviation	167,542	172,877	191,917	174,679
Observations	167,406	167,064	326,370	3,917,366
Number of firms	302	289	307	940

together in the prior year for an employer in the same industry and geography, but with a different firm identification code. We recognize that this procedure may exclude some new start-ups that are founded by groups of people voluntarily leaving the same employer. But the risk of mistakenly treating well-established entities as start-ups seemed more problematic than the exclusion of some employee spin-offs.

Table 1 reports the median, mean and standard deviation of the wages for all employees in each of these size and age categories across the entire period from 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 for the study population of full-time workers between the ages of 18 and 60 during the period from 1991-2006. Looking down the columns, one can begin to see a clear size gradient. Within each age range, larger firms pay more than smaller ones. The smallest employers – those with 1-10 employees – pay their employees 18-23% less than those with 250 or more employees. Looking across the rows, patterns are more difficult to detect. However, within each size category, the oldest firms pay, on average, about 3%-7% less than the youngest ones.

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 that 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 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 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 excluded all individuals listed as founders, employers, or entrepreneurs, as their compensation may involve equity as well as wages. These restrictions reduce our sample size, but still leave 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. The average wages in Table 2 are consistently lower than those in Table 1 which implies that, as one would expect, there are returns to firm tenure. 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 sees a much stronger *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 and the youngest firms paid 9% to 13% more to new 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 differed systematically

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

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees				
Mean	221,932	212,674	207,906	199,925
Median	209,496	202,911	199,888	193,638
Standard Deviation	140,783	128,440	123,108	112,021
Observations	142,098	84,950	105,506	216,144
Number of firms	63,030	40,195	38,360	46,720
11-49 employees				
Mean	259,284	252,437	249,179	235,099
Median	237,420	232,509	230,577	220,225
Standard Deviation	161,576	147,481	140,816	136,324
Observations	70738	59604	95501	306166
Number of firms	7,785	7,480	8,876	14,556
50-249 employees				
Mean	274,507	264,220	266,412	253,423
Median	242,250	234,108	238,448	230,167
Standard Deviation	190,235	172,847	155,758	146,250
Observations	33856	30255	58874	363947
Number of firms	1,118	1,155	1,530	3,818
250+ employees				
Mean	298,975	276,950	267,888	258,301
Median	263,854	250,740	243,772	232,402
Standard Deviation	187,801	141,998	148,116	152,463
Observations	42,794	28,501	62,318	496,044
Number of firms	275	269	298	889

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 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 closed. Although our sample includes the service sector, we refer to this subsample as the "plant closings" group. This again reduces our sample size dramatically from more than 3 million observations to fewer than 220,000. However, this restricted sample allows us to examine the subset of individuals who 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 all of the medians and means increase within this subsample as compared to the full sample of both voluntary and involuntary movers in Table 2. The median difference ranges from 3,971 to 47,672 Danish kroner (DKK) and is, on average, 18,022 DKK, suggesting that who moves exhibits some adverse selection relative to who remains in their job. Within this subsample, the patterns appear similar but the magnitudes change, particularly with regard to the firm size gradient. The youngest employers continue to pay less than the oldest employers, here by 4% to 11%. But the differences in pay between the smallest firms and the largest firms expand to 26% to 33%. Adverse selection in who moves appears more acute among the larger and older firms.

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

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees				
Mean	227,681	216,721	211,877	204,102
Median	211,214	204,931	201,193	195,032
Standard Deviation	168,340	140,978	125,260	146,087
Observations	13,245	6,354	7,694	14,336
Number of firms	10,603	5,421	6,365	10,909
11-49 employees				
Mean	272,540	267,503	264,116	246,621
Median	246,890	241,729	238,519	226,608
Standard Deviation	173,182	159,993	158,393	168,330
Observations	8,506	5,045	7,364	21,467
Number of firms	3,307	2,617	3,645	7,825
50-249 employees				
Mean	288,987	285,392	304,107	278,046
Median	248,929	243,872	262,227	245,134
Standard Deviation	198,038	197,953	202,753	198,380
Observations	4,810	3,475	5,718	30,718
Number of firms	680	668	1,063	3,085
250+ employees				
Mean	322,941	293,463	297,387	305,973
Median	278,377	262,485	270,671	261,039
Standard Deviation	206,841	155,731	150,507	201,852
Observations	11,756	4,115	9,218	59,871
Number of firms	223	207	250	789

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 in fact hire different kinds of individuals therefore remains an open question.

Table 4 reports the sample demographic characteristics for the full sample of new hires and two relevant subsets. The first two columns present averages across all firms for any new hire and for new hires coming from a plant closing. The next two columns compare these individuals to those moving to young firms. The final two columns compare these individuals to those moving to small firms. Interestingly, we see very few differences between the overall population of new hires in the first and second columns and those going to young firms in the third and fourth columns. The average age among all hires is roughly 34 years old and it increases by 3 years to 37 years among the subsample of movers due to plant closings. Approximately one third of the employees are female. We see somewhat larger differences between smaller and larger firms when we compare the first and second columns with the fifth and sixth columns. Here we see that smaller employers, particularly in the plant closing sub-samples, are hiring people who are less educated, have less labor market experience, and have had more months of unemployment.

Table 4: 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)
Prior wage (1000s)	183.52	218.51	184.87	231.96	171.07	188.95
	(1,062.4)	(2,447.3)	(1,739.2)	(4,699.1)	(1,286.9)	(3,881.5)
Observations	2,197,296	213,692	492,796	57,306	1,080,707	84,011

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 estimated a wage equation, effectively assuming that each of the relevant human capital dimensions had additive effects on the expected wage, or its logged value (e.g., Brown and Medoff 2003). But that approach assumes, for example, that the returns to education and experience do not vary across men and women. Substantial research nevertheless suggests that men and women may sort into occupations that have different returns to experience (e.g., Polachek 1981; Benson 2014). These population average adjustments theremore may not capture well the actual differences in human capital across individuals.

Having data on the entire population allowed us to adopt a more flexible and non-parametric approach to adjusting for these 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).

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

[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.

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 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.

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 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 of the observationally-equivalent individuals in our baseline category of large, established firms (those beginning jobs at employers that have 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.[5]

Although matching accounts for differences across employees on some of the most important factors influencing wages, workers likely 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 can therefore use information on the prior wages of individuals to account for these differences. From the set of available individuals who matched exactly on gender, age, and degree, we therefore 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. This procedure yielded statistically indistinguishable average wages across all sets of cases and controls.

We matched with replacement, meaning a control could serve as a match for multiple focal individuals. Matching with replacement can often reduce the level of bias as it ensures closer matches between the 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 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. For each focal individual, we selected the exact match closest but just above the person in earnings

[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.

$(t-1)$ 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.[6]

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 categories). Specifically, we estimated the following equation:

$$W_i = \beta_{as} AS_i + \gamma_j + \epsilon_i, \tag{1}$$

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 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. We repeated this procedure for each of the 15 matched samples.

Table 5 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 range relative to the pay offered in an

[6] In total, we have 15 sets of matched samples (one for each cell in the age-size matrix, except for the baseline category) and obtain 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, which had individuals most dissimilar from those working for the largest, most established firms.

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	-25,484***	-29,668***	-29,920***	-30,194***
	(492.55)	(538.32)	(483.20)	(419.39)
Triads	135,520	81,548	101,582	208,930
11-49 employees	926	-5,613***	-7,915***	-12,108***
	(621.64)	(652.47)	(534.90)	(357.40)
Triads	68,300	57,598	92,550	297,818
50-249 employees	12,600***	6,345***	2,873***	-2,445***
	(1,032.90)	(961.28)	(627.39)	(334.27)
Triads	32,744	29,322	57,147	354,571
250+ employees	26,529***	12,765***	4,993***	
	(1,012.33)	(939.63)	(629.39)	
Triads	41,319	27,603	60,507	

* p <0.1, ** p <0.05, *** 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.

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 roughly 25 thousand Danish kroner (about \$3,850) less per year than a similar individual hired by a large, established firm.[7] Reading across the first row we see a negative wage coefficient for all four firm age categories, revealing sizable wage penalty for employees hired by the smallest firms, which seems to worsen as firm age increases. Employees in a very small established firm (9+ years) earn approximately 30,000 DKK (about \$4,500) less than those in the baseline category, the largest established firms.

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 go up as firm size increases. Moreover, the wage penalty becomes a wage premium as firms become larger within a given age category. Reading across rows, the youngest firms within each of size category paid the highest wages.

Interestingly, however, some startups paid higher wages that the largest, most established firms. Notably, the younger three columns of the two largest size categories (rows) all have positive wage coefficients. Start-ups, particularly those that grow rapidly, therefore would appear to create high-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 that of large, established firms.

[7] Currency conversions made using the average exchange rate for 1991–2006: 6.5 DKK = 1 USD.

Adjusting for industry

Although the adjustments made up until this point 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. 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 reestimated the models above with 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.[8]

Table 6 reports the equivalent of Table 5, adjusting for industry effects. Surprisingly, the addition of more than 500 industry intercepts have relatively little impact on average wages. Even after adjusting for industry effects, younger firms continue to compensate observationally-equivalent individuals better than older firms. Firms with 50 or more employees consistently pay more than firms with fewer than 50 employees. Given the similarities across Tables 6 and 5, it seems unlikely that pay practices across industries, or industry-level

[8]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 6: 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	-24,369***	-27,323***	-28,080***	-28,905***
	(815.17)	(867.63)	(728.98)	(680.07)
Triads	135,520	81,548	101,582	208,930
11-49 employees	-2,335***	-11,583***	-12,744***	-13,735***
	(834.12)	(892.97)	(745.24)	(483.35)
Triads	68,300	57,598	92,550	297,818
50-249 employees	15,082***	6,153***	-1,036	-4,314***
	(1,552.63)	(1,524.38)	(814.89)	(400.62)
Triads	32,744	29,322	57,147	354,571
250+ employees	23,322***	9,306***	2,257**	
	(1,256.31)	(1,234.16)	(878.58)	
Triads	41,319	27,603	60,507	

* p <0.1, ** p <0.05, *** 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 and for each 4-digit industry.

differences in firm characteristics, account for any substantial portion of the observed wage differences across firm age or firm size.

Addressing the endogeneity of moves. As noted above, to disentangle the effects of firm age from those of employee tenure, we have focused our attention only on new hires – "movers" – the set of people who had changed jobs. But, one might worry that job changers differ systematically from the population of employees – particularly when we have no way to distinguish between voluntary and involuntary movers. We therefore repeated our case-control construction and estimation including controls for industry, while restricting the population to individuals who had left their prior employers because the plant or business location at which they had been working closed (involuntary movers).

Table 7: 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	-35,792***	-35,139***	-41,288***	-40,494***
	(3,523.08)	(4,729.80)	(4,027.53)	(3,586.24)
Triads	8,772	4,365	5,278	10,231
11-49 employees	-3,892	-16,742***	-23,611***	-21,339***
	(3,160.99)	(4,138.80)	(3,462.14)	(2,533.31)
Triads	6,257	3,504	5,309	15,737
50-249 employees	15,415***	-4,271	-4,016	-5,884***
	(5,721.84)	(7,001.13)	(5,777.97)	(1,759.21)
Triads	3,614	2,653	4,071	22,670
250+ employees	19,491***	5,072	-4,657	
	(2,992.10)	(6,051.95)	(3,461.82)	
Triads	9,512	3,118	6,582	

* p <0.1, ** p <0.05, *** 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 and for each 4-digit industry.

Table 7 reports the results of these models. Although these estimates are based on much smaller samples, the general patterns of wage penalties and premia with respect to firm age and size remain. But some notable differences appear. As one might expect from the descriptive statistics for this subsample (Table 3), the magnitudes of the differences in pay between the smallest and the largest firms become more amplified. This effect appears here primarily in the first row of coefficients, average wages associated with being employed in the smallest firm category, where the wage penalty for being employed at one of these firms has nearly doubled relative to Table 5.

Is it age and size, or growth?

The fact that the tables consistently reveal wage penalties for smaller and older firms and wage premia for larger and younger firms suggests that the real driver of differences in wages across firms might stem from their growth prospects rather than from their age or size. In order to reach the largest size categories in their first eight years, firms must generally grow quite quickly. Similarly, any firm that has remained under 50 employees for a decade or more has probably grown relatively slowly.

These growth differences might stem from a 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.

Moreover, if individuals can accurately evaluate the growth prospects of their potential employers and if potential employers can similarly distinguish between the more and less productive job applicants, one might then 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.

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). We do not have five-year forward projections for all firms, most notably because many firms fail. We excluded any firms without $t+5$ data from the analysis. To allow for a very flexible relationship between firm growth and wages, we used these future growth rates 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 adjust for any wage differentials associated with being in a particular firm growth decile.

Table 8 reports the results from these models, replicating our baseline table, Table 5, but including the vector of growth decile indicators. 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 20% to 30% – after adjusting for prospective firm growth. Interestingly, nearly all of this adjustment stems from shrinkage in the sizes of the wage penalties associated with the smallest firms. Prospective employees may find it easier to predict which firms will not grow than to guess which of the many firms attempting to grow large will succeed.

We would note that these results seem quite consistent with those of Gibson and Stillman (2009). Using rich and detailed measures of specific worker skills, they found little evidence for the idea that sorting of better employees into larger firms could account for the firm-size

Table 8: Regression matrix: New hires matched on age, gender, education, prior job, and prior earnings and controlling for 4-digit industries and future firm growth rates

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees	-22,345***	-25,734***	-26,549***	-26,220***
	(789.37)	(881.12)	(785.28)	(627.96)
Triads	88,702	53,548	68,565	150,915
11-49 employees	-3,757***	-11,388***	-11,923***	-14,017***
	(908.76)	(922.42)	(764.55)	(452.22)
Triads	47,578	37,901	61,536	208,825
50-249 employees	10,199***	1,326	-3,389***	-4,697***
	(1,183.28)	(1,180.97)	(838.06)	(400.05)
Triads	24,306	20,692	39,165	249,316
250+ employees	19,376***	15,680***	7,907***	
	(1,342.93)	(1,516.13)	(848.16)	
Triads	28,866	17,699	41,224	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. 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 and for firm growth deciles.

wage effect. Our results suggest that sorting also probably does not account for the firm-age wage effect.

Entrants to the labor market

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. While focusing on involuntarily movers

– those working at plants that closed – partially accounts for these differences, experience on the labor market may allow workers to signal better their quality.

We therefore estimated the wage premia and penalties again using only new entrants to the labor market (for a similar approach using Swedish data, see Nystrom and Elvung 2014). Firms hiring new 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. The construction of these samples and estimates mimics those for Table 5 in every aspect except for one: Because these individuals have not had jobs, we could not include nearest-neighbor matching on prior wages.

Table 9 reports the coefficients for labor market entrants. Overall, the patterns of the results nearly perfectly parallel – though somewhat smaller in scale – those observed for individuals with prior experience in the labor market. We should note, however, that these labor market entrants have lower average wages. In percentage terms, the gradients in wages associated with firm age and with firm size are therefore nearly identical in magnitude to those observed above.

Discussion

Do startups create good jobs? 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

Table 9: Regression matrix: New labor market entrants matched on age, gender, and education controlling for 4-digit industries

	1-2 years	3-4 years	5-8 years	9+ years
1-10 employees	-16,589***	-20,852***	-17,717***	-21,507***
	(2,260.01)	(2,268.50)	(2,058.91)	(1,556.19)
Triads	8,804	6,080	9,021	24,290
11-49 employees	-1,466	-6,406***	-10,918***	-9,776***
	(2,652.38)	(2,212.26)	(1,573.19)	(1,042.30)
Triads	4,307	4,343	7,691	28,042
50-249 employees	283	1409	-2,014	-4,039***
	(3,461.23)	(3,505.68)	(1,971.57)	(812.03)
Triads	1,983	2,186	4,800	31,370
250+ employees	7,651*	-8,307*	1,270	
	(3,924.66)	(4,941.56)	(2,400.57)	
Triads	2,162	2,003	5,407	

* p <0.1, ** p <0.05, *** 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.

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 become large quickly overcome the size effect fast enough that they 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 under four years of age.

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. But this explanation seems somewhat inconsistent with the observed patterns of premia. The largest differentials in wages between the youngest and oldest firms appeared in the largest size categories, which presumably had lower odds of failing.

Even if they do represent compensating differentials, one might reasonably ask whether they are large enough. Not only is the constant risk of job loss due to firm failure high among these firms but it probably rises during periods of economic contraction, precisely the times during which laid-off employees would find it most difficult to secure another job.

These open questions point to the fact that, though 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 other 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 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, 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.

In addition to establishing a set of empirical facts about the jobs being created by startups, we believe that the research also offers a methodological contribution. 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.

The increasing availability of longitudinal registry data, however, opens the door for alternative approaches. The Danish registry data, for example, include more than 20 million person-years of information. Instead of adjusting for observed characteristics through regression, we instead 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.

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.

Benson, Alan. 2014. Re-thinking the two-body problem: The segregation of women into geographically-flexible occupations. *Demography* **51**(5) 1619–1639.

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

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.

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

Goldin, Claudia. 2006. The quiet revolution that transformed women's employment, education, and the family. *American Economic Review* **96**(2) 1–21.

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.

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.

Polachek, Solomon. 1981. Occupational self-selection: A human capital approach to sex differences in occupational structure. *Review of Economics and Statistics* **58** 60–69.

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.