Do startup employees earn more in the long run?*

Olav Sorenson[†]
Yale University

Michael S. Dahl[‡]
Aarhus University

Rodrigo Canales[§]
Yale University

M. Diane Burton[¶]
Cornell University

October 30, 2019

Abstract: Evaluating the attractiveness of startup employment requires an understanding of both what startups pay and the implications of these jobs for earnings trajectories. Analyzing Danish registry data, we find that employees hired by startups earn roughly 17% less over the next ten years than those hired by large, established firms. About half of this earnings differential stems from sorting—from the fact that startup employees have less human capital. Long-term earnings also vary depending on when individuals are hired. While the earliest employees of startups suffer an earnings penalty, those hired by already-successful startups earn a small premium. Two factors appear to account for the earnings penalties for the early employees: Startups fail at high rates, creating costly spells of unemployment for their (former) employees. Job mobility patterns also diverge: After being employed by a small startup, individuals rarely return to the large employers that pay more.

*All authors contributed substantially to the paper. Authorship order simply represents a reverse alphabetical assignment. We thank Yale University and the Aarhus University Research Foundation (Grant no. AUFF-E2017-9-23) for generous financial support and Paul Gertler, Olenka Kacperczyk, Petra Moser, Jesper Sørensen, Kim Weeden, Rebecca Zarutskie, three anonymous referees, and the participants of seminars at Boston University, Columbia, Copenhagen Business School, Duke, Gothenberg University, Harvard Business School, INSEAD, McGill, the NBER, Oxford, Stanford, UC Berkeley, UCLA, University of Mannheim, University of Michigan, and Yale for their comments on earlier versions of this paper.

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

[‡]Department of Management, Fuglesangs Allé 104, DK-8210 Aarhus V, msd@mgmt.au.dk

[§]165 Whitney Ave, P.O. Box 208200, New Haven, CT 06520, rodrigo.canales@yale.edu

[¶]ILR School, 170 Ives Hall, Ithaca, NY 14853-7601, burton@cornell.edu

Introduction

Entrepreneurship has captured the imagination. Policymakers and pundits alike see startups as a solution to a variety of economic problems, from slowing economic growth to high levels of unemployment. Legislation favoring fledgling firms has followed. For example, in introducing the JOBS Act, President Obama asserted that "...new businesses account for almost every new job that's created in America" (Liberto 2012). Recent academic research has supported these notions, arguing that startups account for the majority of all net job creation (Haltiwanger et al. 2013, de Wit and de Kok 2014, Lawless 2014).

Individuals worldwide, meanwhile, increasingly see startups as desirable employers and entrepreneurship as an attractive career option (Bosma and Kelley 2018). Jobs in startups and entrepreneurial careers seem more interesting than the more routine roles found in bureaucratic behemoths (e.g., Sheldon 2012). People may also hope to become rich. Stories abound of entrepreneurs and employees—from executives to janitors—who came into a windfall of wealth when their employers went public (e.g., Lien 2017).

But does this enthusiasm reflect reality? Despite the growing interest in entrepreneurship as an engine for job creation, little research has considered the long-term consequences of startup employment.

Whether we would anticipate positive or negative outcomes in the long run depends on how we view these individuals. If we see them as similar to founders, then we might expect startup employment to confer a long-term benefit. Individuals who become entrepreneurs—who build firms that employ others, as opposed to simply being self-employed—earn more than similar peers who are employees (Braguinsky et al. 2012, Manso 2016, Levine and Rubenstein 2017, Sorgner et al. 2017). Even former entrepreneurs who return to paid employment climb faster up the corporate ladder and earn more than otherwise-equivalent peers who never became entrepreneurs (Kaiser and Malchow-Møller 2011, Baptista et al. 2012, Luzzi and Sasson 2016).

If we instead see them more as conventional employees, however, we might expect them to earn less in the long run. In contrast to members of the founding team, startup employees typically own little equity. Their long-term financial rewards therefore depend more on the salaries that they earn and on the implications of their startup employment for the jobs available to them in the future. On those dimensions, startups have some disadvantages. Startups fail at high rates (Freeman et al. 1983, Carroll and Hannan 2000), leaving their (former) employees in search of jobs. Stints at startups may also lead individuals to develop sets of skills not valued in the labor market (Sturman et al. 2008).

We examined these issues by estimating the long-term earnings consequences of startup employment in Denmark, from 1992–2012. Those who became employees of startups—defined as firms that had been operating for four or fewer years—earned about 17% less over the next ten years.

Sorting accounted for nearly half of this difference. Startups systematically hire people who would earn less at any employer.

The consequences of a spell of startup employment also varied depending on the relative success of the startup at the time of hiring. Employees who joined a startup before the firm had grown to 50 employees earned, on average, 10%-15% less over the subsequent decade than observationally-equivalent peers. In contrast, those hired by already-successful startups—firms that had grown to more than 50 employees—earned 2%-4% more.

We explored what might account for the earnings penalty associated with employment at a small startup. Two factors appeared important. First, startup employees, particularly those who joined small firms, had less stable jobs and therefore experienced more and longer spells of unemployment. These spells led both to a short-term loss of income and to slower long-term earnings growth.

Second, those hired by startups before they became large appeared to become almost trapped in small firm employment. Most startups never grow large. Even when employees left small startups, they moved to other small firms. This path dependency in employer size creates a bifurcated labor market, with one set of individuals progressing through careers in large employers and a second set holding precarious positions at a series of startups and small firms (cf. Doeringer and Piore 1971).

Careers and startups

Although a large and growing stream of research has examined the consequences of founding a firm for entrepreneurs, much less attention has been given to the employees of these startups. What limited research has considered startup employees, moreover, has focused almost entirely on their short-term earnings. This research has found that young firms in general, and small startups in particular, pay less than large, established firms (Troske 1998, Audretsch et al. 2001, Brixy et al. 2007). Troske (1998), for example, reported that the youngest manufacturing plants in the United States paid nearly 20% less than the oldest ones.

Some of these differences stem from sorting. Startups, on average, hire younger, less experienced, and less educated employees, who would earn less in any job (Ouimet and Zarutskie 2014).

Even accounting for these differences, startups pay less (Nystrom and Elvung 2014, Ouimet and Zarutskie 2014). Ouimet and Zarutskie (2014), for example, found that startups in the U.S. paid roughly 5% less, on average, than established employers in the United States. Burton et al. (2018) similarly show that the employees of small startups in Denmark earned at least 5% less at the time of hiring than their observationally-equivalent peers who joined established firms.

But employees choose jobs based not only on their short-term rewards but also on how they expect these jobs to influence their career progression and their earnings trajectories (Spilerman 1977, Bidwell and Briscoe 2010). Prestigious internships, for example, pay poorly but offer the promise of attractive future positions. In similar fashion, startups may pay less than other employers, but experience at a startup may nevertheless lead to an earnings premium in the long run.

Researchers have yet to examine the long-term earnings consequences of startup employment but two lines of research—one looking at the effects of entrepreneurship on firm founders, the other focused on the career progression of corporate employees—address related questions. Interestingly, these two literatures point in opposite directions. Seeing startup employees as similar to founders, as proto-founders, suggests that startup employees should enjoy long-term benefits. In contrast, viewing them more as conventional employees suggests that they might suffer long-term penalties.

Startup employees as proto-founders

Firm founders experience greater variance in their earnings than otherwise similar employees (Evans and Leighton 1989, Hamilton 2000, Manso 2016). To compensate for this greater variability, we would expect them to earn more on average. With risk should come return. Consistent with this expectation, research on the returns to entrepreneurship has established that founders earn more in the long run than similar peers who never became entrepreneurs (Braguinsky et al. 2012, Levine and Rubenstein 2017, Manso 2016, Sorgner et al. 2017).[1]

Much like founders, startup employees also experience more variable income streams. They also earn less at the time of hiring (Nystrom and Elvung 2014, Ouimet and Zarutskie 2014, Burton et al. 2018). We might therefore expect them to earn more in the long run as compensation for this uncertainty and for their initial earnings penalties. Such long-term benefits might arise through multiple channels.

Rising tide.

Startup employees get in on the ground floor. Because these firms operate at smaller scale and have fewer resources, individuals enter at higher relative levels in the organizational hierarchy than they could command in a larger organization. Someone fresh out of business school, for example, might become the Chief Marketing Officer (CMO) of a startup instead of a regional brand manager for a large corporation. Such a position would not initially offer the pay or prestige associated with being the CMO of a large firm, but it might come to if the startup grows large, allowing early startup employees to attain senior positions earlier in their careers than they could have through climbing corporate ladders (Rosenbaum 1979, Stewman and Konda 1983).

Early employees of successful startups may enjoy steeper earnings trajectories even if their employers hire people at levels senior to them. As successful startups become more productive, they presumably have the capacity to reward early employees for the initial pay penalties and for joining the firm early despite the uncertainty (Adrjan 2018).

Managers-in-training.

Even if the startup fails (or does not grow large), employees, like founders, may benefit from the experience. Firms start with limited resources and without established roles and routines (Stinchcombe 1965, Freeman et al. 1983, Yang and Aldrich 2017). Startup employees must wear multiple hats. At any given time, startup employees therefore engage in a wider range of activities and develop a broader set of abilities than their peers in larger, more bureaucratic organizations (Sørensen 2007, Campbell 2013).

Startup employees also develop an even broader set of experiences over time. Startups face substantial uncertainty in how their operations will evolve. Roles and routines change frequently during the early stages of a startup, meaning that managers must redeploy employees to new responsibilities on a regular basis (Ferguson et al. 2016, Serra and Thiel 2019).

Startup employees effectively become generalists. Stints in startups could therefore act almost like a rotation program, preparing individuals for management. Senior managers in large corporations require a broad range of expertise (Blau 1970), but employees climbing corporate ladders have few opportunities to develop such breadth. Firms therefore have increasingly been turning to the outside for senior talent (Bidwell 2011). Experience in a startup may provide entry into these managerial positions. Indeed, this manager-in-training effect has been the presumed reason why startup founders rise more rapidly through the ranks if they return to corporate employment (Kaiser and Malchow-Møller 2011, Baptista et al. 2012, Campbell 2013, Luzzi and Sasson 2016).

Entrepreneurship options.

Those with experience in startups also improve their odds of becoming successful entrepreneurs. Firm founders, because they must manage and engage in nearly all of the activities of their firms, benefit from being jacks-of-all-trades (Lazear 2005, Sørensen 2007). Startup experience can help individuals to build this broad range of abilities. Startup employees also more commonly have connections with the external environment—with suppliers and customers—putting them in a better position to spot business opportunities and to mobilize resources (Sorenson and Audia 2000, Sørensen 2007, Sorenson 2017).

Startup employment may also give would-be entrepreneurs access to better organizational blueprints. Past employers serve as organizational models for entrepreneurs—templates for who to hire, how to operate, and how to organize activities within the firm (Freeman 1986, Phillips 2002). The ideal organizational structure for a startup differs from that of an established firm (Sørensen and Phillips 2011). One requires more flexibility, the other more reliability. Individuals who have prior experience in startups may therefore have a better sense of how to organize their own fledgling firms. Consistent with these expectations, studies have found that those employed by startups become entrepreneurs at higher rates than their peers in more established firms (Dobrev and Barnett 2005, Sørensen 2007, Kacperczyk 2012, Rider et al. 2019), and that their ventures survive longer (Sørensen and Phillips 2011, Dahl and Sorenson 2014).

Startup employees as conventional employees

But startup employees are, first and foremost, employees. They differ from founders in several respects. They own less equity. Their long-term compensation depends more on their wages and their wage trajectories. They have less control. They also often become involved with startups less as a matter of choice and more just because they needed jobs. If we instead turn to the literature on careers for guidance, startup employment has at least two notable disadvantages.

Precarious positions.

Startups fail at higher rates than older firms (Freeman et al. 1983, Carroll and Hannan 2000). They also need to correct course more frequently, which often entails changing their internal configurations. In either case, employees can unexpectedly find themselves out of work (Carroll et al. 1992, Haveman and Cohen 1994).

Job loss due to the failure of a startup, moreover, may prove even more problematic than job loss due to a plant closing or a mass layoff. For one, job loss due to firm failure may create a stigma, a belief that the loss reflects negatively on the individual. Rider and Negro (2015), for example, found that the employees of a failed law firm experienced substantial downward mobility, moving to less desirable employers.

Since startups seem particularly vulnerable to economic downturns (Fort et al. 2013), moreover, the (former) employees of startups may find themselves in search of employment precisely when jobs are most scarce. Spells of unemployment, particularly during economic downturns, have significant, negative, and long-lasting effects on earning trajectories (Gangl 2006, Kahn 2010, Schwandt and von Wachter 2019).

Masters-of-none.

Although the limited role differentiation in startups can produce "jack-of-all-trades" profiles, well-suited to managerial roles (Lazear 2005), engaging in a large number of activities also means spending less time learning any one. The inherent instability of startup jobs compounds this issue. Because the needs of the firm shift over time, the scope of jobs within the startup fluctuate. Startup employees may therefore end up disadvantaged in the development of the specialized experience valued by large firms (Rosen 1983). Gathmann and Schönberg (2010) estimate that the accumulation of such specialized human capital accounts for nearly half of all earnings growth.

Startup employees also frequently find themselves performing atypical jobs that aggregate idiosyncratic sets of activities (Robbins 2002, Burton and Beckman 2007). Startups often define jobs based on the backgrounds of the original occupants of those positions, aggregating unusual combinations of responsibilities (Miner 1987). Each employment experience develops two types of human capital, one firm-specific, the other with value to a range of employers (Rider et al. 2019). To the extent that startup employees become experts, they often become experts in unusual sets of activities primarily of value, in combination, to their current employers.

The resulting shallow-and-idiosyncratic human capital profiles may limit earnings growth. Having such unusual sets of experiences may impede individuals from moving into across firms. Employees experience the fastest career and earnings progression when they move up the corporate ladder *across* firms (Bidwell and Mollick 2015). But large, bureaucratic firms have highly-institutionalized expectations for how careers should unfold (Barnett and Miner 1992). Individuals who do not fit those sequences often find themselves screened out in the hiring process (Leung 2014).

Even when these individuals can move to other employers, they may lose much of the value of the human capital that they have built. To the extent that other employers only need a subset of their skills, they may pay them accordingly (Kalleberg et al. 1996, Sturman et al. 2008).

Data and measures

We examined the long-term earnings consequences of startup employment in Danish. Denmark has an economy representative in many ways of other high-income countries, such as Canada and the United States (Bingley and Westergård-Nielsen 2003, Sørensen and Sorenson 2007). It has high levels of labor mobility and of industrial dynamism—entry and exit of firms—on par with the United States (Sørensen and Sorenson 2007). In terms of areal size, population, average income, and the proportion of the labor force employed in agriculture, manufacturing, high tech, and services, it looks much like the state of Massachusetts.

Our data come from the Integrated Database for Labor Market Research (commonly referred to by its Danish acronym, IDA). This matched employer-employee database, which includes all individuals legally residing in the country, contains information on the characteristics of individuals, on their annual earnings, as well as on the characteristics of their employers.

Although IDA begins in 1980, we only analyzed data from 1991 onwards. From the late-1970s to 1990, Denmark underwent a series of regulatory reforms dismantling the centralized wage setting system (Madsen et al. 2001). Denmark now has one of the most liberal labor markets in Europe.

Employees enter our analyses the first time that they change employers after 1991 and remain in our analyses for ten years following that event. Models exploring even longer windows—up to 20 years—nevertheless revealed that the earnings penalties and premiums reported here persist far past this first decade. We restricted our analyses to individuals who changed employers between the ages of 18 and 50. That constrains the age range for our analyses to individuals from 18 to 60, largely removing retirement from the equation. We also limited our analyses to full-time employees to avoid the confounding effects of hours worked on earnings.

Startup employees.

To estimate the long-term consequences of startup employment, we must examine what individuals earn over extended periods, regardless of whether they remain at the startup or move on to other opportunities. We therefore treat joining a startup as an event—a "treatment"—rather than as a state variable. In doing so, our approach parallels studies that examine the long-term career and earnings consequences of events such as marriage, recessions, and incarceration (e.g., Western 2002, Schwandt and von Wachter 2019).

We considered employees to have experienced a spell of startup employment if, between 1992 and 2011, they became an employee of a firm that had been operating for no more than four years. We chose four years as the age threshold for being a startup because that represents the half-life of a firm in Denmark (Dahl and Sorenson 2012). Failure rates decline after that. Based on this definition, startups employ roughly 16% of the population. In this respect, also, the Danish economy appears similar to the U.S., as recent surveys report that young firms employ a similar percentage of Americans (Decker et al. 2014).

To focus on startup *employees* (as opposed to founders), we excluded both individuals with equity ownership and those employed at the organization from day one. We also used ownership information to ensure that foreign subsidiaries, spin-outs, and privatizations would not appear to be startups. As a further precaution against including spin-outs, we also dropped firms in which at least 30% of the first-year employees worked together in the prior year for another employer in the same industry and region.

Earnings.

Our primary income measure includes all hourly wages, salaries, and bonuses received by an individual in a year from all employers. Over the period covered by our analysis, the average employee in Denmark earned just over $42,000 per year. Although this measure does not capture fringe benefits or equity-based compensation, the Danish state provides most benefits, so they vary little from firm to firm. Few companies in Denmark, moreover, reward employees with shares or stock options (Eriksson 2001). Because Denmark taxes equity awards as income rather than as capital gains and has historically taxed them in the year awarded (rather than in the year exercised), such awards end up being unattractive.[2] Most private companies in Denmark therefore use bonuses to pay for performance (Eriksson 2001).[3]

Long-term earnings differentials

We begin by comparing the earnings of those who had a stint at a startup to those hired by more established firms. We do so by regressing logged earnings on an indicator variable for entering our sample as a startup employee (i.e. having been hired by a startup in the first post-1991 job change). For each individual, this indicator variable does not change over time. Our estimates therefore capture the *total* long-term earnings consequences of startup employment from a multitude of channels—disparities in initial earnings, differing slopes of earnings trajectories both at the startup and at potential future employers, and future episodes of entrepreneurship.

Beyond this binary variable, the model includes only a vector of indicator variables for the calendar year and another vector for the year in which each individual entered our analyses (i.e. first changed jobs post-1991). Together these time fixed effects adjust for average differences in earnings over time.

Table 1 reports our regression estimates of these gross ten-year differences in earnings for the employees of startups relative to those of more established firms. On average, those who became employees of startups earned about 17% less—about $58,000 for the average person in the dataset—

Table 1: Regression Estimates of Logged Long-Term Earnings

	Unmatched	Matched	Matched	Matched (primary job)	Matched (w/social pay)
Startup hire	-0.184*	-0.081*			
	(0.002)	(0.004)			
Small startup hire			-0.105*	-0.119*	-0.054*
			(0.004)	(0.004)	(0.006)
Successful startup hire			0.022*	0.018*	0.031*
			(0.007)	(0.007)	(0.009)
Small, est. firm hire			-0.093*	-0.109*	-0.042*
			(0.004)	(0.003)	(0.005)
Treatment year FE	Yes	No	No	No	No
Case-control FE	No	Yes	Yes	Yes	Yes
Calendar year FE	Yes	Yes	Yes	Yes	Yes
Observations	2,017,476	1,440,984	1,440,984	1,440,984	1,440,980

Notes: Standard errors in parentheses; * significant at the 1% level. The first three columns use income from all employers as the dependent variable; in the fourth column, income only includes earnings from the primary employer; the fifth column includes both earnings and social payments in the income measure.

over the subsequent ten years relative to those hired by more established employers (firms that had been operating for more than four years).[4] Because this ten-year window includes earnings from the startup, from other employers if these individuals leave the startup, and from future entrepreneurship, it captures all long-run upsides and downsides to being a startup employee.

These gross differences, however, stem both from selection—differences in who gets hired by startups versus established firms—as well as from treatment (i.e. from having a stint in a startup). But startups do not hire at random. In our setting, they systematically employ younger people, who earned less in their prior jobs (See Appendix A). Since these characteristics would influence the amount that an individual could expect to earn at any job and may also influence earnings trajectories, we need to adjust for this sorting to separate the effects of joining a startup from those resulting from the fact that startups tend to hire different kinds of individuals.

Earnings differentials for matched individuals

We account for the sorting of individuals to employers by using the matching method of case-control triplets introduced by Burton et al. (2018). We pair each focal individual who moved to a startup with two (control) individuals who moved to large, established firms in the same calendar year.[5] For each focal individual, we found all individuals who matched that individual exactly in terms of gender, year of birth, years of education, and prior occupation (occupation at their previous employer), and who moved in the same year as the focal individual but to employers that had been operating for more than four years and that had more than 50 employees. These large, established employers account for just over half of all employment in Denmark.

This first step produces a set of job movers who are exact matches on demographic characteristics. To ensure balance on prior income, within this set, we then selected as controls the two nearest neighbors on the *prior* year's earnings—the closest observation above and the closest below what the focal employee earned in that year.[6] The case-control triplets therefore have identical observable characteristics, all changed jobs at the same time, and earned the same amount prior to joining their new employers. Despite this tight matching, our matched sample retains more than 70% of those who moved to a startup in their first post-1991 job change. Our results for the matched sample therefore should closely reflect the average effects of startup employment in the population.

Our models on the matched sample include "triplet" fixed effects to adjust for all of the common time-invariant observed and unobserved characteristics of these triplets. This approach has a number of advantages relative to including the individual-level characteristics as regressors. Most notably, because the fixed effect for each triplet adjusts for a particular combination of attributes, it effectively allows those attributes, such as education, prior earnings, and experience, to have flexible independent and joint effects in the determination of wages (Burton et al. 2018). In other words, it does not assume any functional forms in the relationships between these factors and income. Although the triplet fixed effects absorb the treatment year fixed effects used in the first model, the matched models include indicator variables for the calendar year, to adjust for factors such as inflation and for changes in average earnings over time.

The second column of Table 1 reveals that a large portion of the observed long-term earnings differential stems from the characteristics of individuals. After adjusting for the sorting of employees to employers, those hired by startups made about 8% less over the next decade (about $27,000 less on average) than their observationally-equivalent peers at large, established firms.

We consider these results conservative. In a series of models, we also explored less restrictive matching schemes. For example, whereas our main matching specification accounts for nearly half of the gross startup employment effect, matching only on demographics—but not on the characteristics of the prior employer—reduces the startup employment effect by only about 20%.

Matching on prior income (in addition to demographics) proves powerful, increasing the amount of the effect explained by sorting to more than 40%. A host of factors can influence individual earnings, many of which, such as social skills or personality traits, elude easy measurement. But, to the extent that they remain relatively stable, these individual-level differences should contribute not only to current and future earnings but also to what individuals made in their previous jobs. By ensuring that the cases and controls have equivalent prior incomes, the matching effectively accounts for many of these difficult-to-observe factors.[7]

By contrast, restricting the control cases to those who moved jobs at the same time as those hired by startups only reduces the estimated penalties associated with startup employment by an additional 10%-15%. Any less restrictive approach to matching would therefore produce even larger estimates of the penalties associated with startup employment.

Success of startup at time of hiring.

The effects may vary as startups grow and become more bureaucratic. The opportunity to benefit from the rising tide, for example, declines but so too does the precariousness of the position. To explore how these earnings differentials depend on startup success, we estimated them based on the size of the startup, in terms of number of employees, at the time that the focal individual joined

Figure 1: Long-term Earnings Differentials Based on Startup Size at the Time of Hiring

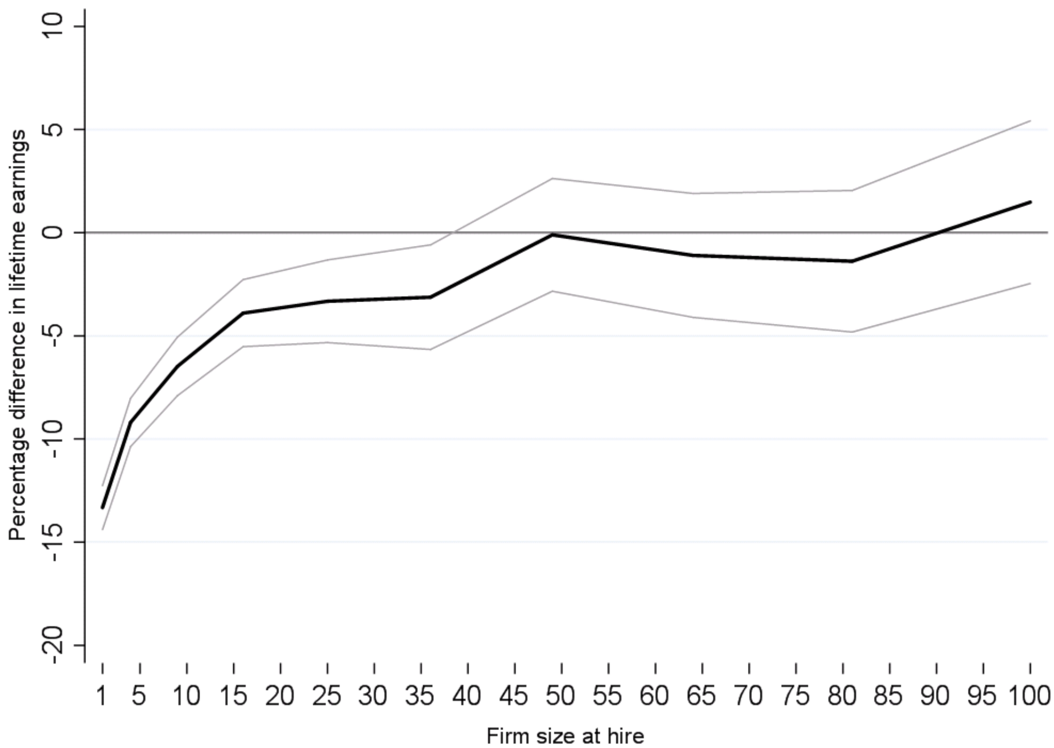

Notes: The plot displays the estimated long-term earnings differentials associated with being hired by a startup based on the number of employees at the startup at the time of hiring. The line has been smoothed by estimating these effects for bins that increase in width as a function of the square of size. The associated gray lines depict the 95% confidence intervals for these estimates.

the firm. Figure 1 describes how long-term earnings differentials vary as a function of firm size at the time of hiring. The gray lines represent the 95% confidence intervals around these estimates. We smooth the spline by estimating the pay penalty for bins of employee size. Because the data become sparser as firm size increases, we increase the width of the bins as a function of the square of the number of employees. Our first bin, therefore, covers firms with 1-4 employees while our last bin includes those with 81-100 employees.

As the figure clearly shows, the long-term pay penalties associated with being hired by a startup decline over this range. The larger the startup at the time of hiring, the smaller the penalty. In fact, beginning with the 49-64 employee size range, the confidence intervals always include zero.

We therefore distinguish between small and already-successful startups in our subsequent analyses. In doing so, we focused on a 50-employee threshold. We considered individuals hired after a startup had grown to 50 or more employees (but before the firm reached five years of age) as "successful startup" hires. Firms with 50 or more employees account for roughly 4% of employers but 58% of employment in Denmark. Only 2% of startups reach that scale before their fifth year of operations. But those firms account for 33% of startup employment.

Returning to Table 1, the third column splits startup employees into these two sets. For completeness, we similarly split established employers into two size categories: those with fewer than 50 employees and those with 50 or more employees. We treat the latter as the baseline. Consistent with the splined estimates in Figure 1, the long-term earnings penalty associated with startup employment stems entirely from the employees who joined smaller startups. Those hired by already-successful startups actually earned a small premium over the subsequent decade compared to similar others who joined large, established firms.

The next two columns explore alternative definitions of the dependent variable. In the fourth column, the earnings variable includes only pay from the primary employer. Although the employees of small and young firms more commonly have secondary employment—hence the somewhat more negative point estimates here—the exclusion of earnings from these secondary jobs only increases the earnings penalties associated with small startup employment, by about 10% across our models.

The fifth column adds nearly all forms of social payment to the dependent variable. Denmark has a generous social support system. The state provides payments to individuals in unemployment, on parental leave, and on long-term disability. These payments mitigate the negative income consequences of being out of the labor force for long periods of time. The inclusion of these social payments reduces the income penalty associated with being a small startup hire and increases the premium for those joining already-successful startups because startup employees in both categories experience more and longer spells of unemployment. Although the inclusion of this income does not change any of the qualitative conclusions of our analyses, across any of our models, including social payments in income would consistently compress the differences between the long-term earnings of the small startup hires versus the hires of already-successful startups by 30%-40%.

Do long-term earnings converge over time?

The estimates in Figure 1 and in Table 1 report average differences over a ten-year period. But wage trajectories matter as well and the groups may converge in their pay over time. Figure 2 examines how the earnings differentials unfold for the matched sample. We created this figure by estimating a separate coefficient for each group for each year post-treatment—ten coefficients for each category (e.g., one year after being hired, two years after being hired, etc.). Overall, the penalties associated with becoming an early employee of a startup appear relatively stable—rather than declining, if anything, they increase over time. In analyses on longer panels, the penalties associated with being an early startup hire persist for at least twenty years (the maximum length that our data allow us to explore). In contrast, the pay premiums associated with being hired by an already-successful startup appear short-lived, converging to a level indistinguishable from zero after only four years.

Does the small startup pay penalty stem from "lifestyle" businesses?

In the literature on the earnings of entrepreneurs, it has been crucial to distinguish individuals who hope to build large enterprises ("growth entrepreneurs") from those interested simply in a more satisfying, and perhaps easier, means of making a living ("lifestyle entrepreneurs"). Research on entrepreneurs that has attempted to distinguish between these two groups has consistently found that, while lifestyle entrepreneurs earn less than they would as employees, growth-oriented entrepreneurs earn more (e.g., Levine and Rubenstein 2017, Sorgner et al. 2017).

Might a similar distinction account for the earnings penalties associated with becoming an employee at a small startup? Perhaps people join small startups with no real expectation that they will grow, attracted instead by non-pecuniary rewards (Roach and Sauermann 2015, Sauermann

Figure 2: Earnings Trajectories

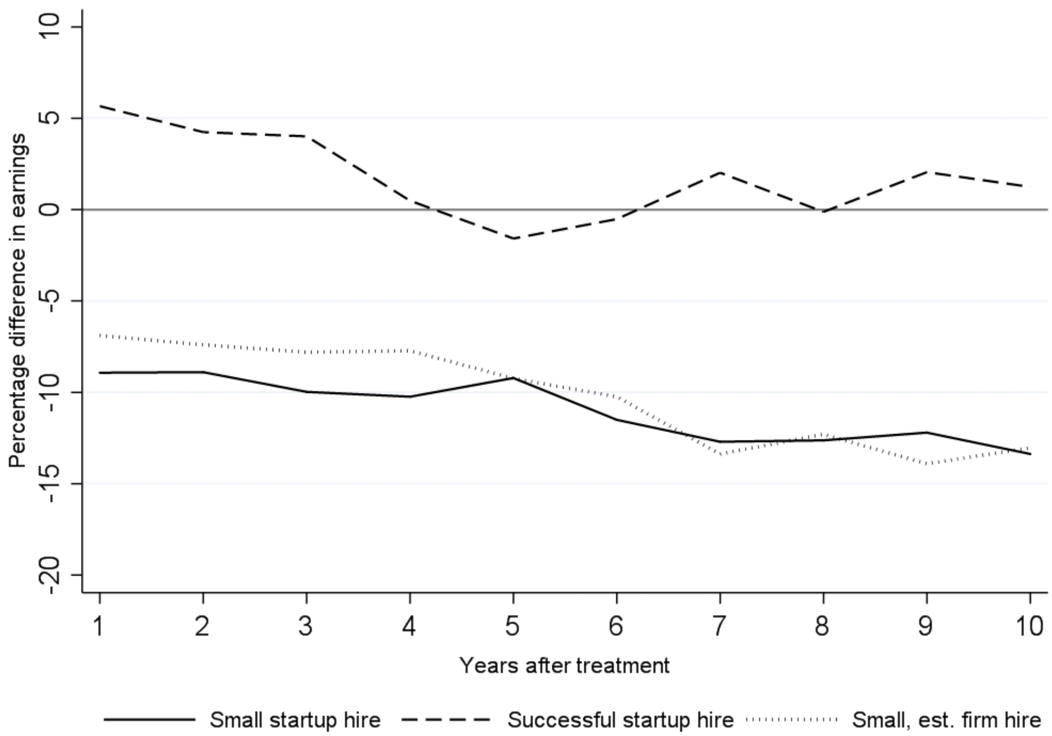

Notes: The plot displays the estimated earnings differential, relative to those hired by large, established firms, for each year after the focal individual joined a startup (or small, established firm). For each group (e.g., small startup), we included a separate indicator variable for each year post-treatment (i.e. 1 year post-hire, 2 years, ... , 10 years). The estimates are sufficiently precise that a 95% confidence interval would not be visually distinguishable from the point estimates.

2018). Individuals who join growth-oriented startups, in contrast, might have different expectations (Kim 2018).

One of the difficulties in trying to assess this question, however, stems from the fact that registry data do not capture information on entrepreneurial ambition. In the studies of entrepreneurs cited above, researchers have distinguished between lifestyle and growth-oriented entrepreneurs on the basis of being incorporated or having hired employees (e.g., Levine and Rubenstein 2017, Sorgner et al. 2017). But on those criteria, *every startup* in our sample would qualify as being growth-oriented. By construction, all of the employees in our sample work at incorporated firms with at least one full-time, non-owner employee.

What other factors might reflect the ambitions of the entrepreneur? Industry is one. Certain businesses, like a retail shop or software programming, can survive while remaining small. But others, such as shipbuilding, require scale. Although their firms might still start small, entrepreneurs in industries with economies of scale probably hope to grow large.

Following this logic, we calculated the median size for all firms in each four-digit industry and restricted the estimation to employees moving to employers in industries where the median firm had at least 50 employees. The first two columns of Table 2 report the average differences in earnings for this subset of industries. Matching appears even more important in this subset. The third column nevertheless reveals that employees who join small startups in industries that require scale also experience a long-term earnings penalty, though one about one-third smaller in magnitude than in the population of startups as a whole. Those hired by already-successful startups in these industries earn an even larger premium over the subsequent decade, of more than 4%. These estimates seem consistent with the premiums reported by Kim (2018) for successful growth-oriented startups in the United States.

Growth-oriented startups also tend to hire from the right-hand tail of the human capital distribution (Kim 2018). We therefore estimated another set of models, including only individuals who held post-graduate degrees (see Table 2). On average, these highly-educated individuals had smaller long-term earnings penalties associated with a stint of startup employment. But much of

Table 2: Regressions Estimates of Long-Term Earnings for Subsamples

	Large-firm industries			Graduate degree holders		
	Unmatched	Matched	Matched	Unmatched	Matched	Matched
Startup hire	-0.187*	-0.034*		-0.161*	-0.115*	
	(0.002)	(0.008)		(0.006)	(0.015)	
Small startup hire			-0.065*			-0.138*
			(0.010)			(0.018)
Successful startup hire			0.044*			-0.019
			(0.013)			(0.018)
Small, est. firm hire			-0.036*			-0.225*
			(0.008)			(0.021)
Treatment year FE	Yes	No	No	Yes	No	No
Case-control FE	No	Yes	Yes	No	Yes	Yes
Calendar year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,154,168	623,660	623,660	137,756	75,026	75,026

Notes: Standard errors in parentheses; * significant at the 1% level. The first three columns restrict the sample to industries with a median firm size of more than 50 employees. The second three columns restrict the sample to employees with graduate degrees.

this difference stemmed from the fact that they had a higher propensity to join already-successful startups (see Appendix A). In the matched sample estimates, the final column of the table, these highly-educated individuals suffered even larger long-term earnings penalties from being hired by a small startup. They also received no premium for joining a successful startup, relative to what their peers made in large, established firms.

Estimating earnings differentials with instrumental variables

The analysis so far has sought to account for selection effects based on differences in individual ability. The triplets approach, which carefully matches individuals to others with similar observables and prior earnings, together with the analyses restricted to growth-oriented startups and individuals with higher human capital, should account for most ability-based sorting of employees to firms. But these results do not allow us to rule out compensating differentials, such as a preference by small startup employees for the flexibility, autonomy, or dynamism associated with being in a small startup (e.g., Roach and Sauermann 2015, Sauermann 2018). We therefore turned to estimation with instrumetal variables (IV) regression as additional evidence.

An instrument uses variation unrelated to the outcome to estimate the effect of a treatment (Morgan and Winship 2007). As instruments for becoming an employee in a startup (or a small, established firm), we used the proportion of hiring occurring in particular categories (e.g., small startup hire) in an industry and region at the time that the focal individual changed employers.[8] To guard against a mechanical relationship, we did not include the job entered by the focal individual when calculating these proportions.

We build on the idea that most people search for jobs in specific regions, at specific times, and in specific industries. The supply of jobs, however, varies across places, types of firms, industries, and time. With a similar structure to a Bartik-style instrument (Bartik 1991), our instrument identifies the effects of becoming a startup employee based simply on the jobs available in a given industry at a given place and time. Although the demography of employers in an industry-region might influence pay (Haveman and Cohen 1994, Sørensen and Sorenson 2007), violating the exclusion restriction, the *flows* of job openings in any particular year, which constitute our instruments, are essentially uncorrelated with the overall demography of employers, the *stock* of jobs ($r \leq .03$).

Table 3 reports the results of these IV estimates on our matched sample. The first three columns present the first stages for these models. In each case, the prediction of the endogenous variable loads primarily on the associated instrument. All of the instruments have t-statistics in excess of 100 for predicting the endogenous variables that they instrument and, overall, the instruments explain roughly 20% of the variance in whether individuals become employees at startups versus more established employers.[9]

The second-stage coefficients indicate that the pay penalties do not stem from who chooses to join a startup. If anything, the estimates in Table 1 may understate the penalties associated with becoming an employee of a small startup. In our instrumental variables estimation, employees who join a small startup suffer an income penalty of roughly 14%. Hires of already-successful startups, by contrast, experience a boost in earnings of about 4%, similar in magnitude to that seen in growth-oriented startups. But the 95% confidence intervals around these point estimates overlap

Table 3: Instrumental Variables Estimates of Logged Long-Term Earnings

	First stages			Second stage
	Small startup	Success startup	Small, est.	
Small startup instrument	0.837*	0.020*	0.338*	
	(0.007)	(0.002)	(0.007)	
Successful startup instrument	0.106*	1.046*	0.103*	
	(0.005)	(0.008)	(0.006)	
Small, est. instrument	0.246*	0.016*	0.874*	
	(0.005)	(0.001)	(0.006)	
Small startup hire				-0.152*
				(0.015)
Successful startup hire				0.041*
				(0.012)
Small, est. hire				-0.159*
				(0.013)
Calendar year FE	Yes	Yes	Yes	Yes
F-statistic	6,853	6,113	8,888	

Notes: Based on 106,741 triplets observed for 1,367,962 person-years. Standard errors in parentheses; * significant at the 1% level.

with the confidence intervals for the coefficients reported in Table 1. We therefore cannot reject the possibility that the instrumental variables produce equivalent estimates of the effect size.

Sources of earnings differentials

Our analyses suggest that we should view startup employees—particularly small startup employees—more as employees than as proto-founders. Any positive rising tide, manager-in-training, or entrepreneurship effects appear overwhelmed by the negative effects of precarious positions and of the master-of-none effect. But these analyses provide limited insight into the mechanisms underlying these pay penalties. Table 4 introduces additional covariates to explore these mechanisms. For reference, we include the results of the baseline model—the third column of Table 1—in Column 1.

Precarious positions.

To account for spells of unemployment, we included a variable that measures, for each individual, the proportion of each year spent in the labor force. *Time in labor force* divides the number of days

Table 4: Regression Estimates of Logged Long-Term Earnings

	1	2	3	4
Small startup hire ($t = 0$)	-0.105*	-0.022*	-0.021*	0.020*
	(0.004)	(0.003)	(0.003)	(0.003)
Successful startup hire ($t = 0$)	0.022*	0.041*	0.041*	0.044*
	(0.007)	(0.005)	(0.005)	(0.005)
Small, est. hire ($t = 0$)	-0.093*	-0.032*	-0.031*	-0.006
	(0.004)	(0.002)	(0.002)	(0.002)
Time in labor market		1.852*	1.851*	1.801*
		(0.008)	(0.008)	(0.008)
Left firm			-0.007*	0.002
			(0.002)	(0.002)
Left failed firm			0.013*	0.011*
			(0.004)	(0.004)
Small startup (t)				-0.130*
				(0.004)
Successful startup (t)				0.004
				(0.002)
Small, est. (t)				-0.035*
				(0.002)
Case-control FE	Yes	Yes	Yes	Yes
Calendar year FE	Yes	Yes	Yes	Yes
Observations	1,440,984	1,393,525	1,393,525	1,317,322

Notes: Standard errors in parentheses; * significant at the 1% level.

worked (across all employers) by the total number of possible workdays for the year.[10] Column 2 of Table 4 reveals that this variable accounts for 80% of the effect of being hired by a small startup. Precariousness also probably accounts for the declining benefits associated with successful startup employment seen in Figure 2.

Importantly, the coefficient for time in labor force substantially exceeds one. The negative earnings effects of unemployment therefore do not stem solely from lost wages—in which case the coefficient would have a value of roughly one (because it has been scaled to the proportion of the year in employment). Unemployment due to job loss imposes a penalty on employees even after they secure another job. Career advancement and pay increases often depend on tenure within a given firm (Bidwell and Mollick 2015). When employees change firms, their tenure clocks reset to zero (Gangl 2006). They must start over. The former employees of failed firms also have no backup option and therefore less bargaining power. They probably end up settling for positions that pay less, not only relative to what their progress would have been in continued employment but perhaps even relative to their previous jobs. Each setback, however small, builds on those before it, leading to increasingly-divergent income paths.

Job loss due to startup failure could be even more costly if it creates a stigma. To capture the reputation effect associated with being employed by a failing firm, we included two additional variables. One, *left firm*, captures the general effects of leaving an employer and of shortened firm tenure. It turns on when individuals leave their original employers and remains turned on in all subsequent years. The second, *left failed firm*, captures the effects of leaving a firm close to its time of failure. It has a value of one—and remains one for all subsequent years—when an individual departs a firm in the same year that the firm ceases to exist. This second variable should allow us to separate the effects of any potential stigma associated with firm failure from the more general consequences of changing jobs. The third column of Table 4 reveals that moving across employers has a negative, though very small (less than 1%), effect on income. The effects of leaving a firm just before or as it fails, however, appear slightly positive. Individuals who remained on board until the very end earned just over 1% more in subsequent years. But the inclusion of these two variables has a negligible influence on the point estimates for the long-term earnings effects of startup employment.

Masters-of-none.

The masters-of-none effect suggests that the human capital of small startup employees develops in ways that limit their ability to move to larger employers. We therefore include covariates to capture the age and size of current employers. Column 4 of Table 4 reports these estimates. Being a current employee of a small startup has a large negative effect on earnings. Once one includes these measures, however, the residual long-term effect associated with having been an employee of a small startup shifts to being positive, though small (about 2%). Put differently, the small subset of startup employees who later end up employed by firms with more than 50 employees, that have been operating for more than four years, experience a small, long-term pay premium (consistent with either the rising tide effect or the manager-in-training effect).

Although the estimates in Table 4 suggest that the characteristics of current employers account for an important piece of the penalty associated with being hired by a small startup, they do not give us a sense of the extent to which these penalties stem from the fact that most startups fail to grow large versus from the restricted mobility of startup employees across firms. We borrow a log-linear modeling technique from stratification research, used to describe educational and occupational mobility, to examine mobility across different types of employers (for details, see the Appendix B). Firm-to-firm mobility depends on both firm age and firm size, but more so on firm size. An individual has about an 85% higher probability than one would expect in random mixing of moving to a firm of similar age but a more than 140% higher probability of moving to a firm of similar size. The employees of small startups, in particular, tend to move to other small startups and to small, established firms. Those employed at large, established firms, in contrast, circulate among these higher-paying employers.

Employment and earnings trajectories

Our analyses thus far have focused on averages but we can use our earnings estimates and our mobility analyses to describe the variety of potential paths experienced by those hired by startups versus by other employers. These paths provide some additional evidence for the mechanisms at play. Table 5 details the most common career paths and the average predicted long-term earnings associated with them.[11]

Table 5: Common Career Trajectories

	Proportion	Earnings
Small startup hire		
Stay at startup, remains small	45%	-1.5%
Stay at startup, grows large	3%	2.0%
Move to a small startup	27%	-21%
Move to a small firm	12%	-12%
Become entrepreneur	7%	-22%
Successful startup hire		
Stay at startup, remains large	43%	4.5%
Move to large firm	25%	3.6%
Move to large startup	16%	4.0%
Move to a small firm/startup	11%	-5.0%
Become entrepreneur	5%	-3.8%
Small firm hire		
Stay at firm, remains small	53%	-4.0%
Stay at firm, grows large	4%	-0.6%
Move to small firm	29%	-10%
Become entrepreneur	5%	-15%
Large firm hire		
Stay at firm, remains large	68%	Baseline
Move to large firm	27%	0%
Become entrepreneur	2%	-9.5%

Notes: The table reports the most common sequences for each group and the proportion of individuals in that category. For transitions across firms, it only considers the characteristics of the second employer in the 10-year observation window. Entrepreneurial transitions include all those who transition to self-employment during the 10-year window. Earnings reports the predicted earnings for the group relative to similar individuals continuously employed by large, established firms.

For employees who join startups that have not yet grown, the most probable sequence, experienced by roughly 45% of these individuals, has the startup surviving but failing to grow, leaving these individuals employed in mature but small firms and suffering a long-term pay penalty of about 1.5%. These individuals have not experienced any job loss but they still suffer a pay penalty, potentially because they have not been able to accumulate valuable human capital through learning-by-doing (Gathmann and Schönberg 2010).

If, however, the startup grows large, these early employees enjoy a boost in their long-run earnings of roughly 2%. In this sense, the popular wisdom about the long-term benefits of getting in on the ground floor—the rising tide effect—appear to hold. But these outcomes are relatively uncommon, occurring in only about 3% of cases.

More common paths for these small startup hires involve transitions to other small employers, either another small startup (27% of cases) or a small firm that has been operating for more than four years (12% of cases). Both of these trajectories predict large long-term pay penalties, of 21% on average for those moving to small startups and of 12% for those moving to more established small employers. These penalties emerge in part because those transitioning to other jobs usually experience a spell of unemployment first. They therefore probably reflect a combination of precarious positions and of having developed human capital of limited value to larger employers.

Even those who become entrepreneurs, roughly 7.5% of early employees, earn about 22% less, on average, than their observationally-equivalent peers in large, established firms. That calculation includes earnings from the dividends and capital gains associated with their startups. Although firms founded by former startup employees have been found to survive longer (Sørensen and Phillips 2011), that experience does not appear to translate into stronger financial performance.

By contrast, the majority of paths available to those joining already-successful startups—remaining at the firm, moving to another successful startup, or moving to a large, established firm—predict long-term earnings premiums. These benefits appear largest for those who remain at the firm, again providing evidence for the rising tide effect. Interestingly, even as entrepreneurs, these individuals have the highest expected financial success. Having access to successful organizational blueprints may trump the breadth of experience gained in smaller firms (Burton et al. 2002, Phillips 2002).

Finally, the consequences of employment at a small, established firm differ from those of being a small startup hire. Employees who join small, established firms tend to stay in their jobs longer and are less likely to experience extended spells of unemployment. If they do move, they tend to go to other small, established firms rather than to startups.

Discussion

For most of the twentieth century, industrial policy focused on fostering the growth of large, established enterprises, titans of the global economy, with the idea that these juggernauts would provide good jobs and promote economic expansion (e.g., Johnson 1982, Dertouzos et al. 1989). But beliefs have changed. Policymakers increasingly see entrepreneurs, not established enterprises, as the engines of economic growth (e.g., OECD 2017). Changes in regulation and public support have followed, directing more and more resources, attention, and favorable regulation toward startups. Largely absent from this shift in orientation, however, has been a consideration of how these changes might affect the career trajectories and long-run earnings prospects of employees.

Using Danish registry data, we examined the long-term earnings associated with being hired by a startup. On average, those hired by startups earn substantially less over the long run than their counterparts at more established firms. Roughly half of this earnings differential, however, stems from the sorting of individuals into organizations. Similar to what one sees in other parts of Europe and in the United States (Nystrom and Elvung 2014, Ouimet and Zarutskie 2014), startups in Denmark hire younger individuals who had been earning less in their previous jobs. In other words, these firms disproportionately hire those who would probably earn less at any employer.

Income trajectories also vary considerably depending on the size of the startup at the time of hiring. Individuals who became one of the first 50 employees of a startup earned 8% to 14% less over the subsequent ten years relative to peers employed at large, established firms. By contrast, those hired by successful startups, firms that had already grown to at least 50 employees, earned a long-term earnings *premium*, of 2% to 4%.

Only a small fraction of startups then provide attractive career opportunities, at least in terms of long-term financial rewards. Even though these successful startups account for almost one-third of startup employment, only about 2% of startups ever reach a size of 50 employees. Much as the literature on "gazelles" points to the fact that a small proportion of startups account for most job creation (Henrekson and Johansson 2010), our results suggest that these same rapidly-growing startups also create the highest-quality jobs, in terms of their effects on long-term earnings.

But, for both policymakers and potential employees, it remains unclear how early one can identify these promising startups. Coad et al. (2016) have argued that "the fog is thick" around future growth and survival prospects. Consistent with the idea that potential employees have difficulty predicting the success of startups, we found large long-term earnings penalties associated with being small startup employees even among the well educated and even in settings where entrepreneurs almost certainly have ambitions to become large.

Two factors, precarious positions and constrained job mobility, appear to account for the pay penalties associated with small startup employment. First, those in startups, particularly small startups, have less stable jobs. They experience more frequent and longer spells of unemployment. Even when they find work, they fall behind in income relative to those continuously employed. Second, those who become the employees of startups before they reach sufficient scale appear to find themselves segregated into employment in small firms. This segregation also helps to explain why those caught in lower-paying jobs do not leave them.

Our data do not allow us to isolate the ultimate sources of this constrained mobility. But the results appear consistent with a master-of-none effect. Because of their small scale, their inchoate roles, and their ever-changing environments, startup employees may develop a broad but shallow-and-idiosyncratic set of skills, closely aligned to the current needs of their organizations but disconnected from the external labor market (Miner 1987, Burton and Beckman 2007). Further progress on understanding these mechanisms will probably require detailed information both on what startup employees do on a day-to-day basis and on their experiences after leaving their original employers.

Our results nevertheless raise an obvious question: Given these long-term earnings consequences, why would anyone become one of the early employees of a startup? We see at least three possible explanations. First, employees may see startup employment as a form of compensating differential. Startups offer their employees more interesting jobs (e.g., Roach and Sauermann 2015). The smaller size of these firms may also foster camaraderie, social cohesion, among early hires. Although our IV estimates might appear to rule out this possibility, individuals could come to appreciate these advantages in the course of being an employee—that is, rather than selecting into these firms as a result of having a particular set of preferences, these tastes may emerge endogenously from their experiences. Startup employees may then happily earn less in exchange.

A second possibility is that people simply do not have a good sense of the pros and cons of startup employment. These earnings differentials accrue over time. Given the variety of career sequences and the time between cause and effect, people might find it difficult to assess the penalties associated with being an employee of a small startup. Noise may drown out the signal. Individuals also probably overweight the positive outliers, which appear to receive most of the press attention. Startup employment can lead to long-term benefits—but only for a lucky few.

A final possibility is that many employees do not choose their employers. Although students about to leave school sometimes entertain multiple job offers simultaneously, post-graduation, job searches generally do not follow a standardized calendar. Job searchers receive offers one at a time. Instead of choosing among a set of options, they decide whether to accept a job or to decline it and continue searching. The asynchronous and sequential nature of this search process introduces considerable friction into the matching of employees to employers (Mortensen 2004). Moreover, even fully aware of the consequences, job searchers may prefer an offer on the table from a small startup to the uncertain prospect of something better. Although distinguishing between these explanations lies outside the scope of the current paper and will require more detailed information on the job search process, it represents an important question for future research.

To exploit high-quality registry data, our research has focused on Denmark. Would other countries exhibit similar patterns? We suspect that many would. Research on the short-term earnings effects of joining startups has notably found quite consistent effects, and even effect sizes, across Denmark, Germany, Sweden, and the United States (Schmieder 2013, Nystrom and Elvung 2014, Burton et al. 2018, Babina et al. 2019).

But at least three country-level factors might influence the size of any earnings penalties associated with being an early employee of a startup. The first concerns the relative stature of startups in the economy (Barbulescu and Bonet 2019). Startup employment may prove more costly when changing employers in countries, such as Spain and Japan, where people have stronger preferences for large, established employers, than it would in places, such as Ireland and the United States, where entrepreneurship has been celebrated. Denmark, on this dimension, appears somewhat in the middle, with Danes giving greater esteem to entrepreneurs than those in Finland, Germany, or Spain but not as much as those in Ireland or the United States (Schøtt 2007).

A second factor involves the strength of the social support system. Denmark has a strong safety net and relatively uniform benefits across employers. But in countries where benefits vary from one employer to the next, one might expect even larger disparities to emerge between the early employees of startups and those in large, established firms. Benefits such as health care insurance and retirement plans in the United States vary with firm age and size, with older and larger firms offering better fringe benefits (Kalleberg et al. 1996, Litwin and Phan 2013). Spells of unemployment also prove more costly when these benefits depend on employment, leading the involuntarily unemployed to settle for less attractive jobs (Nickell 1997).

A final factor may involve the extent to which firms compensate their employees with equity. In Denmark, few startups offer stock or options. But in the United States and some other countries, equity awards have been more common (Hand 2008). Equity has the potential to tie long-term compensation more tightly to the success of startups, perhaps creating a fatter right-hand tail. But, even in the United States, the odds of these awards paying out end up being very low. Even when they do pay out, they often have little value.[12] The press loves to cover the janitor or receptionist who became rich from being employed at a high-tech startup. But these events are as likely and as representative of the common experience of startup employees as is the multi-million-dollar lottery winner among those buying tickets.

To a large extent, this paper represents a first step. We need to understand better how these patterns vary across countries, what ultimate causes account for them, and why some but not others end up being employed in startups. Our findings suggest that all of these topics represent important questions for future research.

But our results also suggest a need for caution. Most people rely on employment for income and wealth. Becoming the employee of a small startup, even one with ambitions to grow, leads to large long-term earnings penalties, similar in size to the penalties associated with not completing a college degree.[13] Our results therefore provide a cautionary tale for job candidates entertaining multiple offers and for policymakers harboring unfettered enthusiasm for entrepreneurship as the primary engine of job creation.

Appendix A: Employee sorting

Prior research has found that startups employ different sorts of workers (Ouimet and Zarutskie 2014, Burton et al. 2018). To assess the magnitude of these differences in our sample, we estimated a set of three logistic regressions, characterizing the extent to which each individual-level characteristic predicts being hired by a particular category of employer. The results of these models appear in Table A1. The first column compares those who joined a small startup as an employee to the employees of large, established firms; the second assesses how the employees of successful startups differ from the employees of large, established firms; and the third compares the hires of small, established firms to those of large, established firms.

Comparing across the columns, the employees of small startups are younger, more commonly men, and have lower prior incomes than the employees of large, established firms. Even successful startups continue to hire more men than large, established firms, though they no longer hire those who had been earning less in their prior jobs (for further evidence on positive selection, see Kim

2018).

Table A1: Logistic Regression Estimates of Correlates of Hiring by Age-Size Categories

	Small startup	Small, old	Successful startup
Age	-0.041*	-0.041*	-0.024*
	(0.001)	(0.000)	(0.001)
Male	0.088*	0.107*	0.136*
	(0.010)	(0.009)	(0.022)
Years of education	-0.004	-0.025*	0.001
	(0.002)	(0.002)	(0.005)
Top management	-0.087	-0.185*	-0.341*
	(0.035)	(0.031)	(0.073)
Upper white collar	-0.219*	-0.336*	-0.184*
	(0.018)	(0.017)	(0.034)
Lower white collar	-0.425*	-0.405*	-0.288*
	(0.016)	(0.015)	(0.031)
Unskilled blue collar	-0.028	0.046*	0.089
	(0.018)	(0.016)	(0.036)
Undefined or missing	0.465*	0.338*	0.354*
	(0.014)	(0.012)	(0.030)
Prior wage	-0.273*	-0.281*	-0.003
	(0.007)	(0.006)	(0.017)
Treatment year FE	Yes	Yes	Yes
Observations	322,617	339,341	280,505

Notes: For each model being hired by a large, established firm serves as the baseline. Standard errors in parentheses; * significant at the 1% level.

Appendix B: Log-linear analysis

To examine further the idea of path dependency in careers, we borrowed a modeling technique, log-linear analysis, from the stratification literature in sociology. Log-linear models have been deployed to analyze occupational and class mobility, both over time and across generations (for a recent review, see Torche 2015). For more technical information about the statistical technique, see Fienberg (2007).

These analyses begin with a cross-classified table of counts, often called a contingency table, with rows that represent the origin state and columns that represent the destination state. Under the null hypothesis of independence, one would expect each cell in the table to contain a proportion of the cases equal to the product of the marginal distributions. In other words, imagine that the rows and columns represented two sequential flips of a coin. In that case, the table would have two rows and two columns, with each state having a probability of 0.5. If one flip did not influence the next one, then each cell should contain roughly 25% of the cases ($= .5 \times .5$).

In the context of mobility, the observed cell counts represent the total amount of class or occupational mobility. Researchers examine the proportions of people who are immobile—beginning and ending in the same category—versus those who are upwardly or downwardly mobile. Stratification researchers have argued that total mobility has two distinct components: a structural component and a relative component. The structural component stems from exogenous factors that shape the marginal distribution over time. For example, the shift from a manufacturing economy to a service economy changes the relative proportions of blue-collar and white-collar jobs. The relative component captures mobility net of structural constraints. It measures the association between origins and destinations, allowing researchers to explore the extent to which those from different origins have equal opportunities.

Log-linear models usefully allow researchers to explore different forms of interdependency between the cases, conditional on the marginal distributions of the origin and destination states. In the coin flipping example, one could estimate the tendency for one side of the coin to repeat. In class or occupation studies, researchers often study inheritance or persistence, sometimes referred to as quasi-independence, which assumes a higher probability of remaining in the origin state than in moving to a different destination.

We are interested in mobility between different categories of firms. At the beginning of a year, employees could be employed in any one of four types of firms: small startup, large startup, small established firm, or large established firm. They could also be employed in any of those four types of firms at the end of the year. Each year, therefore, has a corresponding 4×4 contingency table.

Given our multiple years of data, we could aggregate all of the data to produce a single two-dimensional table (with four rows and four columns). But another option, which we pursue, involves creating a three-dimensional table, stacking the two-dimensional tables for each year. This disaggregated approach has some advantages in terms of better accounting for the marginal distributions, the availability of jobs in particular types of firms from one year to the next. We therefore have a contingency table with 320 cells $= 4 \times 4 \times 20$.

We actually constructed two different contingency tables. In the first, people who stay in the same firm year after year contribute to the cell totals. But they could still move to a different firm type if their employer matured or grew into a different category. The second table only includes individuals who changed employers in the cell counts (and the marginal distributions).

We estimated all of the models using maximum likelihood Poisson regression. We first estimated the independence model. In other words, how much of the variation in the cell counts stems simply from the marginal distributions—the number of jobs in different sorts of firms at the beginning and end of the year? We then tried more complex models to assess the extent to which employees tended to remain within employers of the same age category (0-4 vs. 5+ years) and of the same size category (1-49 vs. 50+ employees). We also allowed age and size to interact in predicting the destination state.

Table A2 summarizes the results. Because our models include such large counts, the traditional measures of model fit for log-linear analyses—chi-squared tests—would almost never suggest that one should prefer a more parsimonious model. We therefore report and focus on the amount of variance explained by each model.

Table A2: Log-linear Analysis of Employer Age-Size Mobility

	Label	Specification	d.f.	Year-to-year Pseudo R^2	Job changes Pseudo R^2
1	Conditional independence	$A_o \times S_o + A_d \times S_d$	26	.576	.760
2	Age inheritance	$[1] + A_o \times A_d$	27	.708	.780
3	Size inheritance	$[1] + S_o \times S_d$	27	.833	.880
4	Additive inheritance	$[1] + A_o \times A_d + S_o \times S_d$	28	.961	.914
5	Joint inheritance	$[4] + A_o \times A_d \times S_o \times S_d$	29	.971	.920

Beginning with the independence model, which essentially tells us the degree to which the observed patterns in the relationships between the origin and destination states stem simply from the distributions of employment opportunities available, can account for roughly 58% of the variance in employer age-size transitions and 76% of the age-size mobility across firms. A very large share of the movement of individuals across employers of varying ages and sizes therefore appears to stem not from sorting but from the opportunity structure—that is, the distribution of jobs available.

The second and third models include terms for age and size inheritance, respectively—the odds that an individual remains with an employer of small size or young in age. Comparing these models, which each use one additional degree of freedom, back to the baseline conditional independence model, we see from the relative increases in the explained variance that firm size predicts more of the variance in employee transitions than firm age. The inclusion of these two terms explains almost all of the remaining variance (Model 4).

More complicated model specifications only add marginally to the fit of the model. Model 5, for example, adds an age-size interaction term, to capture the extent to which transitions occur primarily within an age-size quadrant (e.g., only from small startups to other small startups). Although this term does improve the model fit, it only captures another 1% of the variance above and beyond the main effects of age and size. More generally, Model 4 already explains more than 92% of the variance for job changers using only 29 degrees of freedom, 26 of which simply capture the marginal distributions. Because a fully saturated model would involve interactions between the inheritance terms and each year, it would require roughly 300 additional terms to explain the remaining 3% to 8% of the variance.

Endnotes

[1]Many early studies found that “entrepreneurs” earned less than employees (Evans and Leighton 1989, Hamilton 2000, Moskowitz and Vissing-Jørgensen 2002, Hyytinen et al. 2011), but these studies did not distinguish firm founders from the self-employed. The earnings penalty in these studies stems from the self-employed, those without any employees (Åstebro et al. 2011, Levine

and Rubenstein 2017).

[2]By contrast, the United States both taxes most equity awards at a lower rate than other types of income and usually does not impose those taxes until the equity has been sold.

[3]By not including dividends and capital gains, we could understate the earnings of those who become entrepreneurs. But entrepreneurs with and without prior experience as a startup employee did not differ in their dividends or capital gains. Including these forms of income therefore would not meaningfully change our results.

[4]The coefficient $-.184$ estimates the difference in logged earnings. One can approximate this difference in percentage terms using the antilogarithm (i.e. $= e^{-.184}$). Over the period being analyzed, the exchange rate of the Danish kroner to the U.S. dollar averaged about 6.5:1.

[5]Each focal individual can only appear once in the analyses. If a person has multiple instances of being hired by a startup, we included only the first event as a treatment. Control individuals, however, have been selected with replacement. Matching with replacement introduces some correlation across triplets into the error structure, but it has the advantage of ensuring the retention of more treated cases and therefore of generating estimates more representative of the population average effect (Abadie and Imbens 2006).

[6]Our approach combines coarsened-exact matching with nearest-neighbor matching on income. Coarsened-exact matching guarantees balance between the cases and the controls on all dimensions used for matching. For extended discussions of the advantages of coarsened-exact matching relative to propensity score matching, see Iacus et al. (2012), King and Nielsen (2019).

[7]By contrast, employee-level fixed effects would produce biased estimates because they confound the effects of startup employment on future jobs with time-invariant individual differences.

[8]As regions, we use the 79 labor markets developed by Andersen (2000).

[9]The Angrist-Pischke F-statistic has values ranging from 6,113 to 8,888. Although no critical values are available for this test, it has the same distribution as the F-statistic for which Stock and Yogo (2005) calculated critical values. Our first stages exceed their recommended critical values, of 10-20, by two orders of magnitude.

[10]Unreported models with a quadratic specification for time in the labor force do not improve the overall fit of the model but they do suggest increasing marginal penalties as individuals spend more time in unemployment.

[11]The table does not include all possible paths. Some trajectories that apply to less than 2% of individuals have been omitted.

[12]According to a recent survey by Schwab, for those receiving equity compensation, the median package had a value of only $21,000 (O'Brien 2018).

[13]For estimates of the returns to education in Denmark, see Sorenson and Dahl (2016).

References

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

Adrjan, Pawel. 2018. Risky Business? Earnings Prospects of Employees at Young Firms. Discussion Paper University of Oxford Department of Economics Discussion Paper 852, University of Oxford.

Andersen, Anne Kaag. 2000. *Commuting Areas in Denmark*. AKF Forlaget, Copenhagen.

Åstebro, Thomas, Jing Chen, Peter Thompson. 2011. Stars and misfits: Self-employment and labor market frictions. *Management Science* **57**(11) 1999–2017.

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**(5) 795–821.

Babina, Tania, Wenting Ma, Christian Moser, Paige Ouimet, Rebecca Zarutskie. 2019. Pay, employment, and dynamics of young firms. Working paper, Columbia Univerity.

Baptista, Rui, Francisco Lima, Miguel Torres Preto. 2012. How former business owners fare in the labor market? Job assignment and earnings. *European Economic Review* **56**(2) 263–276.

Barbulescu, Roxana, Rocio Bonet. 2019. Outrunning the Past? Prior Employer Status and Job Matching in the MBA Labor Market. Working paper, HEC Paris.

Barnett, William P., Anne S. Miner. 1992. Standing on the shoulders of others: Career interdependence in job mobility. *Administrative Science Quarterly* **37**(2) 262–281.

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?*. W.E. Upjohn Institute for Employment Research, Kalamazoo, MI.

Bidwell, Matthew. 2011. Paying more to get less: Specific skills, incomplete information and the effects of external hiring. *Administrative Science Quarterly* **56**(3) 369–407.

Bidwell, Matthew, Forrest Briscoe. 2010. The dynamics of interorganizational careers. *Organization Science* **21**(5) 1034–1053.

Bidwell, Matthew, Ethan Mollick. 2015. Shifts and ladders: Comparing the role of internal and external mobility in managerial careers. *Organization Science* **26**(6) 1629–1645.

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

Blau, Peter M. 1970. A formal theory of differentiation in organizations. *American Sociological Review* **35**(2) 201–218.

Bosma, Neils, Donna Kelley. 2018. *Global Entrepreneurship Monitor 2019/2019 Global Report*. Global Entrepreneurship Research Association (GERA).

Braguinsky, Serguey, Steven Klepper, Atsushi Ohyama. 2012. High-tech entrepreneurship. *Journal of Law and Economics* **55**(4) 869–900.

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.

Burton, M. Diane, Christine M. Beckman. 2007. Leaving a legacy: Position imprints and successor turnover in young firms. *American Sociological Review* **72**(2) 239–266.

Burton, M. Diane, Michael S. Dahl, Olav Sorenson. 2018. Do startups pay less? *Industrial and Labor Relations Review* **71**(3) 1179–1200.

Burton, M. Diane, Jesper B. Sørensen, Christine M. Beckman. 2002. Coming from good stock: Career histories and new venture formation. *Research in the Sociology of Organizations* **19** 229–262.

Campbell, Benjamin A. 2013. Earnings effect of entrepreneurial experience: Evidence from the semiconductor industry. *Management Science* **59**(2) 286–304.

Carroll, Glenn R., Michael T. Hannan. 2000. *The Demography of Organizations and Industries*. Princeton University Press, Princeton, NJ.

Carroll, Glenn R., Heather A. Haveman, Anand Swaminathan. 1992. Careers in organizations: An ecological perspective. *Life-Span Development and Behavior* **11** 112–144.

Coad, Alex, Julian S. Frankish, Richard G. Roberts, David J. Storey. 2016. Predicting new venture survival and growth: Does the fog lift? *Small Business Economics* **47**(1) 217–241.

Dahl, Michael S., Olav Sorenson. 2012. Home sweet home? Entrepreneurs' location choices and the performance of their ventures. *Management Science* **58**(6) 1059–1071.

Dahl, Michael S., Olav Sorenson. 2014. The who, why and how of spinoffs. *Industrial and Corporate Change* **23**(3) 661–688.

de Wit, Gerrit, Jan de Kok. 2014. Do small businesses create more jobs? New evidence for Europe. *Small Business Economics* **42**(2) 283–295.

Decker, Ryan, John Haltiwanger, Ron Jarmin, Javier Miranda. 2014. The role of entrepreneurship in U.S. job creation and economic dynamism. *Journal of Economic Perspectives* **28**(3) 3–24.

Dertouzos, Michael L., Richard K. Lester, Robert M. Solow. 1989. *Made in America*. The MIT Press, Cambridge, MA.

Dobrev, Stanislav D., William P. Barnett. 2005. Organizational roles and the transition to entrepreneurship. *Academy of Management Journal* **48**(3) 433–449.

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

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

Evans, David S., Linda S. Leighton. 1989. Some empirical aspects of entrepreneurship. *American Economic Review* **79**(3) 519–535.

Ferguson, Amanda J., Lisa E. Cohen, M. Diane Burton, Christine M. Beckman. 2016. Misfit and milestones: Structural elaboration and capability reinforcement in the evolution of entrepreneurial top management teams. *Academy of Management Journal* **59**(4) 1430–1450.

Fienberg, Stephen E. 2007. *The Analysis of Cross-classified Categorical Data*. Springer Science & Business Media.

Fort, Teresa C, John Haltiwanger, Ron S Jarmin, Javier Miranda. 2013. How firms respond to business cycles: The role of firm age and firm size. *IMF Economic Review* **61**(3) 520–559.

Freeman, John. 1986. Entrepreneurs as organizational products: Semiconductor firms and venture capital firms. G. Libecap, ed., *Advances in the Study of Entrepreneurship, Innovation, and Economic Growth*, vol. 1. JAI Press, Greenwich, CT, 33–58.

Freeman, John H., Glenn R. Carroll, Michael T. Hannan. 1983. The liability of newness: Age dependence in organizational death rates. *American Sociological Review* **88**(5) 692–710.

Gangl, Markus. 2006. Scar effects of unemployment: An assessment of institutional complementarities. *American Sociological Review* **71**(6) 986–1013.

Gathmann, Christina, Uta Schönberg. 2010. How general is human capital? A task-based approach. *Journal of Labor Economics* **28**(1) 1–49.

Haltiwanger, John, Ron S. Jarmin, Javier Miranda. 2013. Who creates jobs? Small versus large versus young. *Review of Economics and Statistics* **95**(2) 347–361.

Hamilton, Barton. 2000. Does entrepreneurship pay? An empirical analyis of the returns to self-employment. *Journal of Political Economy* **108**(3) 604–631.

Hand, John R.M. 2008. Give everyone a prize? Employee stock options in private venture-backed firms. *Journal of Business Venturing* **23**(4) 385–404.

Haveman, Heather A., Lisa E. Cohen. 1994. The ecological dynamics of careers: The impact of organizational founding, dissolution, and merger on job mobility. *American Journal of Sociology* **101**(1) 104–152.

Henrekson, Magnus, Dan Johansson. 2010. Gazelles as job creators: A survey and interpretation of the evidence. *Small Business Economics* **35**(2) 227–244.

Hyytinen, Ari, Pekka Ilmakunnas, Otto Toivanen. 2011. The return-to-entrepreneurship puzzle. *Labour Economics* **20**(1) 57–67.

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

Johnson, Chalmers. 1982. *MITI and the Japanese Miracle*. Stanford University Press, Stanford, CA.

Kacperczyk, Aleksandra J. 2012. Opportunity structures in established firms: Entrepreneurship versus intrapreneurship in mutual funds. *Administrative Science Quarterly* **57**(3) 484–521.

Kahn, Lisa B. 2010. The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics* **17**(2) 303 – 316.

Kaiser, Ulrich, Nikolaj Malchow-Møller. 2011. Is self-employment really a bad experience? The effects of previous self-employment on subsequent wage-employment wages. *Journal of Business Venturing* **26**(5) 572–588.

Kalleberg, Arne L., David Knoke, Peter V. Marsden, Joe L. Spaeth. 1996. *Organizations in America: Analysing Their Structures and Human Resource Practices*. Sage Publications, Thousand Oaks, CA.

Kim, J. Daniel. 2018. Is there a startup wage premium? Evidence from MIT graduates. *Research Policy* **43**(3) 637–649.

King, Gary, Richard Nielsen. 2019. Why propensity scores should not be used for matching. *Political Analysis* forthcoming.

Lawless, Martina. 2014. Age or size? Contributions to job creation. *Small Business Economics* **42**(4) 815–830.

Lazear, Edward P. 2005. Entrepreneurship. *Journal of Labor Economics* **23**(4) 649–680.

Leung, Ming D. 2014. Dilettante or renaissance person? How the order of job experiences affects hiring in an external labor market. *American Sociological Review* **79**(1) 136–158.

Levine, Ross, Yona Rubenstein. 2017. Smart and illicit: Who becomes an entrepreneur and do they earn more? *Quarterly Journal of Economics* **132**(2) 963–1018.

Liberto, Jennifer. 2012. Obama signs bill to help companies go public. *CNN Money* .

Lien, Tracey. 2017. When tech companies go public, employees can strike it rich – or not. and then the trouble starts. *Los Angeles Times* .

Litwin, Adam Seth, Phillip H. Phan. 2013. Quality over quantity: Reexamining the link between entrepreneurship and job creation. *ILR Review* **66**(4) 833–873.

Luzzi, Alessandra, Amir Sasson. 2016. Individual entrepreneurial exit and earnings in subsequent paid employment. *Entrepreneurship Theory and Practice* **40**(2) 401–420.

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

Manso, Gustavo. 2016. Experimentation and the returns to entrepreneurship. *Review of Financial Studies* **29**(9) 2319–2340.

Miner, Anne S. 1987. Idiosyncratic jobs in formalized organizations. *Administrative Science Quarterly* **32**(3) 327–351.

Morgan, Stephen L., Christopher Winship. 2007. *Counterfactuals and Causal Inference*. Cambridge University Press.

Mortensen, Dale T. 2004. *Wage Dispersion: Why Are Similar Workers Paid Differently?*. MIT Press, Cambridge, MA.

Moskowitz, Tobias J., Annette Vissing-Jørgensen. 2002. The returns to entrepreneurial investment: A private equity premium puzzle? *American Economic Review* **92**(4) 745–778.

Nickell, Stephen. 1997. Unemployment and labor market rigidities: Europe versus north america. *Journal of Economic Perspectives* **11**(3) 55–74.

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**(2) 399–410.

O'Brien, Sarah. 2018. Don't overlook the risk that comes with your employee stock options. *CNBC* .

OECD. 2017. *Entrepreneurship at a Glance 2017*. OECD.

Ouimet, Paige, Rebecca Zarutskie. 2014. Who works for startups? The relation between firm age, employee age, and growth. *Journal of Financial Economics* **112**(3) 386–407.

Phillips, Damon J. 2002. A geneological approach to organizational life chances: The parent-progency transfer and Silicon Valley law firms, 1946-1996. *Administrative Science Quarterly* **47**(3) 474–506.

Rider, Chris I., Giacomo Negro. 2015. Organizational failure and intraprofessional status loss. *Organization Science* **26**(3) 633–649.

Rider, Christopher I., Peter Thompson, Aleksandra Kacperczyk, Joacim Tåg. 2019. Experience and entrepreneurship: A career transition perspective. *ILR Review* **72**(5) 1149–1181.

Roach, Michael, Henry Sauermann. 2015. Founder or joiner? The role of preferences and context in shaping different entrepreneurial interests. *Management Science* **61**(9) 2160–2184.

Robbins, Gregory. 2002. Titles and tasks: New jobs for new media in Silicon Alley? Ph.D. thesis, Columbia University.

Rosen, Sherwin. 1983. Specialization and human capital. *Journal of Labor Economics* **1**(1) 43–49.

Rosenbaum, James. 1979. Organizational career mobility: Promotion chances in a corporation during periods of growth and contraction. *American Journal of Sociology* **85**(1) 21–48.

Sauermann, Henry. 2018. Fire in the belly? Employee motives and innovative performance in start-ups versus established firms. *Strategic Entrepreneurship Journal* **12**(4) 423–454.

Schmieder, Johannes. 2013. What Causes Wage Dispersion? Evidence from New Firms. Working paper, Boston University.

Schøtt, Thomas. 2007. *Entrepreneurship in the Regions in Denmark*. University of Southern Denmark, Odense.

Schwandt, Hannes, Till von Wachter. 2019. Unlucky cohorts: Estimating the long-term effects of entering the labor market in a recession in large cross-sectional data sets. *Journal of Labor Economics* **37**(S1) S161–S198.

Serra, Caroline K., Jana Thiel. 2019. Professionalizing entrepreneurial firms: Managing the challenges and outcomes of founder-ceo succession. *Strategic Entrepreneurship Journal* **13**(3) 379–409.

Sheldon, Kerrin. 2012. 8 reasons to choose a startup over a corporate job. *Fast Company* .

Sørensen, Jesper B. 2007. Bureaucracy and entrepreneurship: Workplace effects on entrepreneurial entry. *Administrative Science Quarterly* **52**(3) 387–412.

Sørensen, Jesper B., Damon J. Phillips. 2011. Competence and commitment: Employer size and entrepreneurial endurance. *Industrial and Corporate Change* **20**(5) 1277–1304.

Sørensen, Jesper B., Olav Sorenson. 2007. Corporate demography and income inequality. *American Sociological Review* **72**(5) 766–783.

Sorenson, Olav. 2017. Regional ecologies of entrepreneurship. *Journal of Economic Geography* **17**(5) 959–974.

Sorenson, Olav, Pino G. Audia. 2000. The social structure of entrepreneurial activity: Geographic concentration of footwear production in the United States, 1940-1989. *American Journal of Sociology* **106**(2) 424–462.

Sorenson, Olav, Michael S. Dahl. 2016. Geography, joint choices, and the reproduction of gender inequality. *American Sociological Review* **81**(5) 900–920.

Sorgner, Alina, Michael Fritsch, Alexander Kritikos. 2017. Do entrepreneurs really earn less? *Small Business Economics* **49**(2) 251–272.

Spilerman, Seymour. 1977. Careers, labor market structure, and socioeconomic achievement. *American Journal of Sociology* **83**(3) 551–593.

Stewman, Shelby, Suresh L. Konda. 1983. Careers and organizational labor markets: Demographic models of organizational behavior. *American Journal of Sociology* **88**(4) 637–685.

Stinchcombe, Arthur L. 1965. Social Structure and Organizations. *Handbook of Organizations*, chap. 4. Rand McNally, 142–155.

Stock, James H., Motohiro Yogo. 2005. *Identification and Inference for Econometric Models: Essays in Thomas Rothenberg*, chap. Testing for weak instruments in linear IV regression. Cambridge University Press, 80–108.

Sturman, Michael C., Kate Walsh, Robin A. Cheramie. 2008. The value of human capital specificity versus transferability. *Journal of Management* **34**(2) 290–316.

Torche, Florencia. 2015. Analyses of intergenerational mobility: An interdisciplinary review. *The ANNALS of the American Academy of Political and Social Science* **657**(1) 37–62.

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

Western, Bruce. 2002. The impact of incarceration on wage mobility and inequality. *American Sociological Review* **67**(4) 526–546.

Yang, Tiantian, Howard E. Aldrich. 2017. 'The liability of newness' revisited: Theoretical restatement and empirical testing in emergent organizations. *Social Science Research* **63** 36–53.