Cohort size and transitions into the labour market

Duncan Roth[1]

Abstract

This paper estimates the effect that the size of an individual's labour-market entry cohort has on the subsequent duration of search for employment. Survival-analysis methods are applied to empirically assess this relationship using a sample of apprenticeship graduates who entered the German labour market between 1999 and 2012. The results suggest that apprentices from larger graduation cohorts take less time to find employment, but this effect appears to be significant only for a period of up to six months after graduation. These results therefore do not support the cohort-crowding hypothesis that members of larger cohorts face depressed labour-market outcomes. Moreover, there is no evidence that shorter search durations are the result of graduates being pushed into lower-quality employment. The finding that graduating as part of a larger cohort leads to shorter search durations is in line with those parts of the cohort-size literature that find larger youth cohorts being associated with lower unemployment rates. A possible explanation is that firms react to an anticipated increase in the number of graduates by creating jobs.

JEL classification J21, J64, R23

Key words Survival analysis, entry conditions, cohort size, apprentices, search duration

[1] Institute for Employment Research (IAB), Email: duncan.roth@iab.de

1 Introduction

The extant cohort-size literature has predominantly focussed on how the size of a specifically defined age group affects the wage (Mosca, 2009; Brunello, 2010; Morin, 2015; Garloff and Roth, 2016; Moffat and Roth, 2016a) and (un-)employment outcomes (Korenman and Neumark, 2000; Shimer, 2001; Skans, 2005; Biagi and Lucifora, 2008; Garloff et al., 2013; Moffat and Roth, 2016b) of that group. In contrast, the question how cohort-size shapes an individual's transition into the labour market and subsequent career has so far been left largely unaddressed, although the demographic processes which are projected to lead to reductions in population size and changes in age structures throughout Europe (European Commission, 2014) and in Germany in particular (Statistisches Bundesamt, 2015) would appear to provide a motivation to better understand this relationship.

This paper addresses this question by estimating the effect that an increase in the size of the cohort of graduates from Germany's apprenticeship system has on the duration that apprentices spend searching for employment following graduation. Specifying a cohort-size variable in terms of the group as part of which an individual enters the labour market sets this study apart from the majority of the above-mentioned literature in which cohort size typically refers to the contemporaneous size of an age group. As such, this paper is also related to a recent literature on the effects of the state of the local labour market at the time of entry – usually, based on a measure for the business cycle – on subsequent labour-market outcomes (Stevens, 2007; Kahn, 2010; Brunner and Kuhn, 2014; Cockx and Ghirelli, 2016) since the size of the graduation cohort within a local labour market also represents a feature of the conditions under which labour-market entry takes place.

The use of apprenticeship graduates in this paper as opposed to population-based age groups, which is common in the extant cohort-size literature, also provides a better measure of a group that is relevant to the labour market and therefore allows a better assessment of the consequences of labour-market crowding. It is typically assumed that individuals within a cohort are substitutable for each other, but that there is imperfect substitution across cohorts. This assumption is more likely to hold among apprenticeship graduates since the majority of the former are not only of a similar age but also share a comparable level of qualification which makes it more likely that they will be competing on the same labour market than two individuals that only belong to the same age group. Constructing cohort size from apprenticeship graduates who have completed their training and are therefore ready to enter the labour market should, moreover, reduce the problem of measurement error in this variable. This problem arises when cohorts that are based on young age groups are included, as large parts of the former are likely to be unavailable to the labour market (see Moffat and Roth, 2016b).

From a theoretical perspective the sign of the effect that the size of an individual's graduation cohort has on his subsequent search duration is ex ante unclear. The results of the empirical analysis suggest that belonging to a larger graduation cohort is predicted to reduce search duration. Specifically, the effect of a rise in the size of the entry cohort by one standard deviation is predicted to increase the hazard rate of finding employment by approximately 8%, which is comparable in magnitude to the effect of a corresponding increase in the unemployment at the time of entry. This effect, however, is only significant within a relatively short period following graduation. The empirical analysis therefore does not provide any evidence that members of larger entry cohorts face longer search durations. Moreover, the results do not suggest that shorter search durations come at the price of taking up employment in lower-quality jobs. Alternative explanations for the pattern of the regression results relating to selected migration after graduation or changes in the productivity composition in larger cohorts are also not supported by the data. While offering no direct evidence for the mechanisms suggested by parts of the literature that find that larger youth cohorts reduce the youth unemployment rate, these results are nevertheless in line with the hypothesis that an increase in the size of an entry cohort induces an expansion in labour demand.

The remainder of the paper is structured as follows: Section 2 provides an overview of the extant literature; the empirical analysis is the subject of Section 3, while Section 4 contains the results; Section 5 concludes.

2 Literature and hypotheses

The subject of this paper is related to a large body of literature that analyses the impact that the size of a cohort has on the labour-market outcomes of its members. In this literature the term cohort usually refers to a group of individuals that fall into a specified age range, though in some cases cohorts are also differentiated with respect to educational attainment. The main motivation for defining cohorts in this way is the assumption that differently aged individuals are only imperfectly substitutable for each other and can be thought of as distinct factors of production (Card and Lemieux, 2001). The reason for this assumption is that older individuals tend to have more years of work experience, which in turn makes it more likely that they have acquired more human capital of various types (general, industry-specific, occupation-specific and job-specific). As long as a worker's productivity is related to the amount of human capital he has acquired, it follows that differently aged individuals should only be imperfectly substitutable (a more detailed discussion can be found in Garloff and Roth, 2016, and Moffat and Roth, 2016a).

Most research within this literature has so far concentrated on the effect of cohort size on wages as well as on employment and unemployment outcomes. In the case of wages the benchmark model of a perfectly competitive labour-market predicts that if there is diminishing marginal productivity of labour an increase in cohort size reduces the wages earned by its members (Brunello, 2010), while Michaelis and Debus (2011) show that a similar result holds in the case of an imperfectly competitive labour market in which wages are set by monopoly unions. Findings by Garloff and Roth (2016) suggest that a considerable part of the negative effect can be ascribed to members of larger age groups being more likely to find employment in lower-paying occupations and industries. A substantial body of empirical research from different countries and time periods provides evidence in support of the hypothesis that increases in cohort size reduce the wages of its members (Freeman, 1979; Welch, 1979; Berger, 1983; Dooley, 1986; Wright, 1991; Mosca, 2009; Brunello, 2010; Morin, 2015; Garloff and Roth, 2016; Moffat and Roth, 2016a). However, if age-specific wages are rigid or the number of jobs for members of an age group are limited, changes in cohort size might rather affect age-specific employment or unemployment. The empirical literature provides conflicting evidence on this issue with some studies finding that larger youth cohorts lead to depressed employment and unemployment outcomes (Korenman and Neumark, 2000; Biagi and Lucifora, 2008; Garloff et al., 2013), while others provide evidence of a positive effect (Shimer, 2001; Skans, 2005; Moffat and Roth, 2016b).

One feature of the cohort-size literature is that the former's impact is typically analysed for contemporaneous outcomes. While this paper also utilises the concept of a cohort as a group of individuals with similar characteristics, it differs by defining a cohort-size variable that refers to a specific point in time – the time of entry into the labour market – and estimates its effect on the subsequent duration of search for employment. In light of this set-up, the paper is also relevant to a recent literature analysing the effects that the conditions prevailing at the time of an individual's entrance into the labour market have on subsequent labour-market outcomes. In this literature these conditions refer to the state of the economy when an individual enters the labour market which is typically measured by the local or national unemployment rate, the most commonly used outcome variables being an individual's subsequent wages or earnings, though some studies also consider the effect on annual hours worked or the employment rate. Initially, entering the labour market during an economic downturn has the effect of increasing the probability of being unemployed, while individuals may also be pushed into lower-paying jobs. This initial effect can become persistent if these jobs offer fewer opportunities to acquire productivity-enhancing human capital and if individuals fail to transfer to a higher-quality job at a later stage. Evidence for the hypothesis that labour-market entry during an economic downturn can lead to lasting depressed labour-market outcomes is provided by a number of studies (Stevens, 2007; Kahn, 2010; Brunner and Kuhn, 2014; Cockx and Ghirelli, 2016).

However, as the literature on cohort-size effects suggests, the state of the economy does not necessarily constitute the only factor that is relevant to an individual's labour-market outcomes and the supply of similarly aged and qualified individuals may also represent an important entry condition. So far, evidence on the effects of cohort size at the time of labour-market entry is scarce – Morin (2015) analyses changes in the size of the Canadian school graduation cohorts on subsequent wage outcomes and the quality of employment – and the former's relationship with search duration, which is the subject of this paper, has so far not been studied. The effect that an increase in the size of the entry cohort has on the amount of time that its members have to search before finding employment is ex ante unclear. However, the cohort-size literature and especially the mechanisms underlying the relationship with (un-)employment outcomes provide a basis from which to derive hypotheses. The cohort-crowding argument states that in the absence of a full and immediate adjustment in cohort-specific wages, an increase in cohort size leads to depressed employment and unemployment outcomes due to increased competition. Within such a framework members of larger entry cohorts can be expected to have longer search durations. However, the relationship between entry-cohort size and search duration would become indeterminate if members of larger cohorts avoided prolonged search durations by (temporarily) moving into lower-quality jobs. In such a scenario the effect of increased competition may be fully or partially countered depending on how many individuals would be prepared to select into such jobs and how quickly this would happen.

Finally, a possible rationale for members of larger cohorts having shorter search durations is provided by Shimer (2001) who finds that an increase in the size of the youth cohort reduces the unemployment rate of that age group (as well as of other groups). In his model the primary difference between younger and older individuals is that the former are more likely to be either unemployed or employed but poorly matched and therefore more prepared to either take up or switch jobs. An increase in the size of the youth cohort therefore leads to a larger supply of individuals that can be recruited by firms. The central assumption is the existence of a trading externality: that there is a higher probability of employers and job searchers realising a match if the number of trading partners is large. Given this assumption, firms are predicted to react to an increase in the size of the youth cohort by creating vacancies because the larger number of unemployed or poorly matched young individuals increases the probability of making a match. However since new matches can also be poor matches – in which case an individual would continue searching for other job opportunities – firms have an incentive to continue creating vacancies with the result that the overall unemployment rate and the unemployment rate of the young decreases. Within this framework it is conceivable that members of larger entry cohorts have shorter search durations. In order to assess the validity of the above hypotheses, the relationship between entry-cohort size and search duration is analysed empirically based on a sample of graduates from Germany's apprenticeship system who enter the labour market between 1999 and 2012.

In addition to analysing the effect of changes in cohort size on an outcome that has so far not been considered, this paper is also able to deal with two sources of measurement error which are usually not addressed in the cohort-size literature. First, cohorts are supposed to measure the amount of individuals with similar characteristics that are active on the same labour market. Usually, administrative units at different levels of aggregation are used as the spatial basis from which to construct the cohort-size variable. These units do not necessarily provide good measures of actual labour markets because they are typically not delineated according to economic criteria. As a result, a cohort-size variable derived from administrative units is subject to measurement error because it is likely to group together individuals that are not active on the same labour market. This paper addresses this concern by employing the labour-market regions defined by Kosfeld and Werner (2012), which combine one or more administrative units based on the degree of commuting between these units. By creating as large an overlap as possible between the resident and the working population, these functional entities approximate actual labour markets.

Second, cohort-size variables are usually derived from the size of different age groups. Concerning the fact that members of a cohort are supposed to be available to the labour market, this approach can be problematic if considerable parts of an age group are non-participants and as such do not influence the labour-market outcomes of their age group. This is a particular concern for young age groups as their members are often engaged in education and are therefore not available to the labour market. Moffat and Roth (2016b) show that the inclusion of young age groups in the analysis of the relationship between cohort size and (un-)employment outcomes has considerable implications for size and sign of the cohort-size coefficient. This problem should be less of a concern in this study as apprenticeship graduates should be more likely to be available to the labour market.

3 Empirical analysis

3.1 Data

The empirical analysis of this paper utilises two different data sources. To construct the model's main explanatory variable – the number of graduates from an apprenticeship programme – the Integrated Employment Biographies (IEB) are used. This dataset contains information on all individuals who belong to one of the following groups: employees subject to social security contributions, marginal employees, individuals receiving unemployment benefits, individuals registered as seeking employment and participants in the Federal Employment Agency's (FEA) measures of labour-market policy (groups that are not covered are civil servants and the self-employed). For each individual the dataset consists of different records that correspond to episodes in one of the above-mentioned states with specified start and end dates. Moreover, each episode is supplemented with two different sorts of information: first, characteristics of the individual are provided which refer to the beginning of the episode (among others, these characteristics include sex, nationality, year of birth, place of residence and level of education); second, details are provided that describe the state an individual is in (in the case of an employment episode, information would be available on the average daily wage during the episode, the occupation and industry of employment, place of employment as well as on the type of employment).[1]

Participation in apprenticeship programmes constitutes a separate type of employment (employment subject to social security contributions and marginal employment constitute other major categories) and as such it is possible to determine whether an individual is participating in such a training programme at any given point in time. Because a change in the type of employment – e.g. when an individual completes an apprenticeship and takes up another form of employment – entails that a new episode is defined, it is further possible to identify when participation in an apprenticeship programme has ended. Based on this information, the number of individuals graduating from such a training programme in a given month, year and region can be estimated (Section 3.2. provides further details on the conditions that are imposed for an individual to be regarded as having completed training). Due to its size working directly with IEB records can be cumbersome and therefore the regression analysis of this paper uses a 2% sample, the so-called Sample of Integrated Employment Biographies (SIAB).

3.2 Sample and variables

The sample consists of male individuals aged between 19 and 23 who have completed an apprenticeship. Construction of the sample from the SIAB dataset proceeds as follows: first, those individuals without any episode as an apprentice are removed. For the remaining individuals it is then decided whether the information on the registered apprenticeship episodes also warrants the assumption that training was completed. This is done by imposing two criteria: first, the combined duration of apprenticeship episodes has to be at least 730 days. While completion of training can often require more than two years, it is the case that individuals with a higher secondary education degree are able to complete an apprenticeship faster than those without a comparable schooling certificate. The rationale for setting a comparatively low threshold is thus to avoid excluding those who have completed secondary

[1] Variables differ in the extent to which they are provided. An individual's level of education is an example of a variable for which information can often be missing. Moreover, changes in classifications, e.g. in the coding of occupations, can cause problems in constructing a consistent coding scheme over longer periods of time.

school. On the other hand, the risk of including individuals in the sample who have not completed training appears limited since they have already been participating in training for at least two years and dropping out of such schemes can be expected to typically happen earlier. Second, it is required that any gaps between two apprenticeship episodes are no longer than 100 days. A possible reason for such breaks is that training also includes a coursework component which does not take place within the training company. No additional restrictions are imposed; in particular, changes in the training company, in the occupation or industry during the apprenticeship are disregarded because parts of the training should be sufficiently general so as to be transferable to a different company, occupation or industry.

In order to avoid any confounding effects of selected female labour-market participation, the sample is restricted to men. Moreover, the age range of the sample is homogenised to include only those between the age of 19 and 23 because the majority of graduates complete their training within this age range.[2] Applying this procedure yields a sample of 52,234 individuals[3] who have graduated between January 1999 and October 2012 and for whom transition into employment can be observed.

The model's dependent variable, *search*, is defined as the number of days it takes an individual to find employment after graduating from an apprenticeship programme. Figure 1 shows the distribution of this variable for the sample of individuals described above. The distribution is highly skewed as the majority (61%) falls into the category *No search*, which means that the employment episode of these individuals starts the day after graduation from apprenticeship training. Approximately 80% of graduates are able to find employment within 3 months after graduation, with this figure increasing to over 85% after 6 months.

Despite the large number of individuals who find employment directly upon graduation, this group is excluded from the empirical analysis. This is primarily due to technical reasons for the empirical model of Section 3.3 requires strictly positive durations. Moreover, zero and strictly positive durations may not be the outcomes of the same process. Instead firms may first

[2] SIAB only includes an individual's year of birth. Age at the time of graduation is defined as the difference between the year of graduation and the year of birth. Some individuals who are registered as being 25 upon graduation will therefore actually be between 24 and 26. Out of all male observations in SIAB with a completed apprenticeship for which all control variables are available 79% fall into the age range 19-23. As shown in the Supplementary Material, comparable results are obtained if this restriction is not imposed.

[3] The maximum number of observations that can be used in the empirical analysis decreases to 46,408 due to missing values for the covariates.

decide whether to offer an apprentice a position after graduation from the training programme, with this decision being based on the performance of apprentices during training as well as on the economic condition of the firm. Apprentices are then free to either accept or decline the offer. If no match between training firm and apprentice is reached, individuals enter the labour market and search for employment. The empirical analysis therefore models search duration conditional on an apprentice not having been directly employed by his training firm (or having found employment immediately at a different firm).

Figure 1: Distribution of search durations

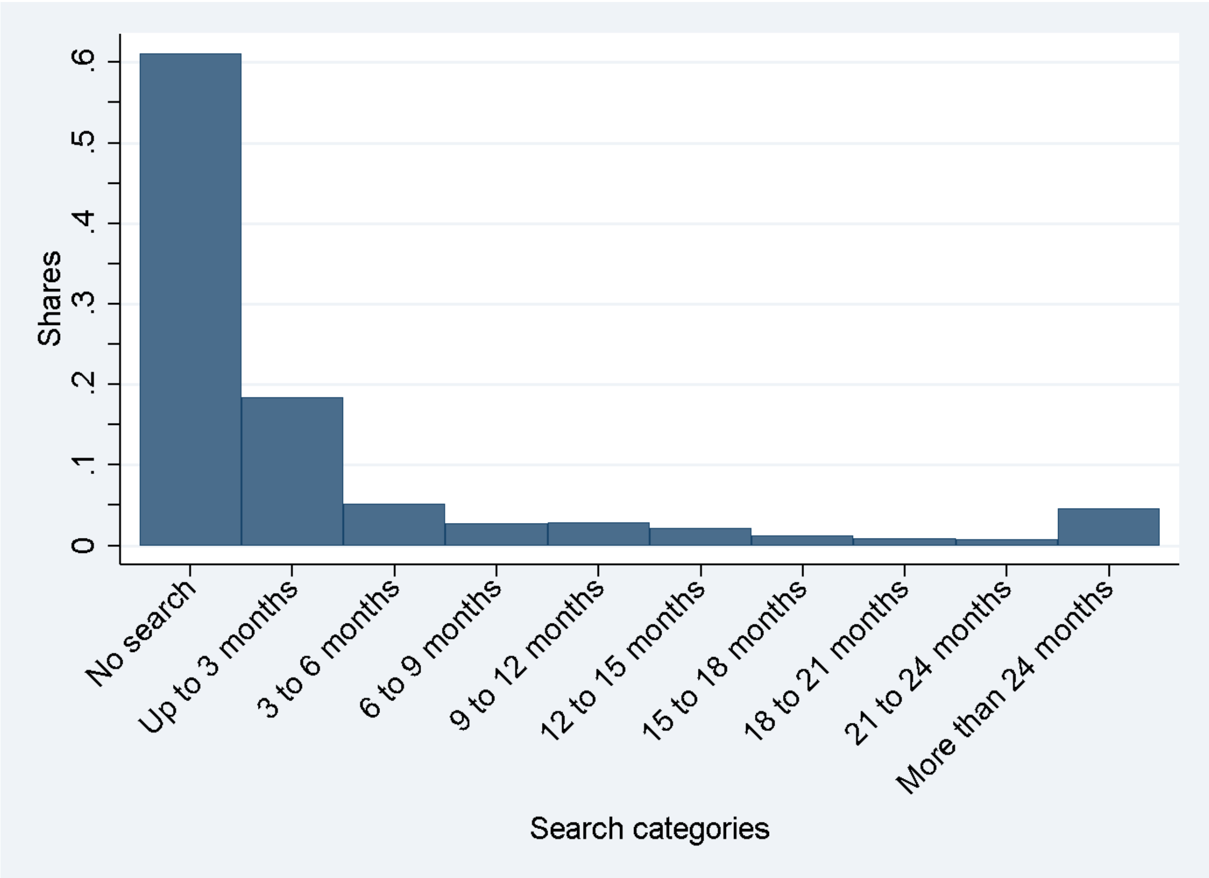

Source: Sample of Integrated Employment Biographies (author's calculations).

The obvious drawback of this approach is that those individuals who are employed directly might not constitute a random sample of graduates. In contrast, it is more likely that firms employ those apprentices which they believe to be especially productive. These individuals might possess characteristics which are not directly observable but which are relevant for on-the-job performance. If these characteristics also increased employability at other firms, graduates who are directly employed would be expected to experience shorter search periods in the counter-factual case of not being directly employed by their training firm.

Table 1 assesses this hypothesis by comparing average values of a number of characteristics between those apprentices who are employed directly and those who experience a strictly positive search duration. The first three variables refer to characteristics of the apprenticeship episode and while the difference in average duration of training and the share of Germans is statistically significant at the 0.01 level, the absolute difference in the variables is very small compared to the mean values of both groups. There is no statistically significant difference in average age at graduation. In contrast, there are sizeable and significant differences between characteristics of the employment spell that follows graduation: average daily earnings are about 20 Euro smaller for individuals who are not directly employed and the share of individuals working in part time is higher by about 15 percentage points. These latter findings suggest that both groups differ with respect to characteristics that are relevant for labour-market performance. Ideally, one would like to explicitly model this selection, but doing so would require an exogenous piece of information that would explain whether an individual is employed directly or experiences a positive search duration. In the absence of a suitable instrument, Section 4.4 provides an alternative way of including individuals with a zero search duration; the results of this analysis suggest that their inclusion reduces the magnitude of the cohort-size effect but does not affect its sign.

Table 1: Comparison of individuals with no and strictly positive search duration

Variable	**Observations**	**Group 1 (*search=0*)**	**Group 2 (*search>0*)**	**Mean difference**
Apprenticeship episode				
Duration of training	52,234	1,095.00	1,086.84	-8.16***
Age at graduation	52,234	21.16	21.17	0.01
German	52,226	0.97	0.95	-0.01***
Employment episode				
Average daily earnings	52,234	61.78	42.32	-19.46***
Part-time share	52,196	0.00	0.16	0.15***

Values derived from a regression on a group indicator as well as dummies for period and region of graduation. Robust standard errors are used. ***/**/* signifies significance at the .01/0.05/0.1 level. Differences in the number of observations are due to missing values of the corresponding variables.

The main explanatory variable, *cohort*, measures the regional supply of apprenticeship graduates and is based on the number of individuals that complete training within a given 6-month period and thus become available to the labour-market. Figure 2 shows the monthly number of graduates for the years 1999-2012. The annual distribution displays two peaks – one in January and another in June and July – which suggests that the bulk of apprentices complete training at two distinct points in time each year. To better reflect this pattern, the size of graduation cohorts is not computed for the whole year, but separately for two periods that cover six months each and that are centred on the peaks: November-April and May-October.

Figure 2: Monthly number of graduates (1999-2012)

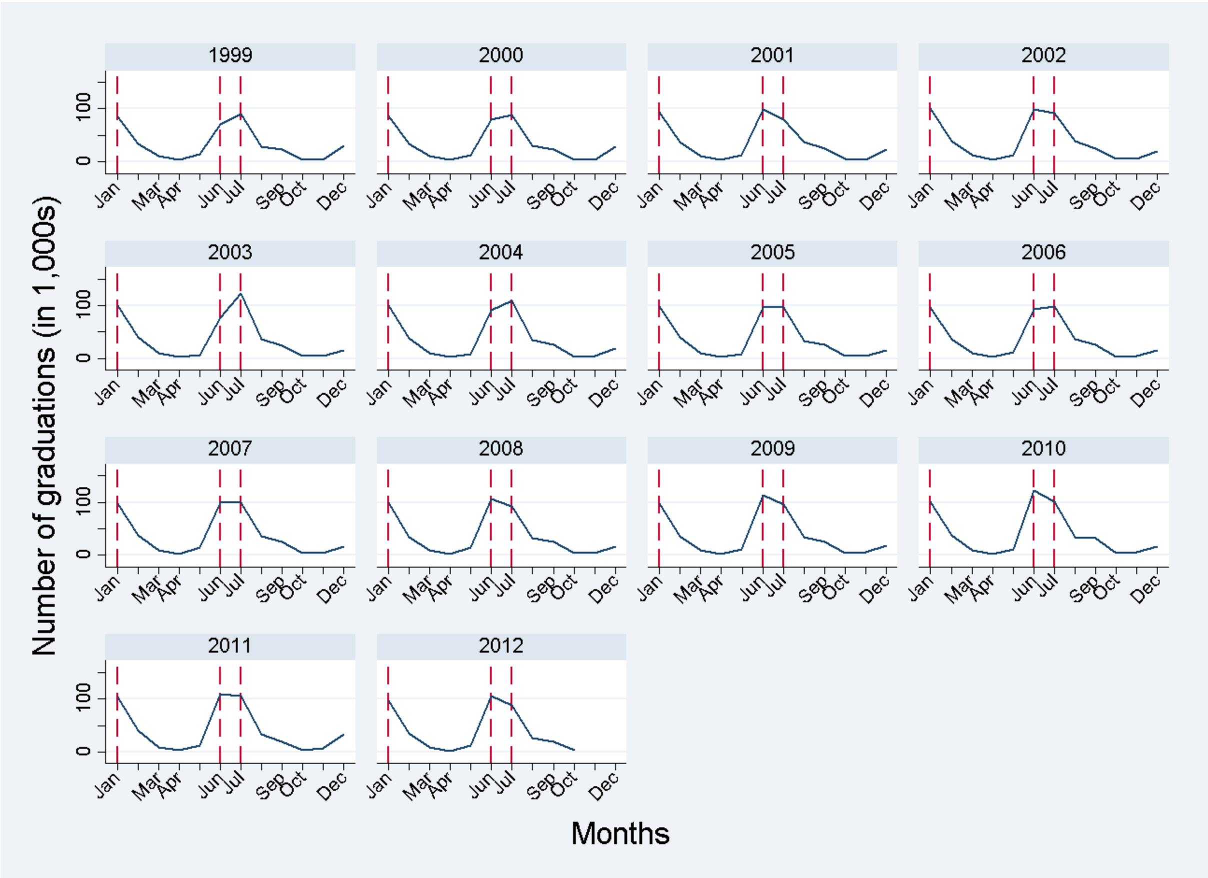

Source: Integrated Employment Biographies (author's calculations).

It is assumed that the duration of search for employment is influenced by the conditions of the labour market that an individual enters after graduation. The variable *cohort* measures the characteristic that is most relevant to this analysis: the degree of labour-market crowding among recently graduated apprentices. In order to avoid measurement error, graduation cohorts are constructed at the level of the 141 labour-market regions that are defined by Kosfeld and Werner (2012). As discussed in Section 2, these entities approximate self-contained units in which the employed population is exclusively recruited from the resident population. Since administrative units are typically not delineated according to economic criteria, they cannot be relied upon to provide an accurate measure of the size of a graduation cohort within an actual labour market. This argument is supported by findings of Garloff and Roth (2016) that the effect of cohort size on wages appears to be biased downwards when measured at the district level as compared to the level of labour-market regions.[4] Finally, to ensure comparability of the size of the graduation cohort across different labour-market regions, this quantity is standardised by total employment in the region.[5]

Additional control variables are given by dummy variables for an individual's age at graduation, for whether an individual is of German nationality, for the occupation of the apprenticeship and industry of the training firm[6] as well as the labour-market-specific unemployment rate. Summary statistics of these variables are given in the Appendix.

3.3 Model

To evaluate empirically the effect that the size of the graduation period has on an individual's search duration, the following Cox model is specified where subscripts *i*, *r* and *p* refer to the individual, the region and the period of graduation:

$$h_{irp}(t) = h_0(t)e^{(\gamma cohort_{rp} + \boldsymbol{\delta' x_{irp}})} \qquad [1]$$

Instead of formulating a relationship between the search duration and covariates, this model is specified in terms of the hazard rate $h_{irp}(t)$, which can be interpreted as the instantaneous probability that an individual realises a transition from search into employment. The term $h_0(t)$ represents the baseline hazard, i.e. the hypothetical hazard rate of an individual for whom all covariates are equal to zero. The Cox model belongs to the class of proportional hazard models meaning that changes in covariates shift the hazard rate up or down relative to the baseline hazard. The variable $cohort_{rp}$ captures the size of the graduation cohort in region *r* and period *p* relative to the number of employed individuals in that region. Sign and significance of the coefficient γ therefore provide the basis for assessing the effect that the size of the entry cohort has on the duration of job search. The vector x_{irp} contains the above-mentioned set of control variables as well as dummy variables for period and region of graduation. The coefficients of the model are derived by maximum partial likelihood estimation (MPLE).[7] To

[4] Labour-market regions refer to the individual's place of employment at the time of graduation. More than 80% of individuals in the sample live and work in the same region.
[5] For the first period (November-April) employment numbers refer to 31 March of the year, while it is 31 October for the second period (May-October).
[6] Occupation indicators are derived from the coding scheme *Klassifikation der Berufe 2010*, while industry indicators are based on the *Klassifikation der Wirtschaftszweige 1993*. Details are provided in the Appendix.
[7] The term partial refers to the fact that in contrast to fully parametric models, information on the search durations themselves is not used in the estimation. Instead, the relationship between the hazard rate and the covariates is derived solely from the ordering of the search durations.

account for the difference in the level of aggregation of the dependent variable and the cohort variable, standard errors are clustered at the level of the labour-market region.

Four different specifications of this model are estimated which differ with respect to the specified period of time during which transitions into employment are observed. An inherent asymmetry in the data is given by the fact that individuals that complete their apprenticeship training earlier can be observed for a longer period of time (up to 31 December 2014) and as such can also accumulate longer search durations. To ensure comparability between graduates from different periods, four common periods of observation following graduation are defined: 3 months, 6 months, 1 year and 2 years. Individuals that find employment after the end of the common observation period are treated as not having realised a transition (i.e. they are right-censored) and their search durations are set equal to the corresponding common period of observation.[8]

To consistently estimate the effect that the size of the graduation cohort has on the duration of search, it has to be assumed that individuals did not systematically select a region in which to undertake the apprenticeship training on the basis of their expectations regarding the probability of finding employment upon graduation. While the absence of regional selection appears unlikely in the context of other studies on cohort-size effects (see Moffat and Roth, 2016a), it is argued that this possibility is less of a concern in this case. First, since individuals are typically young when they start training, the region in which an apprenticeship is being undertaken will usually be determined by the region they live in at that time. Second, it appears unlikely that reliable expectations can be formed about the economic conditions prevailing in a region at the time of graduation. Moreover, if self-selection occurs into regions that constantly provide better employment opportunities for apprentices (and hence shorter search durations), this effect would be captured by the region dummies.

[8] The empirical model is based on two pieces of information: an indicator for whether transition into employment took place and the number of days an individual survived before transition. In the case of a 3-month period of observation an individual who found employment after six months would be recorded as not having experienced transition and his duration of search would be set to three months. Right-censored observations are not dropped from the regression. While they are treated as not having experienced transition, they are included in the 'risk set', i.e. the set of observations that are at risk of realising a transition into employment at each of the recorded transition times. The share of right-censored observations is 53% (3-month period), 39% (6-month), 25% (1-year) and 11% (2-year), respectively.

4 Results

4.1 Baseline results

Table 2 contains the coefficients from estimating the model of Equation 1 for each of the four common observation periods (3-month, 6-month, 1-year and 2-year periods).

Table 2: Regression results

	3 months	6 months	1 year	2 years
Cohort	26.04**	15.16	11.81	5.79
	(11.20)	(9.33)	(9.32)	(8.57)
Unemployment rate	-1.99*	-0.95	-1.50*	-1.31*
	(1.09)	(0.94)	(0.84)	(0.77)
German	0.08	-0.00	-0.07	-0.07*
	(0.06)	(0.05)	(0.05)	(0.04)
Age				
20	-0.04	-0.04	-0.05*	-0.06
	(0.04)	(0.03)	(0.03)	(0.03)*
21	-0.06*	-0.05*	-0.06**	-0.05
	(0.03)	(0.03)	(0.03)	(0.03)
22	-0.13***	-0.11***	-0.12***	-0.09***
	(0.04)	(0.04)	(0.03)	(0.03)
23	-0.06	-0.04	-0.03	-0.03
	(0.05)	(0.04)	(0.04)	(0.04)
Dummies				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-81,452.15	-103,503.43	-126,212.95	-145,330.79
Observations	18,133	18,133	18,133	18,133
Clusters	141	141	141	141
ME(std)	1.07**	1.04	1.03	1.02

All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Breslow method is used to handle tied observations. *ME(std)* shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.

In each case the estimated coefficient for the size of the entry cohort is positive though it is significant only if transitions into employment are counted as such if they take place during the first three months following graduation. For this observation period an increase in the size of the entry cohort significantly increases the hazard rate of finding employment. This means that individuals who complete their apprenticeship training as part of a larger cohort have shorter search durations. As presented in further detail in the Supplementary Material, the finding that belonging to a larger entry cohort is associated with shorter search durations is robust to a number of changes in the sample as well as in the empirical model.

To assess the size of the estimated effects, hazard ratios are computed, which show the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation. This value, which is shown in the bottom row of Table 2, is given by the exponentiated product of the cohort coefficient and the corresponding standard deviation (see Table A1). For the 3-month period such a change is predicted to increase the hazard rate by about 8%. Performing a similar computation for the regional unemployment rate at the time of graduation shows that the effects of both variables are of similar size (though opposite sign) as the hazard rate is predicted to fall by approximately 10% if the unemployment rate increased by one standard deviation.

An alternative way to illustrate the size of the estimated effect is by means of the survivor function, which shows how the share of individuals that have not yet found employment changes with the duration of search. Figure 3 plots the survivor function for the 3-month period and for different values of the entry cohort: the solid line corresponds to the case in which all explanatory variables are equal to zero (for the cohort-size variable and the unemployment rate this implies that they are equal to their mean), while the dashed line above (below) shows the survivor function when the entry cohort is smaller (larger) by one standard deviation. Naturally, the survivor function is decreasing as the share of graduates finding employment increases with time; at the end of the observation period, between 50% and 55% of graduates have taken up employment. In line with the finding that larger entry cohorts increase the hazard rate of finding a job, Figure 3 shows that the share of survivors is generally smaller for larger cohorts. After 90 days the survivor function takes on a value of 52% when all variables are equal to zero, with the corresponding value equal to 54% (50%) when the size of the entry cohort is smaller (larger) by one standard deviation. A change in the relative number of apprenticeship graduates by one standard deviation is therefore predicted to change the share of individuals who are still searching after 90 days by 2 percentage points or, equivalently, by approximately 4%.

Another feature of the results presented in Table 2 is that the size of the cohort coefficient decreases as the period of observation is extended and the estimation results become increasingly affected by graduates with longer search durations (while they are always included in the sample, censored observations only contribute to the estimation by belonging to the set of individuals that are at risk of transition into employment). On the one hand this finding could be seen as evidence that the effect that the size of the entry cohort has on subsequent search durations is not persistent. On the other hand, it is conceivable that for those individuals that require more time to find employment current labour-market conditions matter in addition to the conditions prevailing at the time of graduation. To assess this hypothesis, the model of Equation1 is supplemented with measures of cohort size and the unemployment rate that refer to later points in time: 6 months after graduation in the case of the 6-month observation period, 6 and 12 months for the 1-year period as well as 6, 12, 18 and 24 months for the 2-year period.

Figure 3: Survivor function (estimated at different values of the graduation cohort)

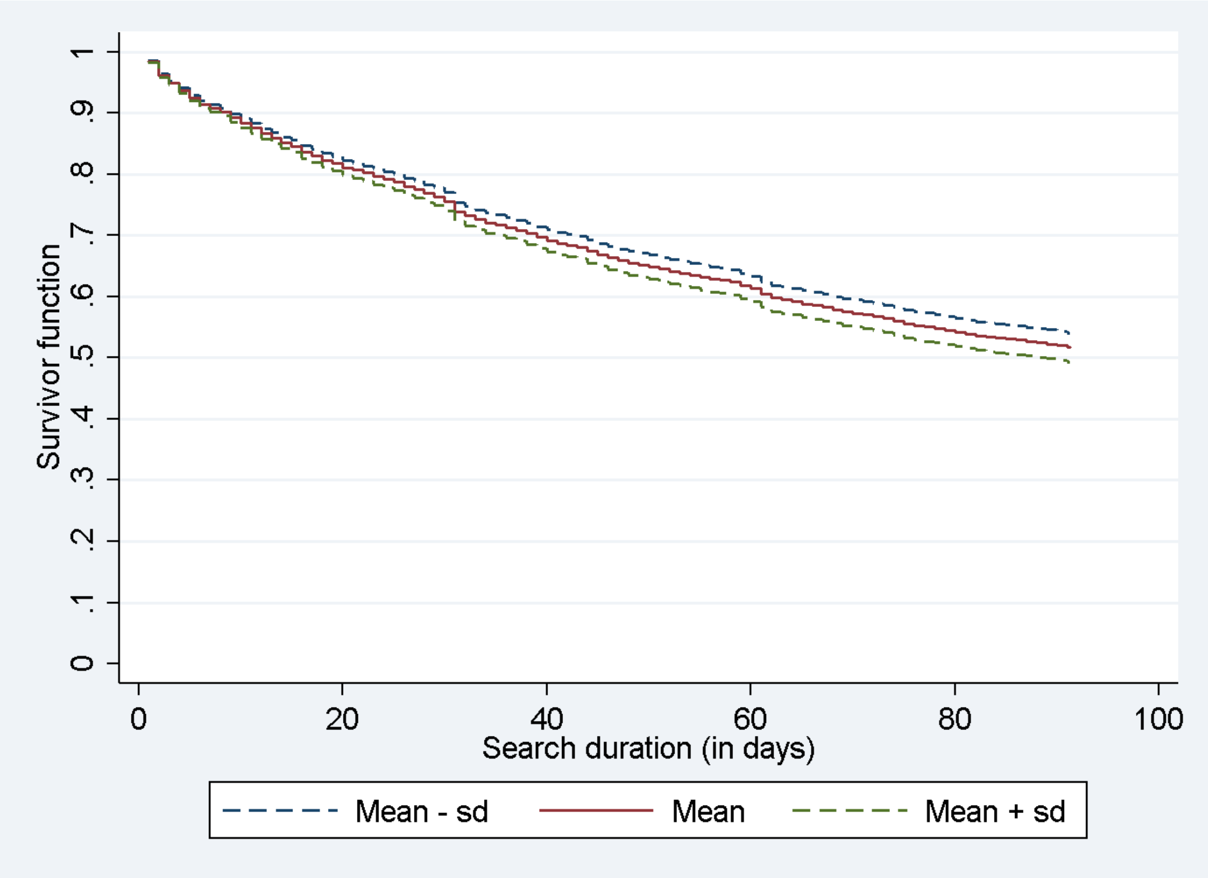

Source: Sample of Integrated Employment Biographies and Integrated Employment Biographies (author's calculations).

Table 3 shows that once these measures for the current labour-market conditions are added, the cohort coefficient in the 6-month period is almost identical in terms of size and significance to the corresponding effect that is measured when only transitions occurring within three months after graduation are treated as such. This suggests that in this case the positive effect of the size of the entry condition on the hazard rate of finding employment continues to exist once current conditions are controlled for. Similar results are, however, not obtained for the two remaining observation periods as, first, the effect of the entry cohort decreases in magnitude and, second, all of the cohort coefficients are individually insignificant, though in the case of the 2-year period they remain jointly significant at the 5% level.

Table 3: Regression results (when current labour-market conditions are controlled for)

	3 months	6 months	1 year	2 years
Cohort	26.04**	26.78**	2.29	-3.94
	(11.20)	(12.77)	(19.63)	(18.97)
Cohort (+6 months)	-	-1.00	-7.54	1.37
		(12.08)	(11.70)	(19.51)
Cohort (+12 months)	-	-	23.30	19.77
			(20.48)	(22.49)
Cohort (+18 months)	-	-	-	-3.42
				(19.70)
Cohort (+24 months)	-	-	-	8.94
				(22.36)
Control variables	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-81,452.15	-103,498.96	-126,203.29	-138,947.97
Observations	18,133	18,133	18,133	17,466
Clusters	141	141	141	141
ME(std)	1.07**	1.07**	1.01	0.99

Variables refer to the time of graduation from apprenticeship training unless indicated otherwise. The set of control variables also includes current values of the unemployment rate. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Breslow method is used to handle tied observations. *ME(std)* shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.

4.2 Discussion of the hypotheses

The results of Table 2 provide no support for the cohort-crowding hypothesis that members of larger entry cohorts have longer search durations; on the contrary, the empirical evidence suggests that graduating as part of a larger group reduces the time required to find a job. A possible explanation for this relationship, as discussed in Section 2, is that in the face of increased competition graduates from larger cohorts choose to take up lower-quality jobs. If entering the labour market as part of a larger group indeed pushes apprentices into jobs that do not match their qualifications, characteristics of the first employment spell should differ between graduates from large and from small cohorts. This hypothesis is assessed by means of two outcome variables: the natural logarithm of the average daily wage earned in the first employment spell and an indicator for whether this spell refers to regular employment subject to social-security contributions.[9] These variables are regressed on the size of the entry cohort as well as on the set of control variables used in the estimation of Equation 1.

[9] The smaller number of observations in the top panel is due to some individuals being assigned wages of a value zero. For approximately 76% of observations with a strictly positive search duration the first employment spell is of the regular type, with 18% being registered as working in marginal employment (*Geringfügige Beschäftigung*) and 5% having started a new apprenticeship.

Table 4 contains the estimated cohort-size coefficients for the full set of observations as well as separately for those individuals that fall into each of the four periods of observation. The top panel reports the results pertaining to average daily wages, while the effects on the probability of being in non-regular employment are recorded in the bottom half. The results do not support the hypothesis that graduates from larger cohorts are pushed into lower-quality jobs. If anything, the findings suggest that the size of the graduation cohort is positively associated with the wages earned in the first job as well as with the probability of being in regular employment, though none of the estimated coefficients is statistically different from zero.

Table 4: Cohort-size effects on wages and regular employment status

	All	**3 months**	**6 months**	**1 year**	**2 years**
Ln(average daily wage)					
Cohort	2.57	11.61	7.19	7.29	1.40
	(6.19)	(9.42)	(8.15)	(7.40)	(6.69)
Control variables	Yes	Yes	Yes	Yes	Yes
Observations	17,995	8,567	10,985	13,586	15,938
Clusters	141	141	141	141	141
Indicator for regular employment					
Cohort	2.91	5.60	3.87	3.42	0.98
	(3.04)	(4.17)	(4.09)	(3.55)	(2.96)
Control variables	Yes	Yes	Yes	Yes	Yes
Observations	18,133	8,605	11,056	13,685	16,057
Clusters	141	141	141	141	141

Estimation by ordinary least squares (OLS). All variables refer to the time of graduation from apprenticeship training. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region.

The positive impact of cohort size on the hazard rate of finding employment and the lack of evidence in support of the hypothesis that graduating as part of a large group drives apprentices into lower-quality jobs leaves the possibility that firms react to changes in the number of apprenticeship completers by creating jobs, though this effect appears to be restricted to a relatively short period after graduation. This explanation would appear to challenge findings by Garloff et al. (2013) whose empirical analysis for West German labour-market regions shows that larger cohorts increase the overall unemployment rate.[10] If firms creating jobs in expectation of large entry cohorts is indeed the explanation for the finding that members of larger cohorts have shorter search durations, the empirical evidence

[10] Their use of the overall unemployment rate as the dependent variable and the focus on the share of young individuals aged between 15 and 24 rather than on the number of graduates from an apprenticeship programme, however, limit the comparability to this paper.

presented in Tables 2 and 3 suggests that these beneficial effects are limited to a period of about six months following graduation.

4.3 Alternative explanations

Two alternative explanations for the findings should be considered which address the role of regional selection following graduation and changes in the composition of the group of graduates. First, the positive relationship between the hazard rate and the size of the graduation cohort could be spurious if it is driven by apprentices that graduate as part of a large group choose to search for employment in regions where search durations are shorter. In the sample of graduates with strictly positive search durations approximately 31% of individuals register their first employment spell in a different region to the one in which they have graduated. If belonging to a large entry cohort induces some individuals to search for employment elsewhere, the size of the graduation cohort and the probability of finding employment in a different region should be positively related.

Table 5 shows the results from regressing a binary dependent variable that takes the value 1 if the region of an individual's first employment spell is not the same as the one in which he graduated on the set of explanatory variables that are used in the estimation of Equation 1. The cohort-size coefficients, however, provide no evidence for the hypothesis that selecting into another region after graduation is the reason for the results of Table 2 as none of the estimated effects are significantly different from zero.[11]

Table 5: Cohort-size effects on the probability of finding employment in a different region

Indicator for employment in a different region	**All**	**3 months**	**6 months**	**1 year**	**2 years**
Cohort	1.56	-1.73	0.07	0.60	1.70
	(3.08)	(4.88)	(3.87)	(3.26)	(3.13)
Control variables	Yes	Yes	Yes	Yes	Yes
Observations	18,133	8,605	11,056	13,685	16,057
Clusters	141	141	141	141	141

Estimation by ordinary least squares (OLS). All variables refer to the time of graduation from apprenticeship training. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region.

Second, the composition in terms of productivity may differ between small and large graduation cohorts. If the number of graduates that are employed directly is fixed, some highly productive individuals will have to engage in job search if they belong to a larger cohort. In

[11] Similar conclusions can be drawn from estimating a logit model instead of a linear probability model. The results are available upon request.

such a scenario the fact that search durations are shorter in larger cohorts might be the result of a change in the productivity composition of the cohort towards more individuals with a higher level of productivity. More productive graduates are likely to find employment and to require less time to do so, which might explain the positive cohort-size coefficients, especially shortly after graduation.

This hypothesis is assessed by estimating the effect of the size of the entry cohort on the probability of having a strictly positive search duration. If the above argument is correct, belonging to a larger cohort should be associated with a higher probability of having to search for employment. Table 6 shows the results from regressing a binary indicator for whether an individual has to search on the same set of explanatory variables as used in Equation 1.[12] Compared to Table 2 the number of observations increases because those individuals with a zero search duration are now also included. The coefficient of the entry cohort is significant only at the 10% level and suggests that belonging to a larger group of graduates reduces the probability of having a positive search duration. This effect, however, appears to be small with an increase in the size of the graduation cohort by one standard deviation being predicted to increase the probability of search by one percentage point compared to a mean value of the dependent variable of 0.39. The hypothesis that the results of Table 2 reflect a change in the productivity composition of the group of individuals in larger graduation cohorts that have to engage in search is therefore not supported by the data.

Table 6: Cohort-size effects on the probability of having a strictly positive search duration

Indicator for having a strictly positive search duration	**All**
Cohort	-4.01
	(2.31)
Control variables	Yes
Observations	46,408
Clusters	141

Estimation by ordinary least squares (OLS). All variables refer to the time of graduation from apprenticeship training. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The smaller number of observations in the top panel is due to some individuals being assigned wages of a value zero.

4.4 Inclusion of individuals with zero search duration

The use of survival models prevents the inclusion of individuals that are employed upon graduation and therefore have a zero search duration. As discussed in Section 3.2, omitting this set of observations potentially raises a problem of sample selection if the two groups of

[12] Similar conclusions can be drawn from estimating a logit model instead of a linear probability model. The results are available upon request.

individuals differ in terms of unobserved characteristics which in turn may have an effect on their employability. In order to assess the impact of this selection on the estimated effect of cohort size, the search-duration variable is adjusted by adding 1 to each value (and adjusting the censoring variables accordingly). Doing so allows the inclusion of those individuals for whom search duration is actually zero in the estimation of the Cox model. The results of this analysis are presented in Table 7.

Table 7: Regression results (when individuals with zero search durations are included)

	3 months	6 months	1 year	2 years
Cohort	11.58***	7.82*	7.04	4.35
	(4.46)	(4.33)	(4.92)	(4.77)
Unemployment rate	-2.46***	-2.12***	-2.24***	-2.10***
	(0.38)	(0.38)	(0.41)	(0.40)
Control variables	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-384,802.52	-406,547.02	-429,408.37	-448,434.23
Observations	46,408	46,408	46,408	46,408
Clusters	141	141	141	141
ME(std)	1.03***	1.02*	1.02	1.01

All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Breslow method is used to handle tied observations. *ME(std)* shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.

In terms of their pattern the estimated coefficients are comparable to the results of Table 2: the coefficients are positive and decrease in size as the observation period becomes longer; moreover, the effects are significant for the 3-month period, but also for the 6-month period. The main difference is that the coefficients are smaller, suggesting that once those individuals are included that are employed directly upon graduation, the strength of the relationship between the size of an individual's graduation cohort and the duration of her search for employment is reduced. A possible explanation for the weaker relationship between cohort size and search duration is that a number of graduates will always be employed directly regardless of the size of their graduation cohort. This explanation is in line with the results of Table 6, which show that the probability of having to search (i.e. of not becoming employed directly) is only marginally affected by the number of apprentices completing training.[13]

[13] An alternative way of including observations with zero search durations is to estimate count-data models. The results of these models, which are available upon request, also providence evidence that members of larger graduation cohorts have shorter search durations.

5 Conclusion

How the size of the cohort that an individual belongs to affects his contemporaneous labour-market outcomes constitutes a widely analysed field of research, with particular attention being paid to the effects on wages as well as employment and unemployment. In contrast, how cohort size measured at a specific point in an individual's career affects future outcomes has so far not attracted a large amount of attention, while there has recently been a substantial amount of research on the effect of the state of the business cycle at the time of labour-market entry on an individual's subsequent wages and employment opportunities. The contribution of this paper to the cohort-size literature is to analyse the effect on the amount of time an individual spends searching for employment after entering the labour market, which represents an outcome that has so far not been addressed. Moreover, in doing so, this paper conceptualises the size of the cohort as a factor affecting the conditions under which an individual's entry to the labour market takes place rather than as a contemporaneous explanatory variable.

From a theoretical perspective the relationship between the size of the entry cohort and an individual's subsequent duration of search can take various forms. Longer durations would be expected if increased competition makes it harder for members of larger cohorts to find employment – a relationship that would be in line with the standard cohort-crowding hypothesis. Individuals, however, may counteract this effect if they are willing to downgrade by taking up employment in a lower-quality job. Finally, if large cohorts indeed lead to lower unemployment rates, as has been argued by parts of the cohort-size literature, a negative impact on search durations is also conceivable. As such the results of the analysis may not only shed light on the relationship between cohort size at the point of labour-market entry and the subsequent duration of search, but may also provide insights into the former's effect on employment and unemployment outcomes. Since economic theory does not provide a clear indication on the nature of the relationship, the above hypotheses are assessed by means of an empirical analysis. The sample is based on register data and consists of graduates from Germany's apprenticeship programme who completed their training between January 1999 and October 2012.

The results of the empirical analysis suggest that the hazard rate of finding employment increases with the size of the cohort as part of which an individual graduates and enters the labour market. While this effect appears to apply only to individuals that find employment within a relatively short period of three months following graduation, once contemporaneous economic conditions are controlled for this effect is also found for individuals who take up a job within six months of graduating. Overall, the empirical analysis provides no evidence to suggest that members of larger entry cohorts suffer depressed labour-market outcomes in terms of longer search durations. Further analyses show that shorter search durations among members of larger graduation cohorts are not associated with employment in lower-quality jobs as there is no empirical evidence for a negative effect of cohort size on wages or on the probability of finding regular employment. The possibility that the observed effects are driven by either selection into regions with better employment opportunities (and hence shorter search durations) after graduation or changes in the productivity composition of those graduates that have to search for employment is also not supported by the data. Finally, the fact that those apprentices who find employment directly upon graduation cannot be included in the baseline results does not appear to materially affect the conclusions regarding the nature of the relationship between cohort size and search duration.

A possible explanation for the positive effect of the size of the entry cohort on the hazard rate of finding employment is that firms anticipate such changes in the supply of young workers and react by creating jobs which in turn causes a shorter duration of search. Such an interpretation would be compatible with the view that larger cohorts also lead to lower unemployment rates, which potentially challenges existing evidence for (Western) Germany that larger cohorts are associated with higher unemployment rates. In light of the demographic processes that are projected to lead to a lower share of young age groups in the population and of a rising preference for tertiary education, future cohorts of apprenticeship graduates may be expected to decrease in size, which, at least according to this analysis, would suggest that search durations might become longer in the future.

References

Berger, M.C. (1983) Changes in Labor Force Composition and Male Earnings: A Production Approach, *Journal of Human Resources,* **18**, 177-96.

Biagi, F. and Lucifora, C. (2008) Demographic and education effects on unemployment in Europe, *Labour Economics,* **15**, 1076-101.

Brunello, G. (2010) The effects of cohort size on European earnings, *Journal of Population Economics*, **23**, 273-90.

Brunner, B. and Kuhn, A. (2014) The impact of labor market entry conditions on initial job assignments and wages, *Journal of Population Economics,* **27**, 705-38.

Card, D. and Lemieux, T. (2001) Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis, *Quarterly Journal of Economics*, **116**, 705-46.

Cockx, B. and Ghirelli, C. (2016) Scars of recessions in a rigid labor market, *Labour Economics,* **41***,* 162-76.

Dooley, M. (1986) The Overeducated Canadian? Changes in the Relationship among Earnings, Education, and Age for Canadian Men: 1971-81, *The Canadian Journal of Economics/Revue canadienne d'Economique,* **19**, 142-59.

European Commission (2014) *The 2015 Ageing Report*, European Commission, Brussels.

Freeman, R.B. (1979) The effect of demographic factors on age-earnings profiles, *Journal of Human Resources*, **14**, 289-318.

Garloff, A., Pohl, C. and Schanne, N. (2013) Do small labor market entry cohorts reduce unemployment?, *Demographic Research*, **29**, 379-406.

Garloff, A. and Roth, D. (2016) Regional population structure and young workers' wages, IAB Discussion Paper No. 06/2016, Institute for Employment Research, Nuremberg.

Kahn, L.B. (2010) The long-term labor market consequences of graduating from college in a bad economy, *Labour Economics,* **17**, 303-16.

Korenman, S. and Neumark, D. (2000) Cohort Crowding and Youth Labor Markets: A Cross-National Analysis, in D.G. Blanchflower and R.B. Freeman (eds.) *Youth Employment and Joblessness in Advanced Countries*, University of Chicago Press, Chicago.

Kosfeld, R. and Werner, A. (2012) Deutsche Arbeitsmarktregionen – Neuabgrenzung nach den Kreisgebietsreformen 2007-2011, *Raumforschung und Raumordnung*, **70**, 49-64.

Michaelis, J. and Debus, M. (2011) Wage and (un-)employment effects of an ageing workforce, *Journal of Population Economics*, **24**, 1493-511.

Moffat, J. and Roth, D. (2016a) The Cohort Size-Wage Relationship in Europe, *Labour*, forthcoming.

Moffat, J. and Roth, D. (2016b) Cohort size and youth labour-market outcomes: the role of measurement error, IAB Discussion Paper No. 37/2016, Institute for Employment Research, Nuremberg.

Morin, L.-P. (2015) Cohort size and youth earnings: Evidence from a quasi-experiment, *Labour Economics,* **32**, 99-111.

Mosca, I. (2009) Population Ageing and the Labour Market, *Labour, **23***, 371-95.

Shimer, R. (2001) The Impact of Young Workers on the Aggregate Labor Market, *The Quarterly Journal of Economics*, **116**, 969-1007.

Skans, O.N. (2005) Age effects in Swedish local labor markets, *Economics Letters,* **86**, 419-26.

Statistisches Bundesamt (2015) *Bevölkerung Deutschlands bis 2060: 13. koordinierte Bevölkerungsvorausberechnung*, Statistisches Bundesamt, Wiesbaden.

Stevens, K. (2007) Adult Labour Market Outcomes: the Role of Economic Conditions at Entry into the Labour Market, available at http://www.iza.org/conference_files/SUMS2007/stevens_k3362.pdf

Welch, F. (1979) Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust, *Journal of Political Economy*, **87**, S65-S97.

Wright, R.E. (1991) Cohort size and earnings in Great Britain, *Journal of Population Economics*, **4**, 295-305.

Appendix

Table A1: Descriptive statistics

Variable	Observations	Mean	Standard deviation	Minimum	Maximum
Search (3 months)	18,133	63.716	33.848	1	91.250
Search (6 months)	18,133	104.494	72.576	1	182.500
Search (1 year)	18,133	162.522	144.184	1	365
Search (2 years)	18,133	222.391	246.219	1	730
Cohort	18,133	0.009	0.003	0.004	0.020
Unemployment rate	18,133	0.120	0.053	0.024	0.317
German	18,133	0.954	0.209	0	1
Age					
19	18,133	0.144	0.351	0	1
20	18,133	0.292	0.455	0	1
21	18,133	0.278	0.448	0	1
22	18,133	0.175	0.380	0	1
23	18,133	0.111	0.314	0	1
Occupations					
1	18,133	0.041	0.199	0	1
2	18,133	0.495	0.500	0	1
3	18,133	0.223	0.416	0	1
4	18,133	0.023	0.148	0	1
5	18,133	0.034	0.181	0	1
6	18,133	0.084	0.277	0	1
7	18,133	0.068	0.251	0	1
8	18,133	0.023	0.149	0	1
9	18,133	0.010	0.101	0	1
Industries					
1	18,133	0.028	0.166	0	1
2	18,133	0.000	0.018	0	1
3	18,133	0.005	0.067	0	1
4	18,133	0.196	0.397	0	1
5	18,133	0.005	0.069	0	1
6	18,133	0.214	0.410	0	1
7	18,133	0.208	0.406	0	1
8	18,133	0.056	0.229	0	1
9	18,133	0.026	0.158	0	1
10	18,133	0.011	0.106	0	1
11	18,133	0.041	0.197	0	1
12	18,133	0.021	0.144	0	1
13	18,133	0.122	0.327	0	1
14	18,133	0.032	0.176	0	1
15	18,133	0.036	0.185	0	1
16	18,133	0.000	0.013	0	1
17	18,133	0.000	0.011	0	1

Table A2: Classification of occupations and industries

Occupations	
1	Agriculture, forestry, farming and gardening
2	Production of raw materials and goods, and manufacturing
3	Construction, architecture, surveying and technical building services
4	Natural sciences, geography and informatics
5	Traffic, logistics, safety and security
6	Commercial services, trading, sales, the hotel business and tourism
7	Business organisation, accounting, law and administration
8	Health care, the social sector, teaching and education
9	Philology, literature, humanities, social sciences, economics, media, art, culture, and design
Industries	
1	Agriculture and forestry
2	Fishery
3	Mining and quarrying
4	Manufacturing
5	Electricity and water supply
6	Construction
7	Sale, maintenance and repair
8	Tourism
9	Transport
10	Financial and insurance services
11	Real estate
12	Public administration and defence
13	Education
14	Health and social work
15	Other services
16	Households
17	Extraterritorial organisations

The occupation and industry classifications are derived from the *Klassifikation der Berufe 2010* and the *Wirtschaftszweige 1993*, respectively.

Supplementary material

This document's purpose is to assess the validity of the paper's empirical model as well as the robustness of the results. This is done by first performing a set of residual-based tests concerning the specification of the Cox model. Second, the sensitivity of the results is analysed by estimating different variations of the initial Cox model as well as by presenting the results from different fully parametric specifications and comparing them with those of the paper's semi-parametric Cox model.

S1 Model specification

The Cox model gives rise to different kinds of residuals which form the basis for testing the adequacy of the specified model (for further detail on the tests and the different residuals see Box-Steffensmeier and Jones, 2004 or Cleves et al., 2010). The first specification test assesses how well the model proposed in Equation 1 of the paper fits the data (this is done separately for each of the common observation periods of 3 months, 6 months, 1 year and 2 years, respectively). This test is based on the Cox-Snell residuals, which can be derived from the estimated coefficients and the estimated cumulative baseline hazard rate and which can be interpreted as the number of transitions that an individual is expected to have experienced (assuming he can repeatedly experience transitions) within the time that it actually takes the individual to find employment. If the model is correctly specified, these residuals will follow a unit-exponential distribution with a hazard rate that is equal to 1. To assess whether this is the case, the cumulative hazard function of the Cox-Snell residuals is estimated (using the Nelson-Aalen estimator) and is then plotted against the residuals. For the correct model specification this estimate will be close to the 45-degree line.

Figure S1 shows the estimated cumulative hazard functions of the Cox-Snell residuals for each observation period against the 45-degree line. In each case the estimate lies close to this line, which indicates that the model provides a reasonable fit. While deviations from the 45-degree line can be found for higher values of the Cox-Snell residuals, this is likely to reflect the fact that the number of individuals for whom many transitions are expected will be relatively small.

Figure S2: Overall model fit (Cox-Snell residuals)

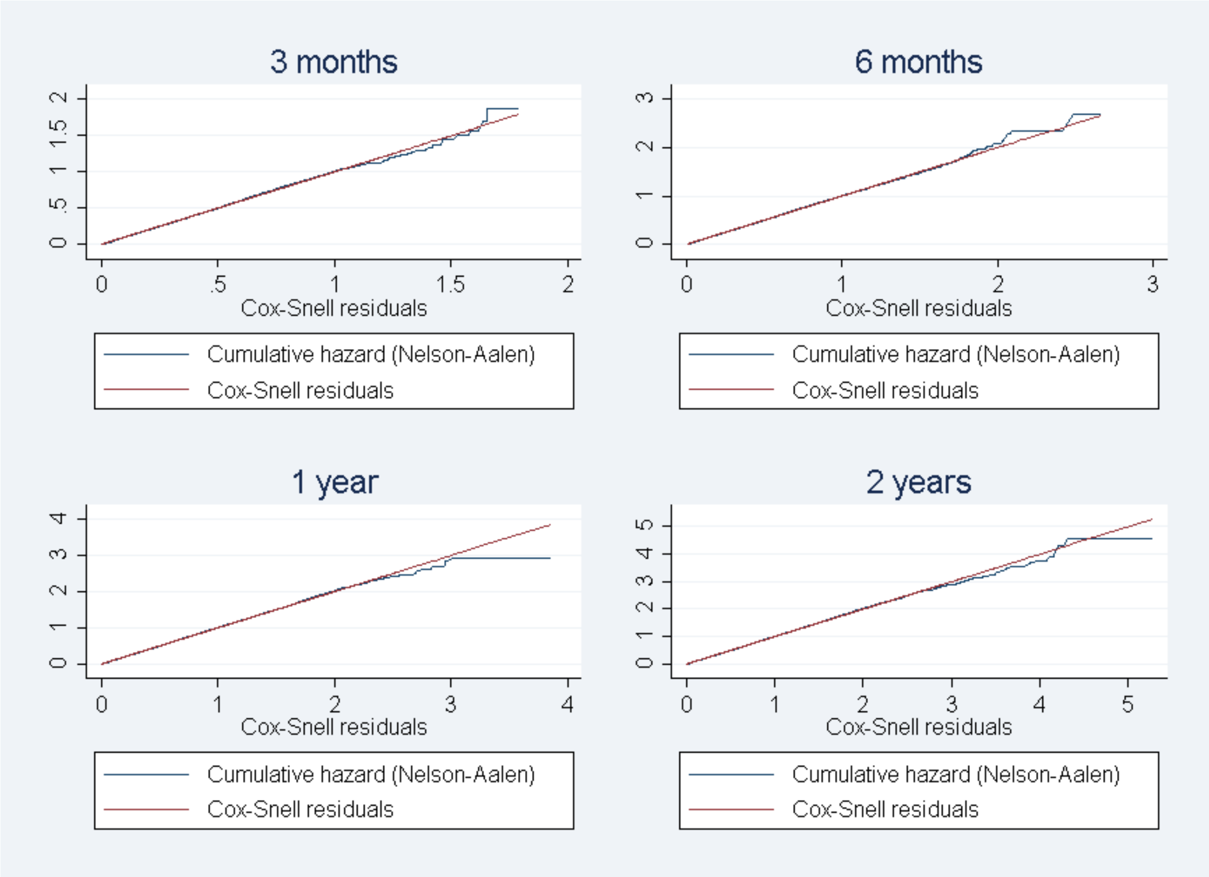

Source: Integrated Employment Biographies and Sample of Integrated Employment Biographies (author's calculations).

Next, it is tested to what extent the coefficient of the cohort-size variable is driven by any single observation. Instead of estimating the model separately after successively excluding one observation, the effect on the coefficient is approximated by multiplying the matrix of score residuals with the variance-covariance matrix of the estimated coefficients. As shown in Figure S2, the largest (absolute) change in the cohort-size coefficient is about 0.5 units when a single observation is dropped, which is a small change in light of an estimated coefficient that ranges between 26 (3-month period) and 6 (2-year period).

Figure S3: Influential observations (Score residuals)

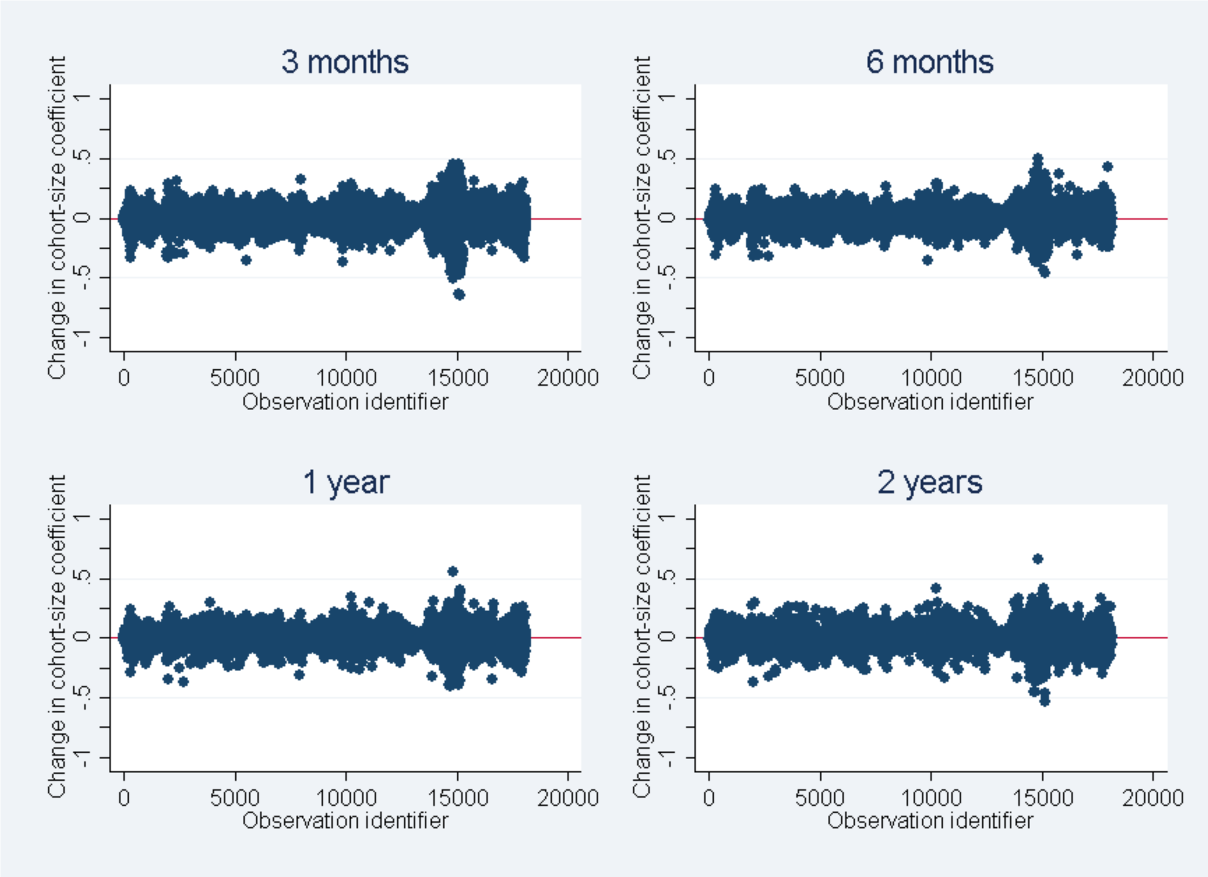

Source: Integrated Employment Biographies and Sample of Integrated Employment Biographies (author's calculations).

The final test concerns the proportional-hazards property. The Cox model belongs to a class of models for which the hazard rate can be decomposed into one component that depends only on time (i.e. the duration of search) and which is given by the baseline hazard $h_0(t)$ and another component which is a function of the model's covariates and their coefficients given by the term $exp(\delta' x_{irp})$. Changes in the covariates are therefore expected to shift the baseline hazard up (or down) in a parallel way. This assumption is testable on the basis of yet another type of residuals, the Schoenfeld residuals. The former are observation-specific and covariate-specific and can be interpreted as the difference between the observed and the expected value of a covariate. The test can be performed globally (i.e. for the whole set of covariates) as well as locally for individual regressors. Table S1 shows the test-statistics and p-values associated with the cohort-size variable as well as with the whole set of regressors. Looking at the 3-month period it can be seen that the null hypothesis of the proportional-hazards assumption being satisfied is not rejected locally for the cohort-size variable. However, the test statistic for the global test is sufficiently large that the null hypothesis is rejected at the 5%-level. For all other observation periods, the results of the tests suggest that the proportional-hazard assumption is violated for the cohort-size variable as well as for the whole set of covariates. A possible response to the null hypothesis of the proportional-hazards assumption being violated is to allow for the effects of the regressors to vary with time by including interactions with search duration (Box-Steffensmeier and Jones, 2004). In order to allow as flexible an approach as possible and to avoid that certain variables pick up the effect of other regressors for which no interactions have been included, a model should be specified that contains interactions with search duration for every regressor. However, given the relatively large number of control variables, the estimation procedure for such a model does no converge. If a model is estimated that only interacts the size of the graduation cohort with the duration of search, the former variable's effect on the hazard rate is found to decrease in magnitude with time spent searching. Due to the above-mentioned concerns about models in which interactions are only included for a subset of regressors (and in this case only for a single regressor) this approach is not pursued any further. In order to justify the use of the paper's results I argue that the main conclusions about the effects of entry-cohort size are based on the 3-month period for which the proportional-hazard assumption appears to be locally satisfied.

Table S8: Proportional-hazard test (Schoenfeld residuals)

	3 months	6 months	1 year	2 years
Local test				
Cohort size	0.50 (0.48)	15.71 (0.00)	15.39 (0.00)	23.93 (0.00)
Global test	170.59 (0.04)	454.85 (0.00)	226.01 (0.00)	252.17 (0.00)

Tests statistics are based on Harrell's rho for the local test and on the Grambsch and Therneau method for the global test. Test statistics are χ^2-distributed. P-values are in parentheses.

S2 Sensitivity analysis

This subsection consists of two parts. In the first part, the robustness of the Cox model's results is assessed against various changes in the sample as well as in the empirical specification. The aim of the second part is to compare the results from the Cox model, which does not make an assumption regarding the specific distribution of the search durations, with a number of models that assume that the search durations follow a particular distribution.

S2.1 Robustness checks on the Cox model

The first set of robustness checks continues with assessing to what extent the results are driven by influential observations. However, in contrast to the above approach in which individual observations were excluded from the sample (see Figure S2) this analysis successively drops all observations from a given graduation period. A similar analysis is then performed in which observations from individual graduation regions are excluded from the sample.

Figure S3 shows the cohort-size coefficient and the corresponding 95% confidence interval for each case in which one of the 28 graduation periods is excluded. In order to allow for a comparison with the results from the full sample, the former's coefficient and 95% confidence interval is included in form of vertical lines. In each case the estimated cohort-size coefficient lies within the full model's confidence interval (represented by the dashed blue lines) and is typically close to the coefficient of the full model. An exception is period 10, which corresponds to the graduation period May-October 2003, where the change is more pronounced and the coefficient becomes insignificant.

Figure S4: Exclusion of single graduation periods from the sample

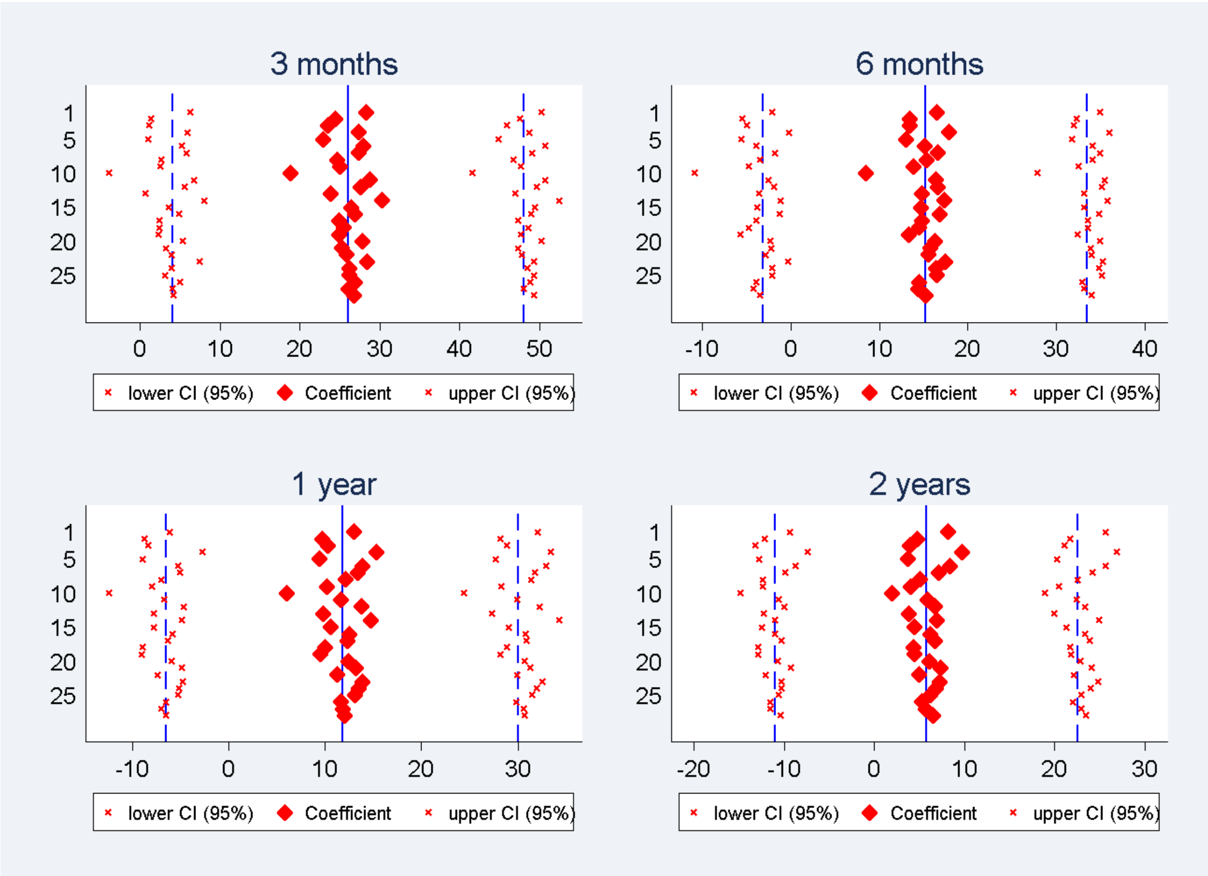

Source: Integrated Employment Biographies and Sample of Integrated Employment Biographies (author's calculations). Cohort-size coefficients are estimated as in Equation 1. The solid blue line represents the cohort-size coefficients from the full model, the blue dashed lines the corresponding 95% confidence interval.

Due to the large number of regions (141) the results from omitting individual graduation regions are shown in form of a histogram which illustrates the distribution of the resulting cohort-size coefficients. As can be seen from Figure S4, there is a certain degree of variation around the coefficient from the full model for each observation period, but these differences are small when compared to the confidence interval of the full model's coefficient. Moreover, for the 3-month observation period all cohort-size coefficients are significant at the 5% level, while in the 6-month period some of the coefficients become significant at the 10% level.

Figure S5: Exclusion of individual graduation regions from the sample

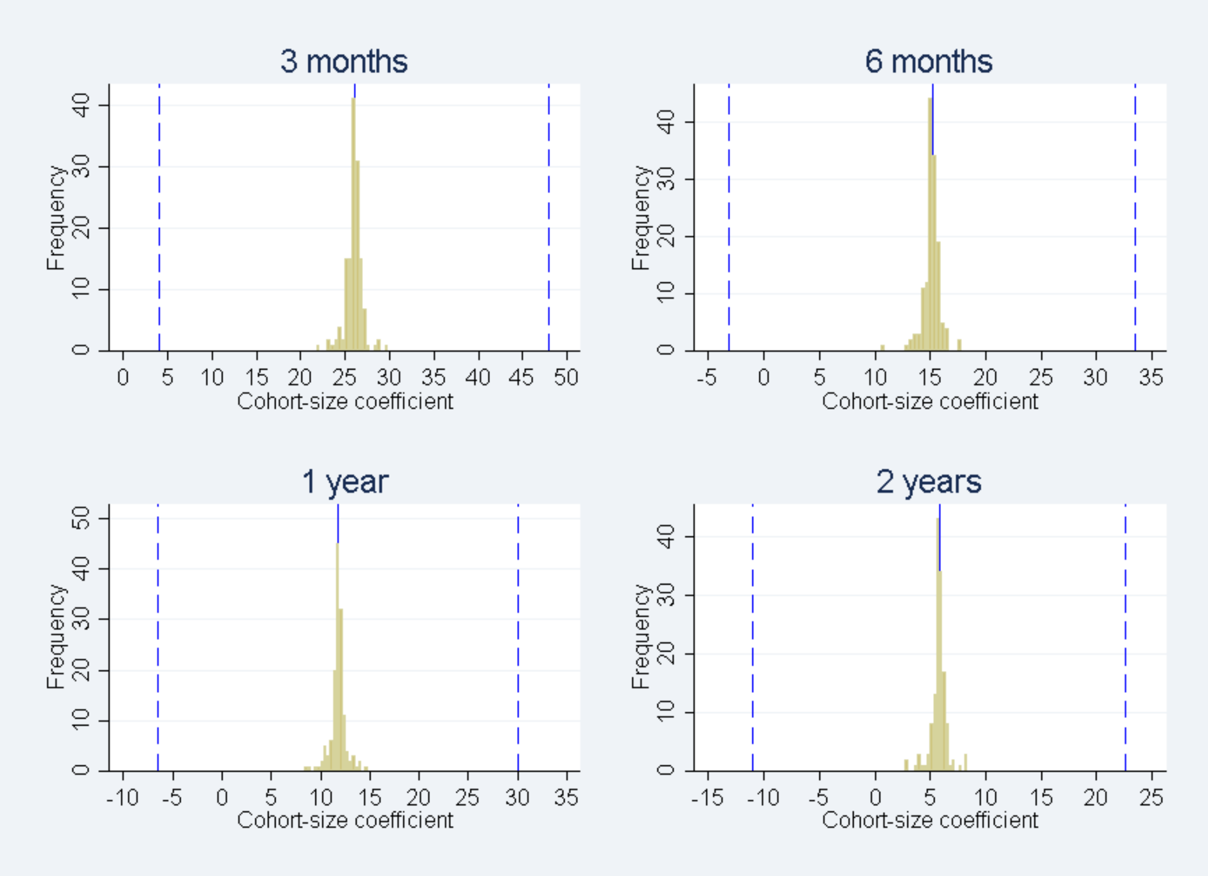

Source: Integrated Employment Biographies and Sample of Integrated Employment Biographies (author's calculations). Cohort-size coefficients are estimated as in Equation 1. The solid blue line represents the cohort-size coefficients from the full model, the blue dashed lines the corresponding 95% confidence interval.

In the paper the sample is homogenised by only including individuals that are aged between 19 and 23 at the time of graduation. As can be seen from Table S2, comparable results are obtained when this restriction is not imposed: the cohort-size coefficient remains positive and significant for the first observation period; moreover, it falls in size and becomes insignificant when the observation period is increased. For the first three periods of observation the coefficients are between 15% and 20% smaller than those of the full model, though the absolute changes are always considerably smaller than the size of the estimated standard errors.

Table S9: Dropping the age restriction

	3 months	6 months	12 months	24 months
Cohort	22.09**	11.95	9.64	3.66
	(9.31)	(8.10)	(8.20)	(7.76)
Dummies				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-105,389.80	-134,403.74	-163,907.26	-187,558.95
Observations	22,828	22,828	22,828	22,828
Clusters	141	141	141	141
ME(std)	1.06**	1.03	1.03	1.01

All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Breslow method is used to handle tied observations. *ME(std)* shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.

Search durations are measured in days and therefore transitions into employment are only observed at discrete points in time (even though the underlying process that generates the search series may be continuous). Under these conditions observations might be tied, i.e. there may be two or more observations with the same observed search duration. In order to account for tied observations the partial likelihood function has to be adjusted. Specifically, an adjustment has to be made to the definition of the risk set, i.e. the set of individuals that are at risk of experiencing a transition into employment at any given value of search duration. Assuming that the true search duration is indeed continuous, two individuals with the same observed search duration will in fact not have experienced transition into employment at the same point in time. If this is the case, both individuals will be in the risk set at the time of the first transition, but at the time of the second transition the individual that found employment earlier will no longer be part of the risk set. The Breslow method for handling tied observations, which has been used in the empirical analysis up to this point, does not make this distinction and instead assumes that the risk set is the same for all individuals sharing the same observed search duration. In contrast, the Efron method takes into account that sequential transitions give rise to different risk sets. Table S3 shows that employing the Efron method yields estimated coefficients and standard errors for the cohort-size coefficient that are very similar to those derived from the Breslow method.

Table S3: Efron method

	3 months	6 months	12 months	24 months
Cohort	26.12**	15.20	11.85	5.82
	(11.25)	(9.37)	(9.35)	(8.60)
Dummies				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-81,407.13	-103,453.55	-126,157.62	-145,272.21
Observations	18,133	18,133	18,133	18,133
Clusters	141	141	141	141
ME(std)	1.07**	1.04	1.03	1.02

All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Efron method is used to handle tied observations. *ME(std)* shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.

S2.2 Parametric model specifications

The empirical analysis of the paper uses the Cox model, which represents an example of a semi-parametric estimation approach as the functional form of the relationship between the hazard rate and the covariates is parameterised, whereas the actual distribution of failure times is left unspecified. Despite not specifying a distribution the Cox model is able to consistently estimate the effect that changes in the covariates have on the hazard rate. Alternatively, a fully parametric approach can be employed, which makes an assumption about the type of distribution from which search durations are drawn. A drawback of this approach is that the validity of the results depends on having chosen the correct distribution function.

In the following, the robustness of the Cox model's results is assessed against specifying a particular distribution. The first three distributions considered – the exponential, the Weibull and the Gompertz distribution – are compatible with the metric in which the Cox model is formulated, i.e. they can also be specified in form of a model of the hazard rate and, moreover, the first two also share the proportional-hazards property of the Cox model. The hazard functions that can be derived from these distributions take the following form:

$$h_i(t) = h_0(t)e^{(\beta cohort_{rp} + \boldsymbol{\delta' x_{irp}})} = e^{(\alpha)}e^{(\beta cohort_{rp} + \boldsymbol{\delta' x_{irp}})} \quad [S1]$$

$$h_i(t) = h_0(t)e^{(\beta cohort_{rp} + \boldsymbol{\delta' x_{irp}})} = pt^{p-1}e^{(\alpha)}e^{(\beta cohort_{rp} + \boldsymbol{\delta' x_{irp}})} \quad [S2]$$

$$h_i(t) = h_0(t)e^{(\beta cohort_{rp} + \boldsymbol{\delta' x_{irp}})} = e^{(\gamma t)}e^{(\alpha)}e^{(\beta cohort_{rp} + \boldsymbol{\delta' x_{irp}})} \quad [S3]$$

These specifications differ from the Cox model in that the baseline hazards are fully parameterised and depend on the constant of the model α as well as, in the case of the Weibull and the Gompertz specifications, on a distribution-specific shape parameter that also has to be estimated. These parameters determine the shape of the baseline hazard, i.e. the predicted hazard rate when all covariates take on a value of zero. Under the Weibull and the Gompertz specification the baseline hazard may be a flat, monotonically increasing or monotonically decreasing function of search duration. In contrast, the exponential model is less flexible in this respect as the baseline hazard is invariably flat, which may not be a realistic prior assumption if the probability of finding employment decreases with the duration of search.

For each of the above models and observation period, Table S4 shows the coefficient of the cohort-size variable as well as the estimated auxiliary parameters. The size of the coefficients can be compared directly with the results from Table 2 in the paper. The results show that the estimated effect of cohort size on search duration is robust to the use of a fully parametric specification, with each of the three distributions yielding coefficients that are similar in size and significance to those of the Cox model. Moreover, the negative sign of the auxiliary parameters in the Weibull and the Gompertz model suggests that the baseline hazard is decreasing with the duration of search, implying that, ceteris paribus, the hazard of finding employment is lower the longer the duration of search.

A prominent feature of the distributions in Table S4 is that they could be expressed in the same metric as the Cox model, i.e. in terms of the hazard rate. However, obtaining such an expression is not possible for all parametric models and the following set of examples – the lognormal, the loglogistic and the generalised Gamma distribution – are instead expressed in the accelerated failure time (AFT) metric: instead of estimating the effect of a change in a specific covariate on the hazard rate, its effect on the survivor function is estimated, which shows the share of observations that have not experienced transition into employment for each value of search duration. Estimates from the hazard rate can be derived from the survivor function and in contrast to the above group of distributions the former may take on a non-monotonic shape. The generalised Gamma distribution in particular allows for more flexibility as its shape is determined by two auxiliary parameters. Since these models are parameterised in terms of the survivor function, the coefficients cannot be directly compared with the results of Table 2. However, if the relationship between cohort size and search durations in these models has the same sign as estimated by the Cox model, the cohort-size coefficients should have the opposite sign as in Table 2: if an increase in the size of the graduation cohort increases the hazard of finding employment, the corresponding effect on the survivor function should be negative since a larger instantaneous probability of finding employment should lead to a smaller share of individuals not having experienced a transition at any given time.

Table S4: Parametric models (Exponential, Weibull, Gompertz)

	3 months	6 months	12 months	24 months
Exponential				
Cohort	27.02**	16.71	12.90	5.61
	(11.85)	(10.50)	(11.12)	(11.28)
Log pseudo-likelihood	-24,167.63	-29,295.60	-33,706.08	-36,401.95
Weibull				
Cohort	26.25**	15.39	11.82	5.63
	(11.38)	(9.61)	(9.61)	(8.90)
Auxiliary parameter: ln(p)	-0.22***	-0.31***	-0.36***	-0.37***
	(0.01)	(0.01)	(0.01)	(0.01)
Log pseudo-likelihood	-23,906.67	-28,519.32	-32,258.50	-34,407.30
Gompertz				
Cohort	26.14**	15.31	12.49	7.08
	(11.28)	(9.33)	(9.43)	(8.81)
Auxiliary parameter: γ	-0.01***	-0.01***	-0.00***	-0.00***
	(0.00)	(0.00)	(0.00)	(0.00)
Log pseudo-likelihood	-23,852.43	-28,331.70	-32,351.38	-34,842.75
Dummies				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Observations	18,133	18,133	18,133	18,133
Clusters	141	141	141	141

All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region.

Table S5 contains the coefficients of the cohort-size variable as well as the auxiliary parameters estimated from a model based on the lognormal, the loglogistic and the generalised Gamma distribution, respectively. In order to compare the size of the estimated effects to those of the Cox model and to those for the first set of distributions, Table S5 also reports the results from a Weibull model, which may also be specified in terms of the survivor function and which, in the hazard-rate metric, yielded coefficients that were very close to those of the Cox model. As hypothesised, the coefficients of the cohort-size variable turn out negative for each period of observation, which implies that, as was the case with the Cox model, increases in the size of an individual's graduation cohort are associated with shorter search durations. Moreover, the size of the coefficient decreases as the observation period becomes longer, which also corresponds to the results from the Cox model. The results from the lognormal, the loglogistic and the generalised Gamma specifications differ, however, in that the coefficients tend to be larger than in the Weibull model, which for the hazard-rate metric produced coefficients that were very close to those of the Cox model: for the 3-month observation period the former are between 15% and 30% larger than those from the Weibull model, with the difference being larger for longer periods of observation. Moreover, the estimated effects of cohort size are also significant in the 6-month and the 1-year period of observation.

Table S5: Parametric models (Weibull, lognormal, loglogisitc, generalised Gamma)

	3 months	6 months	1 year	2 years
Weibull				
Cohort	-32.62**	-21.05	-17.00	-8.19
	(14.16)	(13.14)	(13.83)	(12.94)
Auxiliary parameter				
ln(p)	-0.22***	-0.31***	-0.36***	-0.37***
	(0.01)	(0.01)	(0.01)	(0.01)
Log pseudo-likelihood	-23,906.67	-28,519.32	-32,258.50	-34,407.30
Lognormal				
Cohort	-42.01***	-33.24**	-30.13**	-24.45*
	(16.08)	(14.86)	(14.71)	(13.64)
Auxiliary parameter: ln(σ)	0.64***	0.64***	0.62***	0.58***
	(0.01)	(0.01)	(0.01)	(0.01)
Log pseudo-likelihood	-23,817.50	-28,301.08	-32,032.23	-34,322.21
Loglogistic				
Cohort	-37.02***	-28.23**	-26.27*	-22.63*
	(15.60)	(14.64)	(14.43)	(13.63)
Auxiliary parameter: ln(γ)	0.06***	0.09***	0.07***	0.03***
	(0.01)	(0.01)	(0.01)	(0.01)
Log pseudo-likelihood	-23,847.60	-28,349.18	-32,074.06	-34,370.44
Generalised Gamma				
Cohort	-39.61**	-31.28**	-26.48*	-16.75
	(15.70)	(14.64)	(14.45)	(13.18)
Auxiliary parameter: ln(σ)	0.55***	0.61***	0.57***	0.51***
	(0.02)	(0.01)	(0.01)	(0.01)
Auxiliary parameter: κ	0.25***	0.15***	0.26***	0.43***
	(0.05)	(0.04)	(0.03)	(0.03)
Log pseudo-likelihood	-23,804.83	-28,293.44	-31,997.53	-34,184.58
Dummies				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Observations	18,133	18,133	18,133	18,133
Clusters	141	141	141	141

All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed in the accelerated failure time metric. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region.

References

Box-Steffensmeier, J.M. and Jones, B.S. (2004), *Event History Modeling: A Guide for Social Scientists*, Cambridge University Press, Cambridge.

Cleves, M., Gutierrez, R.G., Gould, W. and Marchenko, Y.V. (2010), *An Introduction to Survival Analysis Using STATA (3rd edition)*, STATA Press, College Station.