Social networks and unemployment exit in Great Britain

Panos Sousounis
Keele Management School, Keele University, Keele, UK, and
Gauthier Lanot
School of Business and Economics, Umeå Universitet, Umeå, Sweden

Structured Abstract

Purpose: The aim of this paper is to examine the effect employed friends have on the probability of exiting unemployment of an unemployed worker according to his/her educational (skill) level.
Design/Methodology: In common with studies on unemployment duration, this paper use a discrete time hazard model.
Findings: The paper finds that the conditional probability of finding work is between 24% and 34% higher per period for each additional employed friend for job seekers with intermediate skills.
Social Implications: These results are of interest since they suggest that the reach of national employment agencies could extend beyond individuals in direct contact with first line employment support bureaus.
Originality/Value: Because of the lack of appropriate longitudinal information, the majority of empirical studies in the area assess the influence of social networks on employment status using proxy measures of social interactions. The current study contributes to the very limited empirical literature of the influence of social networks on job attainment using direct measures of social structures.

Introduction

The role of social networks as informational vehicles in job search was first highlighted in the economics literature by Myers and Shultz (1951) and Rees (1966). Friends' networks among other informal channels contribute to the dissemination of job information through individual inquiries and personal recommendations. It was in the field of sociology however where the effects of social interactions in job attainment were first formalized in the 1970s (Boorman, 1975; Granovetter, 1995).

Sociological interest in the effect of social networks on unemployment exit (and unemployment in general) stems mainly from (1) the emergence of a, so called, "(economic) underclass" (Murray et al., 1990), which implies a degree of social segregation among employed and unemployed individuals, and (2) concerns over "social exclusion", which features very high up in the EU agenda[1] since the late 1980s. The latter, is of course, a very broad term encompassing an array of social processes[2], however, labour market behaviour is considered a catalyst for wider social marginalisation (Russell, 1999).

Differences in job search outcomes among members of similar socio-demographic groups are mostly attributed to information asymmetries (see, for example, Cingano and Rosolia, 2012; Matthews et al., 2009, and references therein). Employees have access to information about

current and upcoming vacancies, which circulate through their social networks. The extent of an unemployed worker's network and the proportion of employed individuals the network contains will then determine the worker's exposure to available employment opportunities. Hence some job seekers receive more information about available jobs than others. Granovetter (1995) makes the distinction between strong and weak ties (close(st) friends and acquaintances, respectively) and argues that weak ties confer more vacancy information since close friends are likely to have access to similar, if not the same, information. Murray et al. (1981), however, argue that strong ties are in fact more effective information providers for job seekers. Yakubovich (2005) finds merit in the influence of both weak and strong ties in job attainment.

Independently of their weak-strong ties composition then, friends' networks appear both efficient and effective in diffusing (Calvo-Armengol and Ioannides, 2005) and filtering information about job vacancies across potential candidates ultimately leading to better matches. Incumbent workers possess specific knowledge of their employer and/or the available jobs and are more likely to have better insight of their friends' (network's) unobservable attributes. They are then able to provide an informal screening mechanism for specific jobs[3]. Employers leverage this informal arrangement to decrease job turnover and reduce recruitment and training costs (Finneran and Kelly, 2003). Individuals on their side benefit through reduced search costs, pecuniary or not, as more and better information about job vacancies flows through their networks.

Another explanation based on an alternative mechanism that delivers a similar outcome could result from increased search effort in the face of peer pressure to avoid the stigmatizing effect of unemployment. This explanation rests on a model of "collective socialization" which posits "the spread of some socially positive behaviour due to interaction of individuals with role models or community networks" (Dietz, 2002, p. 543). In this context, while employed contacts pass on job information, they also offer guidance and advice on job search strategies, on the correct filling of job applications, or on interviewing tips, all which adds to the job seekers' human and social capital. In either case the outcome is socially desirable and beneficial. Consequently, we believe that the umbrella term "job-related" information suits better the description of information circulated within networks.

Hannan (1999) argues that viewing social networks as means of information circulation only is rather limited. While discussing common definitions of 'social capital' in the literature, she supports a much broader definition of social capital, one that encompasses a multitude of social processes partly proxied by or embedded in social networks. The mutual support enjoyed among members of a given social network significantly contributes to its members' psychological well-being. This in turn leads to increased feelings of self-worth and competence in those seeking employment. The latter effect leads to better choices and/or opportunities in the labour market. Hannan describes yet another potential mechanism through which social networks lead to increased probability of unemployment exit. However, Hannan stresses that untangling the influence of each of these channels is impossible with currently available data. Therefore, while noting the potential importance of psychological determinants, we proceed with assuming, as much of the literature does, that information transmission is central to the process that transmits information between firms and employed workers to unemployed workers.

The flow of information from one member of the network to another implies a link between

the labour force status of each member (Jackson, 2011). Models of social networks in labour markets such as those proposed by Calvo- Armengol and Jackson (2004, 2007) predict exactly such patterns. These theoretical models thus imply positive spill-over effects within social networks. Consequently, the larger the network of employed friends one is part of, the higher the probability of exiting unemployment.

Confirming the existence of a positive empirical link between social networks and employment status is of interest to both national employment agencies since it suggests that a network's reach could extend beyond individuals in direct contact with first line employment support bureaus and state agencies concerned with mitigating social exclusion. In most developed labour markets, information about job vacancies is transmitted to job seekers via both formal (e.g. state employment services such as the Jobcentre Plus and Universal Jobmatch in the UK, private employment agencies, and university careers services) and informal (such as referrals and/or exploratory inquiries, and personal contacts) channels. The two operate alongside and amplify the reach[4] of official employment agencies, especially to those more passive job seekers[5]. In that respect, social networks provide a valuable service to local and national employment authorities (such as the Department for Work and Pensions in the UK).

Kramarz and Skans (2014) and Brown et al. (2016) using Swedish and US data respectively, show that social networks have a greater positive effect on the likelihood of exiting unemployment for individuals at the lower end of the skills distribution. This result is echoed by Oesch and von Ow (2017) in their study of Swiss job seekers. The authors argue that low-skilled job seekers are more likely to succeed in obtaining employment through personal contacts since their limited human capital will greatly restrict their choices through formal recruitment channels. In that respect, social networks act as a compensation device for low levels of formal qualifications and/or labour market experience. While this is supported by different empirical studies, the evidence is not yet conclusive (see Oesch and von Ow, 2017, and references therein).

The majority of empirical studies in the area assess the influence of social networks on employment status using proxy measures of social interactions due to appropriate longitudinal data unavailability (Jackson, 2011). In the current study, we extend findings reported by Hannan (1999) and Cappellari and Tatsiramos (2011) to individuals at different parts of the skills distribution. We explore whether increasing the number of employed friends, increases the probability of finding work, for job seekers of differing skills levels (as defined by level of schooling) in Great Britain. Our study hence contributes to the very limited empirical literature of the influence of social networks on job attainment using direct measures of social structures and the stock of studies on the influence of social networks on unemployment exit.

Data and Empirical Strategy

Data

The data used is drawn from the British Household Panel Survey (BHPS), a representative sample of British households surveyed first in 1991 and followed up annually since. Exiting from unemployment to employment is the event of interest. We use the reported current (period) labour force status to identify transitions from being at risk i.e., unemployed and actively looking for a job[6], to experiencing the event i.e., finding employment. Transitions to other states (self-employment or inactivity, including retirement, FT education/training, maternity leave/caring)

from unemployment are excluded[7].

Each individual contributes T_i rows of observations, where T_i is the number of years individual i is at risk of finding a job. Individuals subsequently either experience the event or are censored at the end of the study period – the dataset is thus right-censored, which we deal with by appropriate modifications to the likelihood function as suggested by Jenkins (2005). Observations are interval-censored, in the sense that we only know that the event happens (if it does) within a specified time interval, and not the exact (continuous) time. We deal with interval censoring by specifying a discrete- time duration (survival) model. The sample is also left-censored, common in observational studies, in the sense that some individuals become at risk before t_0, the start of the observation period. For left censored observations, we follow Fichman (1989) and Singer and Willett's (2003) advice and carry out the analysis by (1) redefining the beginning of unemployment to coincide with the start of the observation period and (2) by excluding all left censored observations. In either case the results remain qualitatively unaltered[8] so we proceed with estimation on the entire sample. In the current setting, censoring is assumed non-informative.

Starting in wave 2 (1992), and every other year since[9], the employment status of the respondents' three closest friends[10] is recorded in three categorical variables. We determine whether the respondents' first, second and third closest friend are either employed (=1) or unemployed (=0). We then use these variables to construct a cumulative measure of friends in employment (full- or part- time). This measure becomes the main substantive predictor of our models. We further explore the existence of non-linear effects by including a set of three dichotomous variables indicating the exact number of employed friends (with no employed friends as the reference category). In our "extended" specifications, the vector of explanatory variables contains further demographic, socio-economic and individual respondents' characteristics and regional effects. Finally, we sort the respondents according to their skill level as defined by their formal qualifications at the time of interview (if any). We classify individuals to three broad groups[11] according to whether they have no listed qualifications (low-skilled), listed school qualifications (middle skilled), or listed post-school qualifications (high-skilled). We estimate the discrete-time hazard for each group, in order to explore possible variations in the use of social networks among job seekers of different skills. Table A1, in the appendix, presents sample descriptive statistics (mean/proportion and the respective standard deviations). We test the missing completely at random (MCAR) assumption for missing values in our sample (<1%) using Little's (1988) test, which allows us to also test the covariance-dependent missingness (CDM) assumption. Little's test gives a distance of 10.97 with degrees of freedom 11 and value 0.0044, hence there is evidence that the missing data are not MCAR. We then test the CDM assumption by including the covariates into the test as auxiliary variables. In this instance, the distance is 138.7 with degrees of freedom 162 and a p-value of 0.9070, implying that even though the dependent variable cannot be assumed MCAR, the missing data mechanism can be seen as CDM provided the effect of the explanatory variables is accounted for. Estimates from specifications without conditioning variables should therefore be viewed as valid under the less restrictive, yet plausible, missing at random (MAR) assumption.

Figure 1 presents the distribution of employed friends according to skills level and labour force status in the aggregate person sample. We observe that among employed participants, high-skilled individuals are more likely to have larger networks of employed friends, which is not surprising. The distribution of employed friends across unemployed respondents is somewhat

different, suggesting that having employed friends is associated with one's own labour force status, which is the main hypothesis in the current study, or that friendships form within the workplace and/or industry (causality is reversed). We argue that while the latter is possible, in reality (close) friendships are unlikely to break down (relatively soon) when a worker experiences a transition out of the labour force or when an employer-employee separation occurs. Hence, it is plausible within the span of a year, the friends' labour force status in a network is systematically not determined by a worker's employment status. While the issue of endogeneity (or simultaneity) remains a concern, we are unable to address it in the current set up[12] (see, for example, Varum et al., 2014, and references therein for a discussion of the limitations of regressor endogeneity treatment in discrete time hazard models). Consequently, our results are reliable under the (plausible in the current context) assumption of exogeneity between the dependent (left-hand side) variable and the social network measure.

The estimated survival probability for each category of social network is depicted in figure 2. For the average skill group, it is clear that a larger number of employed friends is associated with a reduced survival rate i.e. exit to unemployment takes place earlier. For the low skilled, we observe that those with three employed friends find jobs faster, on average. This is an interesting observation, suggesting that being part of a network of employed individuals lessens discouragement. For the high skilled, larger networks of employed contacts reduce considerably the probability to remain unemployed. Estimating a median lifetime in our sample is difficult due to the survey design and the interval censored observations. A number of individuals are likely to move in and out of employment during the observation window, making their time spent searching and subsequently exiting unemployment, different to the one "observed". In this setting, we are unable to construct and utilise a complete working life diary, hence, we are limited in the characterisation of the event occurrence distribution.

Figure 1 [here]: Distribution of the number of employed friends according to skill level and labour force status.

Figure 3 plots the estimated hazard probability of exiting unemployment for individuals in our sample given the number of employed members in their social network. For medium and high skill workers having a larger number of friends appears to be associated with higher exit rates. This is not true for the low skilled workers. Noting the potential caveats due to the interval censored nature of our data, this is an encouraging observation. Despite any drawbacks, we could still see from the hazard and survivor functions that larger networks of employed friends are associated with higher hazards of the individual experiencing the event i.e., exiting unemployment.

Figure 2 [here]: Predicted probability of remaining unemployed according to the number of employed friends in the respondents' social network.

Figure 3 [here]: Discrete-time hazard of exiting unemployment according to the number of employed friends in the respondents' social network.

Figure 3 suggests that the proportional hazards assumption implicit in our modelling strategy (described below) may not be satisfied exactly– the main predictor's effect appears more distinct in some years than others. Before we proceed further, we test the proportionality assumption by comparing the goodness-of-fit between a main effects hazards model with models including interactions with time. We do so following Singer and Willett's (2003) suggestion of

5

comparing the difference in the models' deviance statistic (= –2×Log-likelihood) to a χ^2 distribution with degrees of freedom equal to the difference in the number of parameters in each model (Singer and Willett, 2003, p. 455). If the models with the time interactions (non-proportional) provide significant improvements over the main effects hazards model (proportional) then the former specifications are preferred.

The full set of results is not reported here in order to conserve space but is available online (Table B1). In the general interaction with time model, we allow the effect of the number of employed friends to be the same in each period. Comparing the deviance statistics between the main effects hazards and the models which include a interaction with time suggests no significant improvement ($\chi^2 = 28.8$, d.o.f.$= 18, p > 0.05$). Prompted by the irregular evolution of the parameter estimates for the interaction terms, we further allow for the effect of employed friends to interact linearly with time. Again, our testing approach suggests insignificant improvement by the inclusion of the interaction term ($\chi^2 = 27.5$, d.o.f.$= 17$, $p > 0.05$). Comparing the model with a general interaction with time to the linear time model we conclude that the simpler specification fits, at least, as well as the general one ($\chi^2 = 3.8$, d.o.f.$= 1, p > 0.05$). To sum up, we cannot reject the proportionality assumption in any case and hence adopt the main effects hazards model as our preferred specification.

Empirical strategy

We start by specifying a discrete-time survival model in terms of the conditional probability that the event occurs at time t, given that it has not yet occurred. Estimates of the hazards could be (and commonly are) obtained through a logistic regression of an indicator variable (dependent) for individual i finding employment at time t, on a set of time-varying covariates and time-invariant predictors. The relationship between the discrete-time hazard and the time period is left unspecified but a set of period dummies, which describes the baseline hazard when all explanatory variables are set to zero, is included. Such model is semi-parametric in nature since one does not choose a functional form for the baseline hazard. The semi-parametric model provides flexibility at the cost of efficiency loss compared to a correctly specified, fully-parametric model.

Singer and Willett (2003) argue for using a complementary log-log (clog-log) link function instead of a logit function when the data is interval-censored (finding a job happens in continuous time but only the discrete time interval in which it happens is observable due to survey design). Further, the clog-log link implies a proportional hazards assumption (while the logit link implies a proportional odds assumption) and hence provides a direct analog to the popular Cox regression model[13]. So, if the proportional hazards assumption holds in continuous-time and survival times are interval-censored, as in our data, we can specify a complementary log-log discrete-time survival model,

$$h_{it} = 1 - \exp\{-\exp(\boldsymbol{\alpha}_T\boldsymbol{\delta}_{it}, \boldsymbol{\beta}\mathbf{x}'_{it}, \boldsymbol{\gamma}\mathbf{x}'_i)\} \quad \text{(1)}$$

where $\mathbf{x}_{it}$ is a vector of time-varying covariates[14], $\mathbf{x}_i$ a vector of time-invariant predictors and $\boldsymbol{\delta}_{it}$ includes a complete set of period dummies (no intercept is included in the model). The complementary log-log model, again, makes no assumption about the shape of the discrete-time baseline hazard function among years[15].

6

A key feature of this modelling approach is how one specifies the baseline hazard function i.e. the main effect of time. Including a set of indicator variables for each discrete time period yields a general specification, which is sufficiently flexible yet potentially inefficient. Assuming that time is linearly related to the hazard, even though unemployment exit is modelled in a discrete-time framework, is achieved by substituting $\boldsymbol{\delta}_{it}$ with δ_i, a linear trend of the time individual i is at risk of finding a job. Another popular approach involves specifying time as a polynomial of certain order. Often, a set of ordered polynomials are evaluated against the general specification and the preceding lower order polynomial specification. We follow such an approach here. We compare polynomial specifications of time up to order five with the general specification using the deviance statistic and the difference in deviance between specifications. Our results suggest (not reported here but available online, see Table B2) the use of polynomial specifications with multiple stationary points (four or greater). However, as (Singer and Willett, 2003, p. 412) suggest, in such instances, the general specification should be used in light of the complexity of interpretation of the higher order polynomial.

Accounting for the possible effect of unobserved heterogeneity among individuals is achieved by including a random intercept, ζ_i, in (1):

$$h_{it} = 1 - \exp\{-\exp(\boldsymbol{\alpha}_T\boldsymbol{\delta}_{it}, \boldsymbol{\beta}\mathbf{x}'_{it}, \boldsymbol{\gamma}\mathbf{x}'_i, \zeta_i)\} \quad (2)$$

which can also be written as:

$$y^*_{it} = \boldsymbol{\alpha}_T\boldsymbol{\delta}_{it} + \boldsymbol{\beta}\mathbf{x}'_{it} + \boldsymbol{\gamma}\mathbf{x}'_i + \zeta_i + \varepsilon_{it} \quad\quad\quad (3)$$

Where we assume that $\zeta_i \sim N(0, \sigma_\zeta)$ and $\varepsilon_{it}|\zeta_i, \mathbf{x}_{it}, \mathbf{x}_i$ follows the standard extreme-value type-1 (Gumbel) distribution. Furthermore, we assume that ζ_i and ε_{it} are independent of the elements of $\mathbf{x}_{it}, \mathbf{x}_i$, and each other, and the errors, ε_{it} are also mutually independent. y^*_{it} is a continuous latent response determining the observable event variable y_{it} (=1 if employed and =0 if not) according to

$$y_{it} = \begin{cases} 1, & if\, y^*_{it} > 0 \\ 0, & otherwise. \end{cases}$$

The lack of evidence against $H_0: \sigma_\zeta = 0$ implies that the individual heterogeneity (or *frailty*, in survival analysis lingo) may be ignorable i.e. $\zeta_i = 0$ and hence the individual-specific survival probability is similar to the population-averaged one.

Results and discussion

We start by estimating two alternative (linear) specifications of model 1, presented in table 1, where model A includes the effect of friends in employment and model B additionally includes further socio-economic, household and personal respondent's characteristics[16] in the vector of explanatory variables[17].

In both specifications, the coefficient estimates are maximum likelihood estimates of the baseline complementary log-log hazard function. The (exponentiated) coefficient estimates for the period indicators decline over time, with irregular increases in periods five and seven, suggesting that the chances of exiting unemployment increase over time. Table 1 also reports

goodness-of-fit statistics at the bottom panel - the deviance statistic and the AIC and BIC information criteria all suggest the better fit of the extended (model B) specification.

Our results from both specifications suggest that the size of one's social network has a statistically significant effect for those individuals with school- level qualifications only. From model A, we get that the hazard rate for these respondents increases by approximately 34% for each additional employed friend one has. This effect remains statistically significant but reduces in magnitude to 24% in the unrestricted specification (model B). This result is in line with other findings in the area e.g. Kramarz and Skans (2014), Brown et al. (2016) and Oesch and von Ow (2017) who find significant positive effects (lower in magnitude to ours) of social networks for those with fewer formal credentials, and with broader studies such as Korpi (2001) who finds hazard ratios of unemployment exit to employment in Sweden of 14% to 24% across specifications. Cappellari and Tatsiramos (2011) also report increases in the job finding rate of 13% for every additional employed friend a job seeker in the UK has. Gush et al. (2015) report, from a qualitative study, the importance of social networks in job search in the UK. The authors report personal accounts of interviewees on how they utilised their social connections after, and even ahead of (foreseen), redundancy to find alternative employment.

Our findings lend support to the hypothesis that social capital, as embodied within social networks, can compensate for (the lack of) human capital. Assuming that formal qualifications signal ability or other innate attributes, lower-skilled job seekers have to rely more on personal referrals and recommendations to pass through the initial screening processes of employers. This process is, of course, supply as well as demand driven. Recruiters' behaviour and attitude towards personal recommendations is critical. However, in low-skill occupations, with high turnover and relatively low hiring and firing costs, employers are likely to be more receptive of informal hiring practices.

The implications to society are significant. Social networks can match lower skilled job seekers to employers and jobs which would otherwise be out of reach due to lack of either formal credentials or knowledge of the opportunities. While higher skilled individuals with more qualifications may have greater commitment and perseverance in the labour market, lower skilled job seekers are likely to be easier discouraged following possibly multiple rejections. This would lead to labour market (and subsequently social) marginalisation. Among the lower skilled then, social networks could provide an escape route from unemployment. Social networks could contribute to amplify the indirect effect of unemployment agencies to other members of a successful job seeker's network.

We also explore the presence of non-linear effects of social networks on the hazard rate of unemployment exit, as well as possible state dependence of the time-varying measure of employed friends. It is plausible that the effect of one's network of employed friends manifests with a time lag rather than contemporaneously. Table 2 presents our estimates of the different models[18]. In the non-linear specification we, once again, only find statistically significant results of employed friends' influence on the hazard rate of finding a job for middle-skilled individuals in the sample. Having two or three employed friends almost doubles one's chances of finding a job. In the dynamic linear specification we find statistically significant improvements in the hazard rate for those with post-school and school-level qualifications, suggesting that the size of one's network of employed friends is exerting a positive effect in subsequent periods. Assuming that higher skilled individuals are after jobs with formal recruitment processes, the positive effect of the social network is more likely to operate through the dissemination of job

8

search techniques and strategies (post-experience) rather than referrals. Nonetheless, encouragement and support, as well as the emulation of good practice cannot be ruled out as the causal mechanism. This could arguably also be due to the nature of our interval censored data rather than a true state-dependence effect and should be the focus of future research.

Table 1 [here]: Linear specifications of the complementary log-log model to the time of unemployment exit to employment.

Table 2 [here]: Nonlinear and dynamic specifications of the complementary log-log model to the time of unemployment exit to employment.

Since the parameter estimates in Table 1 ignore possible unobserved heterogeneity, the estimated effects may be biased. According to (Jenkins, 2005, p. 81), not accounting for the effect of unobserved heterogeneity leads to either over- or under-estimating the actual hazard rate. We estimate conditional hazard ratios from our model which now includes a subject-specific random component to account for unobserved heterogeneity (frailty). In both specifications, the within-subject correlation is very low and the random intercept insignificant, suggesting that most of the dependence among subjects is captured by the model's fixed component. Due to the very low within-subject correlation, the estimated hazard ratio (1.257) is very close to that from the model without a random intercept and we therefore do not report the estimates here in order to conserve space.

As a final check, we carry out a graphical inspection of the deviance residuals (Model 1) as described in Singer and Willett (2003) p.463-464. The authors suggest that "extreme[19]" values of the residuals, in absolute terms, indicate cases for which the model does not adequately fit the data. Figures A1 and, more prominently, A2 show that our model fits the data well, there are instances of "large" deviance residuals but arguably none of the identified cases are too "distant" from the bulk to raise serious concerns. All estimates remain robust to the removal of these respondents from the estimation sample.

Conclusion

The size of social networks plays an important in the job search of the unemployed. Recent theoretical contributions in the economics literature highlight the role that the social network of an unemployed worker can play in spreading and distributing information about job vacancies. Alternative transmission channels have been proposed in the literature, however disentangling the effects of all the potential mechanisms with the data currently available is nearly impossible. The difficulty rests with the lack of specific measurement of the characteristics of social networks and their use by unemployed workers.

Our current study adds to the limited empirical evidence by using the reported employment status of a respondent's three closest friends. This allows us to construct a direct measure of the "quality" of his/her social network. Admittedly, restricting the size of one's network to maximum three individuals is a limiting factor of the analysis, however, we are still able to provide insights into the effect friends have on the probability of exiting unemployment. We estimate the hazard rate of exiting unemployment given the number of employed contacts one has for job seekers of different skills groups, namely high, middle and low, as defined by the type of qualifications held (and the lack of for the low skilled). We estimate both static and dynamic (where we allow the effect of the social network to operate with a period lag)

specifications of our main model. Given the structure of our dataset and the discussion of potential caveats in the dynamic model our preferred specification is the linear static model (Model B in Table 1).

Our estimates suggest that for individuals in the middle of the skills distribution, for each additional employed friend, the hazard of exiting unemployment is around 24% higher. The corresponding estimate for the high-skilled is more moderate in magnitude at 8%, yet not generalizable. For the low-skilled, our model suggests that there is a marginal reduction in the hazard rate for every employed friend one has. However, as in the case for the high-skilled, the effect is not precisely measured. The direction and significance of these results are maintained in the non-linear specifications. In our dynamic specification, we find positive contributions of the social network for both high and middle skilled individuals comparable in size to those in the static case. For the low skilled, although the direction of the effect changes suggesting a moderate positive effect of 6%, the estimate remains imprecise. While we do not obtain compelling evidence on the effect of social networks for either low or high skilled job seekers, the evidence points to a positive effect for middle skilled individuals. These findings are in line with other studies in the area and lend support to the hypothesis that there is considerable substitution of formal educational credentials (human capital) with social capital (even at a broad definition of the latter). Although our *a priori* expectations were that the effect on the low skilled would have been more pronounced, our results are encouraging nonetheless.

For individuals with higher skills who are usually looking for job vacancies with (more) formal recruitment processes, we argue that the influence of social networks is likely to work through a support or share of good practice avenue. In the absence of more detailed data we are unable to explore this hypothesis further. For the low skilled in our sample, a stronger effect may be easier to uncover, if we follow an occupational rather that educational group classification. It is indeed more likely that individuals in low skill occupations utilize social networks more intensely (than middle and/or high skill job seekers who are after jobs with more formal recruitment channels). Ponzo and Scoppa (2010) report that in Italy, social (informal) networks are used more by workers at the lower end of the skills distribution. First, if we try to restrict the analysis to those in low skill occupations, the resulting estimation sample would be very low in our current dataset. Secondly, and plausibly interacting with the first, the setup of our data is such that it probably masks short transitions in and out of employment. Such transitions are, again, more likely in low skill occupations. This is arguably a main reason for the lack of conclusive effects for the low-skilled job seekers. We are therefore confident that in future research with a more populous and/or richer dataset the effect of social networks will be more prominent. Therefore, the results obtained here are interesting and could act as a springboard for future research.

These results are important for national employment agencies since they suggest that their reach could extend beyond the individuals in direct contact with first line employment support bureaux through the information flow between members of the same social networks, and thus be more effective than a simple count of "users" would suggest.

References

Beaman, L. and J. Magruder (2012). Who gets the job referral? Evidence from a social networks experiment. American Economic Review, 102 (7) 3574-93.

Boorman, S. A. (1975). A combinatorial optimization model for transmission of job information through contact networks. The Bell Journal of Economics, 6 (1), pp. 216-249.

Bramoulle, Y. and G. Saint-Paul (2010). Social networks and labor market transitions. Labour Economics 17 (1), 188 - 195.

Brown, M., Setren, E. and Topa, G. (2016). Do Informal Referrals Lead to Better Matches? Evidence from a Firm's Employee Referral System, Journal of Labor Economics 34, no. 1, 161-209.

Calvo-Armengol, A. and Y. M. Ioannides (2005). Social Networks in Labor Markets. Discussion Papers Series, Department of Economics, Tufts University 0517, Department of Economics, Tufts University.

Calvo-Armengol, A. and M. O. Jackson (2004). The efects of social networks on employment and inequality. American Economic Review 94 (3), 426- 454.

Calvo-Armengol, A. and M. O. Jackson (2007). Networks in labor markets:
wage and employment dynamics and inequality, Journal of Economic Theory 132 (1), 27 - 46.

Cappellari, L. and K. Tatsiramos (2011). Friends' Networks and Job Finding Rates. Discussion Papers in Economics 11/40, Department of Economics, University of Leicester.

Cingano, F. and A. Rosolia (2012). People I Know: Job Search and Social Networks. Journal of Labor Economics 30 (2), 291 - 332.

Dietz, R. D. (2002). The estimation of neighbourhood effects in the social sciences: An interdisciplinary approach. Social Science Research 31 (4), 539 - 575.

EuropeanUnion (2015). 2015 report of the social protection committee: Re- view of recent social policy reforms. Report, European Union: The Social Protection Committee, Luxembourg: Publications Ofce of the European Union.

Fichman, M. (1989). Attendance makes the heart grow fonder: A hazard rate approach to modeling attendance. Journal of Applied Psychology 74 (2), 325 - 335.

Finneran, L. and M. Kelly (2003). Social networks and inequality. Journal of Urban Economics 53 (2), 282 - 299.

Gallie, D., S. Paugam, and S. Jacobs (2003). Unemployment, poverty and social isolation: Is there a vicious circle of social exclusion? European Societies, 5 (1), 1-32.

Granovetter, M. (1995). Getting a Job: A Study of Contacts and Careers (2 ed.). University of Chicago Press.

Gush, K., J. Scott, and H. Laurie (2015). Job loss and social capital: The role of family, friends and wider support networks. Working Paper Series 2015-07, Institute for Social and Economic Research, University of Essex.

Hannan, C. (1999). Beyond Networks: Social Cohesion and Unemployment Exit Rates. Unpublished manuscript, Institute for Labour Research, University of Essex.

Jackson, C. K. and H. S. Schneider (2011). Do social connections reduce moral hazard? evidence from the New York City taxi industry. American Economic Journal: Applied Economics, 3 (3), 244-67.

Jackson, M. O. (2011). Chapter 12 - an overview of social networks and economic applications. In A. B. Jess Benhabib and M. O. Jackson (Eds.), Handbook of Social Economics, Volume 1, pp. 511 - 585. North-Holland.

Jenkins, S. P. (2005). Survival analysis. Unpublished manuscript, Institute for Social and Economic Research (ISER), University of Essex.

Korpi, T. (2001). Good friends in bad times? Social networks and job search among the unemployed in Sweden. Acta Sociologica 44 (2), 157-170.

Kramarz, F. and Nordström Skans, O. (2014). When Strong Ties are Strong: Networks and Youth Labour Market Entry, The Review of Economic Studies, Volume 81, Issue 3, 1164–1200.

Little, R. J. A. (1988). A test of missing completely at random for multivariate data with missing values. Journal of the American Statistical Association 83: 1198–1202.

Matthews, R., R. Pendakur, and N. Young (2009). Social capital, labour markets, and job-fnding in urban and rural regions: comparing paths to employment in prosperous cities and stressed rural communities in canada1,2. The Sociological Review 57 (2), 306-330.

Murray, C., I. of Economic Afairs (Great Britain). Health, and W. Unit (1990). The Emerging British Underclass. Choice in welfare. Institute of Economic Afairs.

Murray, S. O., J. H. Rankin, and D. W. Magill (1981). Strong ties and job information. Work and Occupations, 8 (1), 119-136.

Myers, C. A. C. A. and . Shultz, George P. (George Pratt) (1951). The Dynamics of a labor market : a study of the impact of employment changes on labor mobility, job satisfactions, and company and union policies. N.Y.: Prentice-Hall.

Oesch, D. and von Ow, A. (2017). Social Networks and Job Access for the Unemployed: Work Ties for the Upper-Middle Class, Communal Ties for the Working Class, European Sociological Review, 33(2), 275-291.

Ponzo, M. and V. Scoppa (2010). The use of informal networks in Italy: Efficiency or favoritism? The Journal of Socio-Economics 39 (1), 89 - 99.

Rabe-Hesketh, S. and A. Skrondal (2012). Multilevel and Longitudinal Modelling Using Stata, Volume II: Categorical Responses, Counts, and Survival (3 ed.). Applied Longitudinal Data Analysis: Modelling Change and Event Occurrence. Stata Press.

Rees, A. (1966). Information networks in labor markets. The American Economic Review 56 (1/2), pp. 559-566.

Russell, H. (1999). Friends in low places: Gender, unemployment and soci- ability. Work, Employment c Society 13 (2), 205-224.

Singer, J. and J. Willett (2003). Applied Longitudinal Data Analysis: Mod- eling Change and Event Occurrence. Oxford University Press, USA.

Varum, C., V. C. Rocha, and H. V. da Silva (2014). Economic slowdowns, hazard rates and foreign ownership. International Business Review 23 (4), 761 - 773.

Yakubovich, V. (2005). Weak ties, information, and influence: How workers find jobs in a local Russian labor market. American Sociological Review 70 (3), 408-421.

Appendix

Figure A1: Index plot of deviance residuals for the period of unemployment exit (person-period).

Figure A2: Index plot of sum of squares of deviance residuals for the period of unemployment exit (aggregate person).

Table A1: Summary statistics (person-period sample)

[1] The European platform against poverty and social exclusion is a 'flagship' initiative of the Europe 2020 strategy (European Union, 2015).

[2] Gallie et al. (2003) argue, for example, that the effect of unemployment on social isolation operates through poverty rather than directly.

[3] Beaman and Magruder (2012) conduct a laboratory experiment which shows that in response to defined objectives, incumbent employees do in fact screen their networks for casual workers who would enhance the firm's productivity. Jackson and Schneider (2011) studying New York City taxi owners' decision to hire drivers, show that members of the owners' social network, as defined by ethnicity, are preferred in order to mitigate problems of moral hazard.

[4] Consider a network of unemployed friends, one of which finds a job (say, through his/her local employment agency) and is hence able to pass on job-related information to others thus increasing their respective chances of exiting unemployment.

[5] People who are perhaps discouraged through long unsuccessful periods of job search. Following up a friend's suggestion is more likely than actively going through job postings.

[6] The unemployed in the sample only include individuals reported as such in the survey i.e. individuals on long-term leave, maternity leave, in FT education etc are excluded.

[7] Estimation of the models with exit from unemployment to a state other than employment has produced qualitatively unaltered results. These estimates are available from the authors upon request.

[8] Only a small proportion of our sample is left (10.1%) or right (5.5%) censored.

[9] We assume that a respondent's social network's composition does not change in the span of one year and hence carry over this information in the intervening years – this results in using all available 18 waves of the survey. The results remain unaltered if we instead consider each discrete time period to cover the span of two calendar years.

[10] We could consider these friends to be strong ties in the strong-weak ties classification made in much of the literature; however, in the current context such distinction would be rather inconsequential.

[11] A finer classification results in very low cell numbers in our dataset.

[12] Cappellari and Tatsiramos (2011) use a different methodological approach to ours and provide a detailed discussion of treatments for possible endogeneity in the context of their empirical strategy.

[13] Rabe-Hesketh and Skrondal (2012) show how the proportional hazards assumption for continuous-time survival data extends to the discrete-time hazard, which is parameterised using a generalised linear model with a complementary log-log link.

[14] We assume that the time-varying covariates are constant within the time intervals under consideration (Jenkins, 2005, p.44)

[15] The fitted hazard function, including the fitted baseline hazard function, can subsequently be recovered by re-writing 1 in terms of the original probability (Singer and Willett, 2003, p. 423).

[16] We sequentially drop redundant variables from the model in order to conserve degrees of freedom. The same applies to controls for friends' characteristics, such as gender and whether they are related to the respondent or not (23% of the sample report that one close friend is a relation, 9.6% two friends, 3.1% three and 64.1% report none).

[17] The full set of results is available online (Table B3).

[18] The full set of results is available online (Table B4).

[19] In excess of two standard deviations from the mean.