People I know:
social networks and job search outcomes.

Federico Cingano[‡] and Alfonso Rosolia[‡]

Bank of Italy,
Economics Research Department

PRELIMINARY

This version: December, 2003.

Abstract

We use a detail dataset of individual working histories to assess to what extent informal channels of job search affect the probability of exiting unemployment of previously displaced workers. We evaluate the effects of several measures of informal networks and find that all play a strong and significant role in determining the exit from unemployment.

Keywords: Social networks, job search, unemployment duration.

JEL codes: E24, J23, J64.

[‡]We thank Antonio Ciccone, Piero Cipollone and Sandro Secchi for useful discussions. Many thanks to Giuseppe Tattara and Marco Valentini for supplying and helping us with the data. We are responsible for any mistakes. The views expressed here are our own and do not necessarily reflect those of the Bank of Italy. Correspondence: Economics Research Department, Bank of Italy, Via Nazionale 91, 00184 - Roma - Italy. Ph.: (++39)0647922459, Fax: (++39)0647923720, email: federico.cingano@upf.edu, rosolia.alfonso@insedia.interbusiness.it.

1 Introduction.

There is widespread consensus that people find jobs and firms fill vacancies through informal channels. In particular, casual evidence and a number of studies show that jobs are often found through friends and relatives [cite Granovetter, Addison and Portugal, Blau and Robins, Montgomery, etc.].

The reasons put forth are manifold. On the one hand, a network of contacts allows individuals to collect information otherwise not easily accessible: agents are thus informed about employment opportunities by their acquaintances. On the other, if employers have imperfect information about the applicants they can improve on it by asking (references) to employees networked to the applicant.

An assessment of the workings and strength of these effects appears to be of interest for several reasons. First, if people do find jobs through a network of contacts, this may prevent them from moving to apparently better faring labor markets because of the (potentially high) cost of rebuilding such a network. Therefore, it could potentially account for (part of) the geographical differences in unemployment rates and wages. Second, it may explain why more concentrated labor markets (say, industrial districts) usually display better employment performances than the average. Third, knowing how relevant such a channel is may help design more efficient labor market policies that exploit this *externality*.

Economists have recently paid a great deal of attention to the fact that individuals are embedded in social networks and to its consequences on various aspects of one's life. As concerns the labor market, Bertrand, Luttmer and Mullainathan (2000) document how individuals belonging to groups relying more on welfare tend to use welfare more; Lalive (2003) shows that individual labor supply tend to change with that of a reference group; Topa (2001) develops a model where individuals know of job openings through their network of contacts. He then tests the empirical predictions of the model and finds a significant amount of social interactions among neighboring groups. It is not surprising then to find that people typically find their job through their network of social ties. This fact has been extensively documented in studies that investigate the methods of search for a job and their efficiency. [discuss the main contributions of the literature]

Although the idea that a larger and better network would improve on one's search outcome is always present a formal investigation of this aspect is hardly found in the empirical literature. This is generally due to lack of suitable data. In particular, it is very hard to recover information about *individual* social networks.

The purpose of this paper is to estimate how relevant are network size and quality in determining search outcomes. In particular, we investigate how both unemployment duration and entry wages (and ex-post job tenure) are affected by the extension of the social network. (Here we should say that our results do say something about whether larger networks are more efficient than smaller ones, a thing which cannot in general be stated) To this purpose we use a detailed dataset of individual social security records covering the period 1974-1997 to build several indicators of network extension and quality. We then explore the existence of systematic relationships between these indicators and the variables of interest.

To our knowledge only Wahba and Zenou (2003) develop an investigation such as ours. They build a model where [illustrate]. Using data from the 1988 Egyptian Labor Force Survey and using local population density as a measure of the extension of the network of weak ties, they find that living in a more densely populated area increases the probability of having found the current job through social ties. They take this as compelling evidence that larger social networks imply a more efficient flow of information about vacancies. They do not know for how long individuals have been searching. Moreover, since their measure of network extension has no individual variability, the coefficient they estimate may be the outcome of an omitted variable bias. For example, if highly skilled individuals move to cities and if employers hire through employees' referrals, than it is more likely that such a referral will have a positive outcome in a city than in another place since the quality of the referee is higher.

The paper is organized as follows. The following section reviews some relevant theoretical contributions on this issue and builds a workhorse model that sums up the main features of these models and will guide the empirical investigation. Next, we turn to the data. We first discuss how we use it to recover measures of individual social networks. We present the main results in section(4) and discuss the main methodological points; several robustness checks and additional results are also presented. Section (5) sums up the results and concludes.

2 A stylized model.

In this section we develop a workhorse model that will guide the following empirical analysis.

Let an unemployed agent i be embedded in a network of social ties of size N_i. Among the ties there are $U_i = qN_i$ unemployed individuals and $E_i = (1-q)N_i$ employed ones. Each individual receives a job offer drawn from a distribution $F(w)$ with probability p; the job offer is accepted provided the wage is sufficiently high. Otherwise the information is spread through the network. Therefore an employed contact will pass information about job openings on to agent i with probability $p_E = pF(w_E^*)$ while an unemployed one will do it with probability $p_U = pF(w_U^*)$ where $w_E^* > w_U^*$ are respectively the reservation wage of an employed and of an unemployed contact[1].

The following proposition can then be proved[2].

Proposition 1 *The probability of receiving at least an offer is* $P = 1-(1-p_U)^{qN_i}(1-p_E)^{(1-q)N_i}$. *Everything else equal, it increases with*

- *network extension* N_i,
- *the employment rate in the network* q,
- *the reservation wage of either the employed or the unemployed contacts,* w_U^*, w_E^*.

We then ask how the expected entry wage depends on one's network's features. To keep things simple assume that the unemployed individual collects all offers from her contacts, selects the best one and accepts it. Therefore, given she receives x offers from the unemployed and z from the employed contacts, the expected value of the best wage offer can be written as:

$$E\left(\max_{j\in\{X,Z\}}\{w_j\}\right) = \int_0^{w_U^*}\left(1-\frac{F(w)^{x+z}}{F(w_U^*)^x F(w_E^*)^z}\right)dw + \int_{w_U^*}^{w_E^*}\left(1-\left(\frac{F(w)}{F(w_E^*)}\right)^z\right)dw \qquad (1)$$

It is clear from (1) that the more offers are collected the higher the expected value of the best offer. Moreover, the expected value of the best offer increases more if the additional offer is received from an employed tie. an increase of the reservation wages of the (un)employed also improves the expected value of the best offer. Intuitively, better job offers will be passed on to the unemployed individual by her contacts if their reservation wage is higher.

Before we go on let us define the probability of collecting x offers from the pool of unemployed contacts U and that of collecting z offers from the pool of employed ones as follows:

$$\begin{aligned} P(x,U) &= \binom{U}{x} p_U^x (1-p_U)^{U-x} \\ P(z,E) &= \binom{E}{z} p_E^z (1-p_E)^{E-z} \end{aligned} \qquad (2)$$

[1]For simplicity we assume that the reservation wage only varies with the occupational status.

[2]We are neglecting the usual congestion externality due to the fact that a larger network implies also more competition for a given information. We basically assume that all information is spread through the network and then workers queue for a job.

then the expected entry wage is:

$$E(w; U, E) = \sum_{x=0}^{U} \sum_{z=0}^{E} P(x, U) P(z, E) E \left(\max_{j \in \{X,Z\}} \{w_j\} \right) \tag{3}$$

where X and Z are respectively the set of x and z offers collected from unemployed and employed contacts. The following properties can be shown to be true:

1. $E(w; U + s, E) - E(w; U, E) > 0$ and $E(w; U, E + s) - E(w; U, E) > 0$;

2. $E(w; U - s, E + s) - E(w; U, E) > 0$.

Property (1) states that a larger network improves the expected entry wage, independently of the composition of the network; property (2) states that, given network size, a higher employment rate also increases the expected entry wage.

Figure (1) shows, for a specific parametrization, how the expected entry wage moves with network size and the number of employed ties.

Figure 1: Expected entry wage.

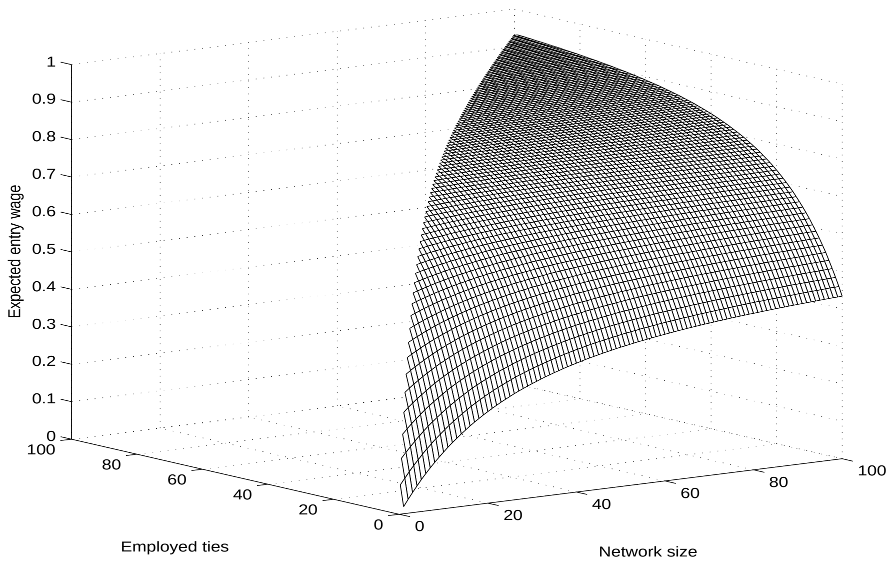

Summing up, our stylized model has several predictions. First, a larger network increases both the entry rate and the expected entry wage. Second, given network size, entry rates and entry wages are higher the more the employed contacts.

We now turn to the data to test these implications.

3 The data.

We use a detailed dataset of worker-firm matched social security records spanning the period 1975-97. The dataset covers the complete working histories of all individuals working at any moment between 1975 and 1997 in private non-agricultural firms located in two large provinces of the north-east of Italy, Treviso and Vicenza. Each record represents a year-worker-position and includes a string telling what are the months the worker was in that position. Therefore we are able to closely follow individual working histories at a monthly frequency. For example, we can establish whether in a given year an individual changed position within the firm (in this case there would be two records in our data, one for position a for the relevant months and one for position b). For each worker we also have information on several individual characteristics as well as characteristics of the firms she has worked in. Overall, we have information on 1.2 millions individuals and 105.000 firms.

We focus only on those individuals entering unemployment for exogenous reasons to control for possible selection bias. Typically, social security records do not collect information on why an individual left a job. Yet, the matched structure of our data allows us to isolate firm closures and consequently all workers displaced by the closure. In our data there are more than 2000 closures for a total of nearly 40.000 displaced workers.

For each displaced worker we build several measures of network size and quality, which will be detailed as we go on. Figure 2 illustrates how we construct some of the most important variables for the analysis. Let i_{lt} be an individual displaced at date t from firm l. The structure of our data allows us to precisely identify for each i_{lt} all individuals she has worked with prior to displacement. We simply look into all firms visited by i_{lt} in a given period prior to t and select all workers that where there while the firm was visited by our displaced individual i_{lt}. We restrict our attention to workers met up to five years before the date of displacement. In the rest of the paper we refer to this period as "*network building*" period (NB). Of course, for each coworker we are able to recover the same information available for the displaced individual (sex, age, place of birth, residence, qualification, working history, wages, characteristics of all the firms they have worked for, etc.).

Our most general definition of network thus includes all individuals one has worked with in the five years prior to displacement. This is clearly an approximation: we may not observe many other social ties which may turn out to be relevant in determining one's job search outcomes. Yet, we think our

Figure 2: The network of coworkers.

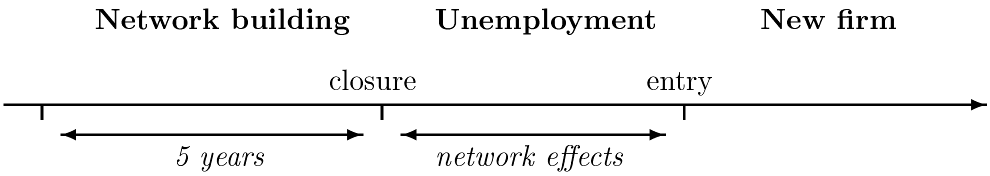

definition is reasonable for several reasons. First, the areas where these workers typically live and work are the same [illustrate discussing mobility in LLMs]. Second, past coworkers are likely to follow employment tracks very informative for our displaced individual: for example, they are likely to be employed in firms suited also to the displaced agent.

Once we have identified all of i_{lt}'s potential contacts in the five years prior to displacement, our data allow us to construct a large set of statistics on the characteristics of the network of potential acquaintances, ranging from the simple count of contacts to a careful description of their wage distribution, from the count of months i_{lt} actually overlapped with any of them in the same firm to the composition of the network by any of the individual characteristics covered by our data.

By following the displaced individuals after displacement date, we recover information about their entry wage, the duration of their unemployment spell, tenure in post-displacement job, characteristics of the firm she reenters to.

The next section lays out the empirical model and the definition of the measures of network size we use.

4 Results.

Our purpose is to identify (if any) the effect of the extension (and quality) of the network of social contacts on the outcome of the job search process.

We begin by considering, as measure of network extension, the overall number of people met by a worker during her network building period (henceforth, gross network). We relate it to both the (log of) individuals' unemployment duration (in months) and the (log of) post-displacement wage (at constant prices) with the following empirical specification:

$$y_{il} = \alpha + X_i\beta + Z_i\gamma + \epsilon_{il} \tag{4}$$

where X_i is a set of individual controls, including a quadratic of age at time of closure and of pre-displacement tenure, gender, qualification, type of contract and Z_i is the (log of) individual i's network extension; ϵ_{il} is an error term[3].

The first row in table (1) reports OLS estimates of γ both for unemployment duration and entry wage. The effect of a larger network on unemployment duration is positive and strongly significant, the opposite of what theory suggests; as to the response of entry wages, the OLS regression does not show any significant effect. Yet, a simple OLS regression has several problems. First, to build our sample we exploit firm closures occurred over a very long time span and regarding firms with different characteristics; therefore there are many factors which could potentially bias OLS estimates obtained from a simple comparison of individuals displaced from different firms. For example, consider two workers displaced from firms operating in different sectors: sector a experiencing a rise in employment in the period considered and the declining sector b, where no vacancies are opened in the same period. Assume that network forces are actually at work and affect post-displacement outcomes, and that the two workers are perfectly identical in any other respect. In particular, they have accumulated the same gross network. Still, if there are barriers to mobility across sectors, the worker exiting from b will experience longer unemployment (or accept lower entry-wage) simply because of the worse conditions in the relevant labor market. To control for this and similar problems arising from the structure of our sample we have estimated the same regression (4) allowing for a *closing firm* fixed effect. Results, reported in the second row of table (1) go in the expected direction. In particular, the coefficient on network size in the unemployment equation switches sign and is still strongly significant: doubling the network size shortens unemployment duration by 10%, reducing the average unemployment spell by nearly one moth. As concerns entry wages, we still find a positive coefficient, now estimated with considerably more precision: raising the network size by 100% increases entry wages by about 1.5%.

At this point one might be worried that the results presented above are actually driven by individual unobserved heterogeneity. For example, (explain with an example).

Let us give more structure to the error term in equation (4) and assume that it is the sum of two components, a firm effect (μ_l) and an individual effect (η_i):

$$y_{il} = \alpha + X_i\beta + Z_i\gamma + \mu_l + \eta_i \tag{5}$$

[3]We do not estimate the model by SUR since the information set is common to both equations, granting that SUR is fully equivalent to separate LS. As regards unemployment duration, the estimated regression preserves consistency of coefficient estimates while implying only a loss of efficiency as compared to the estimation of a duration model.

For unobserved individual heterogeneity to bias our estimates of γ it must be that individual controls X_i do not fully proxy for η_i and that the measures of network extension Z_i are correlated to η_i, that is workers whose unemployment duration is on average shorter or who earn higher wages also have larger networks. As concerns wages, this may happen if, say, larger firms hire better workers, so that a more productive worker also has a larger gross network; with respect to unemployment duration, this correlation may arise if, say, a worker whose unemployment spells are on average shorter also tends to work with people who are more likely to be employed so that her employed network is also larger. Note that, in order to work, these two examples require that workers are sorted across firms according to some unobserved characteristic that also affect wage or unemployment duration: in the first one, better workers work in larger firms; in the second one, workers with shorter unemployment duration tend to work together. Therefore, if *good* workers only work with comparably *good* workers at all times in equilibrium all firms (and thus networks) would be made of individuals sharing similar (unobserved) characteristics and therefore performances [reference to the literature on sorting]. Yet, if this is the case then it is also true that in our cross-section firm characteristics are correlated to individual ones, so that allowing for a *closing firm* fixed effect allows to capture individual unobserved heterogeneity. On the other hand, if the labor market does not sort workers in this way then the issue of spurious correlation (omitted variable bias) is groundless.

Still, one could argue that the very occurrence of a closure might have been caused by the fact that workers were not rightly sorted and the firm turned so unproductive that it exited the market. In this case the *closing firm* fixed effect would not capture the individual unobserved heterogeneity: η_i might still be correlated with network measures Z_i since *in the past* the worker was rightly sorted (but then one should ask why she changed job: recall it was not because of a closure!). In this case our estimates would still be subject to the criticism that unobservable fixed individual characteristics (as ability) which affect the speed of reallocation (or entry wage) might be correlated to the measures of network extension. To tackle this issue we include in the control set also measures of individual past employment experience and pre-displacement wage. The idea is that inclusion of these two controls should proxy for any unobserved fixed individual characteristics which systematically affect wages or individual turnover[4]. Results for this extended control set are reported in the last panel of table (1),

[4]Robustness checks have been conducted on an even more conservative formulation of the control set. We have further exploited our rich dataset to account empirically for unobserved individual heterogeneity affecting both the speed of reallocation (or the entry-wage) and the type of network established in the labour market, but that might not be captured by firm fixed effects if workers were not sorted in the typical way at the time of closure. Specifically, for

where we also report the coefficients on the two additional controls. The results suggest that indeed there may have been a problem of unobserved heterogeneity since the estimated coefficient on network size in the unemployment equation, although still strongly significant, is about a half of our previous FE estimate; the one in the entry-wage equation becomes non significant. The additional controls also have the expected signs, suggesting that indeed they capture unobserved individual heterogeneity: attachment to employment, as captured by past participation to the labor market, and wage at displacement are negatively correlated with the length of the unemployment spell and positively with subsequent wages. Our results imply that shifting network size from the 25th percentile to the 75th of the distribution of gross network cuts unemployment duration down by about 7.5% (shortening the average unemployment spell by three weeks approximately.

Table 1: Network extension.

	Unemployment duration			Entry wage		
	Coeff.	t-stat	R^2	Coeff.	t-stat	R^2
				OLS.		
Gross Network	0.0847	8.4	0.014	0.0047	1.37	0.19
			Closing firm fixed effects.			
Gross Network	-0.0984	-6.52	0.017	0.014	2.66	0.09
	Closing firm fixed effects and extendend controls.					
Gross Network	-0.055	-3.45	0.024	-0.003	-0.6	0.1
Exit wage	-0.152	-6.26	-	-0.118	13.1	-
Participation	-0.35	-7.55	-	0.139	8.17	-

4.1 Appropriateness of our measures of network

The main drawback with our measures of network extension is that it is likely to be affected by serious measurement problems. As discussed above, it is obtained by simply counting the number of people

each displaced worker we computed the entire wage and unemployment profiles during the network building period (the relevant five-years time spell in our exercise). Results were basically unaffected.

met during the network building period. Therefore it may imprecisely measure the number of *true* contacts since these are likely to be only a subset of all co-workers met in each single firm, the more so the larger the firm visited or the more firms visited during the network building period. Yet, if this error is a standard measurement problem (i.e. the error itself is uncorrelated with the variable of interest and with the residual in equation (4), it simply biases the estimates of the coefficients of interest towards zero. If this was the case we would expect to find a stronger effect of network extension when taking care of the measurement issue. One could still argue that since the measurement strategy is such that the size of the network will be, all else equal, larger the larger the firms visited or the more firms visited in the past, the measure may still capture some unobserved characteristics that are also likely to affect our dependent variables. For example, if larger firms are able to attract better workers by offering higher wages good workers will have larger networks on average. By the same token, an individual who changes job more often will on average be assigned a larger network; if this *mobility feature* reflects some unobserved characteristic, the estimated coefficient could possibly reflect still some omitted variable bias.

We have argued in the previous section that the omitted variable bias is dealt with by including variables such as past wages and past unemployment experience that should proxy for all kinds of unobserved heterogeneity affecting the dependent variable. For example, if larger firms attract better workers by offering higher wages inclusion of past earnings controls for the problem. Yet, one could again counter-argue that larger firms attract better workers by offering, for example, non monetary benefits so that earnings do not fix the problem.

In principle, we could control for this problem by including in our estimating equation a measure of the average size and the number of firms visited in the past. In this case, though, we would not be able to cleanly interpret the coefficient on network size because the three measures are clearly related to each other: as we measure it, network size increases either because one visits more firms of because these firms are on average larger.

To tackle the measurement issue and the potential residual effects of endogeneity we exploit the richness of our dataset and further refine our measure of network size along reasonable dimensions. The implicit assumption that a displaced worker knows all individuals met in the five years prior to displacement is clearly too ambitious, the more so the larger the firms visited. By the same token, it is unlikely that all individuals in the network play the same role in affecting one's job search outcome. We therefore propose several alternative refinements of the measure of network size and perform a first robustness check of our previous results. In particular, we split our measure of gross network along some reasonable dimensions and estimate a *substitution* effect between the average network member and the ones belonging to a specific subset. Formally, we estimate the following equation:

$$y_{il} = \alpha + X_i\beta + \gamma \log(N_i + \lambda(Z_i - N_i)) + \mu_l + \eta_i \tag{6}$$

where gross network Z_i has been split into individuals belonging to a specific subset N_i (henceforth core network) and those not belonging to it. Therefore, λ captures the substitutability of these two types of network members.

We start out investigating four alternative dimensions of selection, discussed below.

As discussed in section (3), the available observations concern a very integrated area, characterized by high commuting and self-contained labor markets. This area is split into small administrative units (municipalities) where most of the people in the dataset lives. These units are relatively small [illustrate with statistics by age cohort, place of work/residence], with people commuting to the workplace [illustrate using census and multiscopo] and characterized by a high population density. Table (2) reports population and family density for the area covered by our dataset and for Italy as a whole: both densities are far above the average. Wahba and Zenou (2003) use a similar information about city density to capture the tightness of social ties, arguing that more densely inhabited area expose individuals to more social contacts, extending their network of ties.

Table 2: Population and family density (per km^2).

	Population density	Family density
Sample	304.2	113.1
Italy	186.9	71.4

Therefore, we think that focusing on the members of the network who live in the same municipality as the displaced individual, that is those with whom she is more likely to have tighter ties (commute together, meet at the bar, children in same school, etc.) could provide more clear-cut evidence on the workings of the network of contacts. Again, we must control for measurement errors: for example, an individual living in a larger municipality will on average have a larger imputed network; to this purpose we include population density in the municipality where the displaced individual lives. Moreover we retain the original network measures in the regression. A second dimension of analysis is individual's age. We split gross network according to displaced worker's age, a_i: N_i includes all network members whose age is within the interval $[a_i - 4, a_i + 4]$. Third, we investigate job qualification (blue vs. white collar) and let N_i include only network members sharing the same job qualification as the displaced worker (this should proxy for education). Fourth, we split gross network according to gender: N_i only includes network members of the same gender as the displaced individual.

While we expect individuals belonging to the same municipality or to the same age cohort to establish tighter ties with the displaced worker and thus potentially play a more relevant role in her job search process, we do not have a clear prior as to the job qualification split. On the contrary, we do not expect gender to be a relevant dimension for one's job search.

Table (3) reports estimates for γ and λ in the unemployment duration equation along with the implied effect on duration of a shift of network size from the 25th to the 75th percentile of its distribution for any of the network splits introduced above.

Table 3: Network composition and substitutability.

Core Network	γ	λ	Implied effect
Municipality	-0.056	0.7	9.8%
Age cohort	-0.056	1.7	10.4%
Qualification	-0.054	3.5	10.5%
Sex	-0.057	1	9.9%

Most of the results on network composition are in line with our expectations. We find a strong difference between network members living in the same municipality and those outside it, with the former having a stronger effect on one's job search process. Our estimate of the substitution parameter λ implies that, to the purpose of reducing the displaced's unemployment spell, increasing the network size measure by one neighbor is as effective as adding 1.4 individuals living in a different municipality. The implied effect of our thought experiment is a reduction of about 10% per cent of unemployment duration. As to age cohort, results show that having more individuals belonging to the same age group has a negative impact, since they substitute unfavorably with individuals of other age classes. This might be due to the fact that age also is a proxy for experience: since most of our displaced workers are younger than 30 we might be mixing up the effects of the tightness of social ties, likely to be stronger with people of your age cohort, and the efficiency at helping in one's job search, likely to be increasing with labor market experience and thus age. As expected, the gender split turns out to be ineffective: for one's job search it doesn't matter whether your acquaintances are female or male.

We further refined our network measure combining the two main results emerging from the previous exercise, and pointing in particular to the importance of co-workers living in the neighborhood and belonging to different age classes with respect to the displaced worker. Accordingly, we estimated a version of equation (6) where N_i includes the subset of i' s contacts who live in the same municipality and belong to a different age cohort. The results are reported in the second row of table 4 (the first row just replicates table 3) indicate that, among neighbors, network members of similar age do not help one's job search process. This can be inferred both from the fact that the "residual" network becomes a worse substitute of the core (λ falls from 0.7 to 0.5), and that the implied effect of an increase in network size is slightly larger than before. To see this even more clearly we allowed for the following decomposition of the gross network

$$\widetilde{Z} = N_i + \lambda_1 A_i + \lambda_2 (Z_i - A_i - N_i)$$

where N_i is the core of the network, A_i is a specific subset of i' s contacts we are interested in (the "focus" group) and the last term is the residual network. This specification allows to estimate the substitutability of the "focus" group separately and increases the flexibility of our empirical approach. The result ($\lambda_1 = 0$) confirms the fact that neighbors of the same age cohort as the displaced workers have no role in explaining network effects in our framework.

Table 4: Further refinements in the network measure.

Core Network (N_i)	Focus group (A_i)	λ_1	λ_2	Implied effect
Municipality	-	0.7		9.8%
Different age cohort in munic	-	0.5		10,1%
Different age cohort in munic	Same age cohort in munic	0	0.5	10,9%

4.2 Network quality

A second exercise consists on focussing on the "quality" of the Network by considering how the effects of our measure change once we allow for a different impact of members according to their work status. As largely acknowledged in the literature, the probability to find a job through personal contacts is decreasing in the network rate of unemployment, since unemployed contacts are less likely to pass the information on existing vacancies than employed workers. Our data allow to split the gross network measure into members who are employed at the time of i's displacement and those who are not. This provides us with a precise index of the network unemployment rate that varies at the individual level. We proceeded along the same lines illustrated in the previous section, and estimated equation (6) considering as core network N_i the subset of employed contacts of each displaced individual. The first row in table 5 reports the estimated elasticity and substitution coefficients. Not only we find that, in line with theoretical predictions, the information transmitted by unemployed contacts is very poor, but we actually estimate no role for such contacts in terms of unemployment spell reductions. Our estimates imply that the usual 25th-to-75th percentile increase in the size of *employed* contacts (as opposed to an increase in network size holding the unemployment composition of the network constant) reduces unemployment duration by more than 12 per cent. Hence, the possibility to asses the "quality" of the network, as measured by its employment share, is very relevant for the assessment of the effects of social networks on labor market outcomes.

As a final exercise, and given that proximity has proven to be a relevant dimension for the refinement of our measure, we considered a decomposition of the employed network according to the members' place of residence. The results, reported in the second row of table 5, are in line with our previous findings. First, the implied effect of an increase in size is larger when social networks are measured by the count of employed contacts only rather than as networks with the average unemployment rate. Second, the unemployed contacts are not substitute of the employed contacts: no matter how large the increase in their number they are not effective in terms of reducing experienced unemployment in our sample. Finally, geographical proximity matters even among the employed: the estimated rate of substitution implies that the information provided by neighbors is nearly 25% more valuable than that of other employed network members.

Table 5: Network quality.

Core Network	Focus group (A_i)	γ	λ_1	λ_2	Implied effect
Employed	-	-0.075	0		12,2%
Empl. in municipality	Empl. not in municipality		0.8	0	12,4%

5 Conclusions.

Casual evidence and sociological studies indicate that individuals rely on acquaintances when looking for a job. This has led to the development of theoretical models linking (endogenous) social networks and job search outcomes. These models generally conclude that, all else equal, individuals belonging to larger and "better" networks find jobs more quickly and end up earning higher wages. In this paper we address the relationship between network size, composition and quality on the one hand and unemployment duration and entry wages on the other using a comprehensive dataset of matched worker-firm social security records. To our knowledge, there has been no empirical analysis on the nature and strenght of these relationships, mostly due to the difficulties in measuring the network individuals have access to.

Our preliminary results are in line with theoretical predictions and point to statistically significant and robust network effects on unemployment duration. Findings on post-displacement wages are somehow weaker. The quantitative importance of networks seems to be rather limited however. Throughout the empirical analysis we exploit our data on individual working histories to argue that these effects are not likely to be driven by individual heterogeneity, we discuss the appropriateness of our network measures and deal with potential sources of endogeneity.

References

Bertrand, M., Luttmer, E. F. and Mullainathan, S. (2000), 'Network Effects and Welfare Cultures', *Quarterly Journal of Economics* **11**, 1019–1055.

Lalive, R. (2003), Social Interactions in Unemployment, *mimeo*.

Topa, G. (2001), 'Social Interactions, Local Spillovers and Unemployment', *Review of Economic Studies* **68**, 261–295.

Wahba, J. and Zenou, Y. (2003), Density, Social Networks and Job-Search Methods: Theory and Application to Egypt, *mimeo*.