The Role of Coworker-based Networks in the Labour Market

Albrecht Glitz [1]

Introduction

In 1974, Mark Granovetter published his highly influential monograph "Getting a Job: A Study of Contacts and Careers" (Granovetter 1974) in which he showed that in the sample of professional, technical and managerial workers he surveyed in the Boston area, 56 percent found their current job through personal contacts. Similar numbers have since been documented for many other countries, identifying social networks as one, if not the most important channel through which workers find their jobs. The prominent role played by informal contacts in the job search process naturally raises the question of whether their use leads to better labour market outcomes for workers. On the one hand, social networks may foster the exchange of information about suitable employment opportunities and improve the match between workers and firms. On the other hand, if social contacts are maintained primarily for purposes other than providing information on jobs, they may only help workers to find jobs in a limited set of occupations or segments of the labour market – those that are prevalent in their social network – and thus prevent them from fully exploiting their productive potential (Bentolila, Michelacci and Suarez 2010). In addition, favouritism towards social network members may lead to job matches that are worse than those obtained through formal channels. A priori, it is therefore unclear whether the extensive use of informal contacts in the labour market is economically beneficial.

The theoretical ambiguity regarding the role of social networks in this context has motivated a substantial body of empirical analyses (see Topa 2011, for a comprehensive survey of the literature). The first such studies were usually based on data from surveys, in which respondents were directly asked about the channel through which they obtained their jobs. Comparing different labour market outcomes like wages or job stability across different search channels would give a first indication of their relative effectiveness in creating good job matches. While the evidence is not unambiguous, a majority of such survey-based analyses suggests that informal job finding methods play a positive role in workers' labour market outcomes (Ioannides and Loury 2004). However, a fundamental problem in these types of studies is that workers of different abilities tend to self-select into different job search methods, making it hard to distinguish the effect of a particular job search method from that of the workers' abilities themselves.

In the absence of direct survey information on job finding methods and social interactions, an alternative approach followed in the literature is to focus on the specific dimensions of one's social network that are likely to proxy for such interactions, and indirectly test for the presence of network effects in the labour market. The network dimensions studied include such diverse social groups as neighbours (Bayer, Ross and Topa 2008; Damm 2014; Schmutte 2015), individuals with the same ethnic background (Munshi 2003; Edin, Fredriksson and Aslund 2003; Dustmann, Glitz and Schönberg 2011), close friends (Cappellari and Tatsiramos 2010) and family members (Kramarz and Skans 2015). The evidence from these studies, which typically relate the employment status of a worker to the prevailing employment rate in his network, again points to a positive role of social networks in the labour market.

Despite its extensive coverage, an important dimension of a worker's social network that has long been neglected in the empirical literature is work-related contacts. In fact, in Granovetter's original study, 69 percent of the social contacts through which workers found their jobs were actually work-related contacts such as former colleagues or bosses, and only 31 percent were family members or friends. The strong reliance on work-related contacts is not surprising given that former coworkers are likely to have relatively accurate information about a worker's skill set by having worked alongside this work-

[1]   Humboldt University Berlin.

er for some time. Because they are usually attached to a similar segment of the economy, former coworkers are also more likely to hear about relevant job openings than, say, neighbours or family members.

Despite their potential importance, until recently, a lack of suitable data prevented a systematic analysis of the role of coworker-based networks in the labour market. Only the access to large administrative data sets that was provided to the scientific community in several countries over the last decade has allowed researchers to accurately measure workers' networks of former coworkers and analyse their impact on labour market outcomes. In the following section, I will briefly outline the two main theoretical mechanisms through which coworker-based networks may affect labour market outcomes.[2] In the remainder of this report, I will then summarise two of my own recent studies (Glitz 2013; Glitz and Vejlin 2015) and relate them to the still relatively nascent literature on coworker-based networks. Both studies show that coworker-based networks contribute positively to workers' labour market outcomes, both through their positive effect on the flow of information about job opportunities, and through their ability to reduce uncertainty by means of referrals.

Theoretical mechanisms

The theoretical literature distinguishes two primary mechanisms through which social networks may affect workers' labour market outcomes. The first set of models views informal contacts as additional sources of information about job opportunities. A key reference in this literature is the work by Calvó-Armengol and Jackson (2004, 2007) who embed a model of informal job search into an explicit network structure. The basic idea is the following: in every period, agents hear about job openings with an exogenous probability. If the agent is unemployed, he takes the job. If the agent is employed, and if the job opportunity does not dominate his own current job in terms of wages and job characteristics, the agent forwards the information about the job to one of his unemployed network contacts. The network thus effectively increases an unemployed worker's job offer arrival rate, which leads to faster transitions out of unemployment and higher starting wages due to higher

reservation wages. In Glitz (2013), I investigate precisely this mechanism by comparing the speed with which unemployed workers in Germany exit unemployment as a function of the number of employed contacts in their coworker-based network.

The second theoretical mechanism through which social networks affect labour market outcomes is through their ability to reduce uncertainty. The starting point here is that firms are not able to perfectly observe a job applicant's skills before hiring, and that workers, in turn, do not have full information about an employer's characteristics. In this situation, social networks, by means of referrals, can provide additional information about the suitability of a worker for a given job. There are two ways in which the reduction in uncertainty in this type of framework is modelled. In Montgomery (1991), the starting point is the assumption that employers are able to identify high-ability workers in their own workforce. In the presence of assortative matching in social networks, which implies that high-ability workers tend to be connected to other high-ability workers, firms then have an incentive to ask their high-ability workers for referrals, since this increases their chances of hiring other high-ability workers. By relying on referrals, firms are thus able to mitigate the typical adverse selection problem they face when hiring new workers. Using Swedish matched employer-employee data, Hensvik and Skans (2013) systematically test this particular mechanism of employee referrals in the context of coworker-based networks, providing strong support for the key predictions of the model.

In contrast to Montgomery (1991), Simon and Warner (1992) argue that referrals more generally reduce the initial uncertainty about the match-specific productivity when workers and firms meet, allowing the latter to hire workers that are, on average, better matched than those hired through external channels. This higher match quality is reflected in initially higher wages and greater job stability for referred workers relative to externally hired workers. However, a key empirical prediction of this type of model is that, due to the subsequent on-the-job learning about workers' true match-specific productivity, the initial differences in wages and job stability diminish over time. This is because those workers who turn out to be bad matches, most of which were hired through the external market, are successively laid off over time, so that in the long run only sufficiently good matches remain in the firms, independent of their initial hiring channel. In Glitz and Vejlin (2015), we systematically test these predictions using Danish administrative

---

[2] There are additional conceptualisations of the role of social networks and referrals in the labour market that have been studied in the literature such as their mitigating influence on moral hazard through monitoring (Heath 2013) and their impact through favouritism towards network members (Beaman and Magruder 2012).

Table 1

Summary statistics					
	Mean	Standard deviation	10th percentile	50th percentile	90th percentile
Before closure					
No. of coworkers (last 5 years)	132.6	228.1	9	43	379
Share of coworkers working in t (in %)	58.4	18.0	35.1	60.0	79.1
Duration of cowork	3.0	2.9	1	2	7
Last log wage in closing establishment	4.37	0.47	3.90	4.35	4.91
After closure					
Share of coworkers in mass-layoffs (in %)	71.6	45.1	0	1	1
Log wage in t+1	4.38	0.44	3.93	4.35	4.90
Instrumental variable					
Share of coworkers in mass-layoffs (in %)	4.2	10.9	0.0	0.0	10.0

Source: Glitz (2013), table 2.

data, proxying whether or not a worker obtained his job through a referral by the presence of former coworkers in the firms by which they were hired.

Coworker-based networks as providers of information about job opportunities

The objective of the study in Glitz (2013) is to assess the role of coworker-based networks as providers of information about job opportunities, building on the aforementioned theoretical work by Calvó-Armengol and Jackson (2004, 2007). As a natural starting point for the empirical analysis, I focus on a sample of displaced workers who are pushed into unemployment as the result of a firm closure, and thus start looking for a new job through both formal channels and their social network at the same time. The theory predicts that those workers who are embedded in a strong network characterised by a high employment rate should exit unemployment more quickly than workers who, at the time of displacement, are embedded in a weak network characterised by a low employment rate. Furthermore, conditional on working, the starting wages of workers with a strong network should be higher than those of workers with a weak network because the former set higher reservation wages due to their higher job offer arrival rate.

The set-up of the empirical analysis is closely related to Cingano and Rosolia (2012) who were the first to operationalise the notion of coworker-based networks using comprehensive social security data and to estimate their impact on labour market outcomes. For my analysis, I exploit high-quality German administrative data cover-

ing the universe of workers in the German labour market between 1980 and 2001 to obtain an accurate picture of each worker's coworker-based network. Focusing on the four largest metropolitan areas in West Germany (Hamburg, Cologne, Frankfurt, and Munich), I start by selecting all establishments that closed down in either 1995 or 1996 and had between 5 and 50 workers in their last year of operation. This gives a sample of 1,814 closing establishments from which 10,916 workers are being displaced. For each of these workers, I then find every other worker whom they worked with in the same establishment during the so-called network building period, which, in this paper, is defined as the five-year period prior to the year of displacement. These former colleagues form the respective coworker-based network of each displaced worker. Table 1 provides a selection of summary statistics. On average, a displaced worker's network consists of around 133 former coworkers, with a median number of 43, indicating that the network size distribution is strongly right-skewed. At the time of displacement, the average employment rate in these networks is 58.4 percent, but again there is substantial variation, with some networks featuring employment rates of less than 35 percent (10th percentile) and others boasting employment rates of over 79 percent (90th percentile). The empirical question is whether these differences in network employment rates at the time of displacement have a significant effect on a displaced worker's re-employment probability and starting wages in subsequent years.

Figure 1 illustrates the empirical set-up. In this stylised example, in which the horizontal axis represents the time line, the sample of displaced workers consists of

two workers, red and blue, who become unemployed in 1995 as the result of a firm closure. During the network building phase between 1990 and 1994, both established links with two other workers (co-displaced workers are excluded from the network). The red worker met one of his network contacts in Firm A (first row from the top) and the other contact in the closing firm (third row from the top), just before that contact left for Firm B (second row from the top) in 1993. The blue worker met both his contacts while working in the closing firm, before they then left for Firm C (fourth row from the top) in 1991 and for Firm B in 1993. The contact moving to Firm B is thus a member of both displaced workers' networks, so there can be partial overlap in the networks of former coworkers. However, due to their distinct employment histories, different displaced workers generally tend to have different networks of former coworkers. Now, at the time of displacement in 1995, both network contacts of the red worker are employed, so that the employment rate in his network is 100 percent. By contrast, one of the contacts of the blue worker is unemployed in 1995, so that the employment rate in his network is only 50 percent. According to the theory, the red worker should therefore be more likely to be employed and earn higher wages in 1996 (and potentially subsequent years) than his blue counterpart.

To assess these predictions, I first estimate a standard OLS regression, regressing either an indicator variable for being employed one year after displacement or the log wage in the new job on the network employment rate, the log size of the network and a comprehensive

set of individual control variables. To account for any unobserved shocks at the time of displacement that may be specific to the workers of a given establishment, the specification also includes a full set of closing establishment fixed effects, ensuring that the comparison in post-displacement labour market outcomes is between workers who are laid off from the same establishment and therefore face presumably comparable labour market conditions.

However, even after conditioning for establishment fixed effects and a comprehensive set of control variables, it may still be that the relationship between a worker's outcomes and the prevailing employment rate in his network is a reflection of common unobserved skills that are shared by these workers. If high-ability workers tend to work in the same establishments, then a positive correlation between the network employment rate and the labour market outcomes of a displaced worker may simply reflect these common unobserved skills, rather than a causal network effect. To deal with this endogeneity issue, as well as potential measurement error that may bias the estimates towards zero, I exploit the occurrence of mass-layoffs – the reduction in a large firm's workforce by over 30 percent from one year to the next – as exogenous shocks to the employment rate in a network. Figure 1 again illustrates the key idea. After moving to Firm C, one of the blue worker's network contacts was laid off from that firm in 1994 as part of a mass-layoff and, as a result, continued to be unemployed in 1995. Under the assumption that, conditional to establishment fixed effects and observable characteristics, mass-layoffs are unrelated to the displaced workers' unobserved skill sets, they can serve as exogenous shocks for a given worker's network employment rate, qualifying them as suitable instruments in an instrumental variable (IV) framework.

Table 2 presents the most important empirical results from both the OLS and IV regressions. The OLS estimates show a small positive effect of the network employment rate on both the re-employment probability and the starting wages of displaced workers. According to the specification that includes establishment fixed effects reported in columns (2)
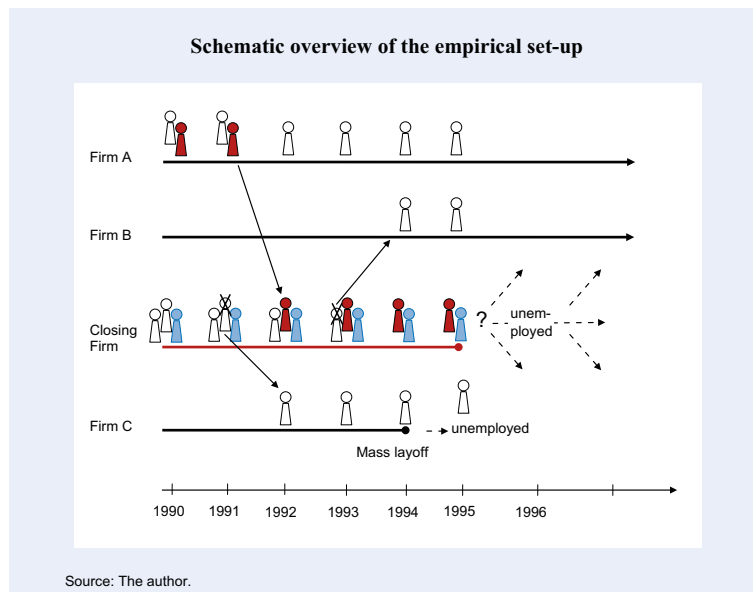
Figure 1

Schematic overview of the empirical set-up

Source: The author.

Table 2

	(1)	(2)	(3)	(4)	(5)	(6)
	\multicolumn Employment and wage effects					
	Employment probability			Log wages		
	OLS		IV	OLS		IV
Employment rate coworkers	0.073**	0.051	0.753*	0.054**	0.067*	0.174
	[0.031]	[0.048]	[0.387]	[0.023]	[0.036]	[0.264]
Log number of coworkers	0.002	0.004	-0.015	0.006	0.008	0.005
	[0.005]	[0.007]	[0.013]	[0.004]	[0.005]	[0.009]
Establishment fixed effects		yes	yes		yes	yes
Observations	10.916	10.661	10.659	7.789	7.435	7.433
1st stage statistics employment rate coworkers						
Share coworkers in mass-layoffs			-0.131***			-0.138***
			[0.016]			[0.021]
F-stat 1st stage			65.1			42.7

Note: The instrumental variable for the coworker employment rate is the share of working former coworkers who, after separation, worked in a large establishment (>50 employees) and separated from that establishment as a result of a mass-layoff. Standard errors are robust and clustered at the closing establishment level. A (*) denotes statistical significance at the 10% level, a (**) at the 5% level, and a (***) at the 1% level.

Source: Glitz (2013), table 2.

and (5), a 10 percentage point increase in the network employment rate is associated with a 0.51 percentage point and 0.67 percent increase in the re-employment probability and starting wages of displaced workers, respectively (with the coefficient for the employment effect, however, not being statistically significant at conventional levels).

In the presence of endogeneity or measurement error in the network employment rate, the estimated OLS parameters are likely to be biased. Columns (3) and (6) show the results when the network employment rates are instrumented with the shares of coworkers who were themselves part of a mass-layoff after separation from the displaced workers.[3] As the summary of the first stage regression at the bottom of the table shows, this share has a strong negative effect on a network's employment rate in the year of displacement: a ten percentage point increase in the share of laid-off coworkers during the network building phase reduces their employment rate in 1995/1996 by 1.3 percentage points. Exploiting the variation in employment rates that is driven by these mass-layoff shocks, the second-stage results indicate a

strong positive effect of the network employment rate on the re-employment probability of a displaced worker: a 10 percentage point increase in the employment rate, which corresponds to about half a standard deviation, leads to a 7.5 percentage point higher probability of being employed in the first year after displacement (p-value 0.052). This estimate is larger in magnitude than both the main estimate in Cingano and Rosolia (2012) and the corresponding figure in a recent study for Austria by Saygin, Weber and Weynandt (2014), both of which, however, restrict their analysis to the baseline OLS regression.[4] In contrast to the positive employment effect, the IV estimate for the effect on starting wages reported in column (6) of Table 2 is not estimated precisely enough to draw any strong conclusions.

Further analysis shows that these results are very robust to changes in the way the instrumental variable is constructed, as well as different functional form assumptions and sets of included control variables. More importantly, a number of placebo estimations provide strong suggestive evidence of the exogeneity of the instrumental variable, by showing that contemporaneous mass-layoffs are unrelated to past labour market

---

[3] It is worth noting that, in the absence of a suitable separate instrument, the network size is included as a potentially endogenous control variable. I provide more details of the conditions under which this leads to consistent estimates of the coefficient on the network employment rate in the paper.

[4] For a recent analysis that focuses explicitly on the role of coworker-based networks for immigrants' labour market outcomes in Italy, see Colussi (2013).

outcomes and that future mass-layoffs are unrelated to current labour market outcomes. Although statistically inconclusive, there is also some indication that low-educated workers, immigrant workers and young workers benefit relatively more from increases in the employment rate in their networks. Looking in greater detail at heterogeneity across different sources of information, the evidence also suggests that female coworkers, coworkers from the same age cohort as the displaced workers, and coworkers with whom prior interaction was more intensive are particularly important for post-displacement employment outcomes. Finally, an analysis looking at the longer run effects reveals that the positive impact on the employment probability of a displaced worker is short-lived, only persisting for the first year after displacement.

Overall, the findings of this study suggest that a strong network of coworkers provides valuable information about labour market opportunities and can serve as a useful resource to accelerate transition out of unemployment in times of economic distress.

Figure 2

Hourly wage and job turnover convergence

Wage difference

Turnover difference

Source: Glitz and Vejlin (2015), Copenhagen area.

Coworker-based networks as providers of information on match quality

Rather than focusing on job offer arrival rates and transitions out of unemployment, Glitz and Vejlin (2015) investigate the extent to which coworker-based networks help reduce uncertainty about match-specific productivity in the labour market, building on the theoretical framework by Simon and Warner (1992) and its extensions by Dustmann et al. (2011). By providing information on network contacts' otherwise unobservable skills to their employers, incumbent workers can, through referrals, help to increase the match quality between new hires and their firms, ensuring that the right worker is matched to the right firm. As discussed above, in a world of initial uncertainty about match-specific productivi-

ty and subsequent learning by the employers about the quality of a match, being hired through a referral is predicted to lead to higher starting wages and more job stability in the initial phase of the employment relationship than being hired through external channels, but that this difference subsequently declines with tenure in the firm.

To take these predictions to the data, we use comprehensive administrative records on the entire Danish population, spanning the period 1980 to 2005, and, similar to Glitz (2013), define each worker's network as every person with whom s/he worked together in the same establishment during the network building phase (which, in this study, comprises the previous ten years). In the absence of direct information about referral use – a typical feature of this type of administrative data – we argue that observing a worker who follows one of his/her former coworkers into the same establishment can be used as a proxy for a referral (see also Hensvik and Skans 2013 and Saygin et al. 2014, for a similar strategy). In the data, over 30 percent of new hires start at
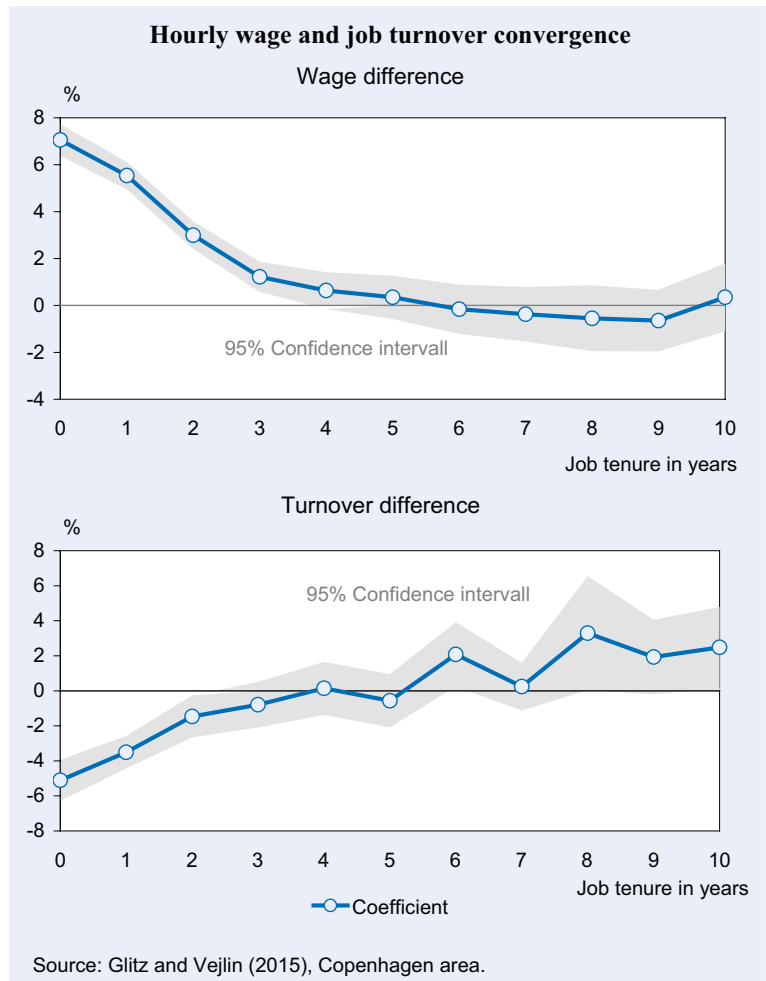
a firm with at least one of their former colleagues already present. Obviously, such mobility patterns could also arise in the absence of an actual use of referrals, for example by chance when labour markets are very thin or because workers who used to work in the same establishment share similar skills which, in turn, make them prone to look for jobs with similar employers. However, we show that only about half of the observed propensity of workers to move where their former coworkers are already working is due to either random meetings or similarity in workers' characteristics, so that this type of mobility can indeed be viewed as being driven, at least in part, by personal interactions of workers with their network members.

In the next step of the analysis, we estimate several OLS regressions in which we relate either the log wage of a worker or an indicator variable for leaving one's firm by the next period to a dummy variable that takes the value one if there is at least one former coworker present in the firm at the time of hiring (our proxy of having obtained the job through a referral). Besides the dummy for the presence of a former coworker, we also include its interaction with individual dummy variables for each year of tenure in the firm. The estimated coefficients of these interaction terms allow an assessment of the convergence predictions derived from the theoretical model, according to which the initial wage and job stability advantage of referred workers declines with tenure in the firm; and converges eventually to zero due to employers' learning about match-specific productivity. Besides these key variables of interest, the regressions, which we initially only estimate for the largest regional labour market in Denmark (the Copenhagen area), also control for a comprehensive set of worker characteristics and, importantly, firm fixed effects to account for permanent differences in firms that hire more or less intensively through referrals.

Figure 2 plots the parameter estimates of our referral proxy for each year of tenure, as well as their confidence intervals, both for log wages and the job turnover probability. Focusing on wages first, the depicted profile shows that having started a job in a firm with a former coworker present increases the initial wage by around 7 percent. In subsequent years, and as predicted by the theory, this advantage declines steadily until, after about four years of tenure, the wage profiles between referred workers and externally hired workers have fully converged. The corresponding findings for job turnover show the expected opposite pattern. In the first period, referred workers are about 5 percentage points less like-

ly to leave their firm compared to externally hired workers. Once more, this initial difference then converges to zero, suggesting that the learning process about the true match-specific productivity of workers is completed within about three to four years.

Further analysis documents these convergence patterns for various subgroups of the population, with particularly large initial gains from referrals for young workers and those with low education levels, suggesting that for these groups of workers, uncertainty may be particularly large; so they have the most to gain from the information-relaying property of referrals. In addition, there is some evidence that referrals by former coworkers with whom the prior personal interaction was more prolonged or recent are particularly efficient in reducing match-specific uncertainty. Overall, this study documents how the use of referrals within coworker-based networks positively affects the matching of workers to firms, giving rise to higher paid and longer lasting jobs and increasing overall efficiency in the labour market.

Conclusion

With the arrival of administrative data that provide comprehensive longitudinal information on individual work relationships, the systematic analysis of coworker-based networks and their role in the labour market has experienced an enormous boost in recent years. Moving ever closer to capturing actual personal interaction, several researchers have studied this important dimension of social networks within the context of established theoretical models, providing strong overall evidence that coworker-based networks play a positive role in labour market outcomes. My own findings show that this happens both through these networks' positive effect on the flow of information about job opportunities (Glitz 2013), and through their ability to reduce uncertainty about match-specific productivity (Glitz and Vejlin 2015).

References

Bayer, P., S. E. Ross and G. Topa (2008), "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes", *Journal of Political Economy* 116 (6), 1150–96.

Beaman, L. and J. Magruder (2012), "Who Gets the Job Referral? Evidence from a Social Networks Experiment", *American Economic Review* 102 (7), 3574–93.

Bentolila, S., C. Michelacci and J. Suarez (2010), "Social Contacts and Occupational Choice", *Economica* 77 (305), 20–45.

Calvó-Armengol, A. and M. O. Jackson (2004), "The Effects of Social Networks on Employment and Inequality", *American Economic Review* 94 (3), 426–54.

Calvó-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 (2010), "Friends Networks and Job Finding Rates", *IZA Working Paper* No. 5240.

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

Colussi, T. (2013), "Migrant Networks and Job Search Outcomes: Evidence from Displaced Workers", *mimeo.*

Damm, A. P. (2014), "Neighborhood Quality and Labor Market Outcomes: Evidence from Quasi-Random Neighborhood Assignment of Immigrants", *Journal of Urban Economics* 79, 139–66.

Dustmann, C., A. Glitz and U. Schönberg (2011), "Referral-based Job Search Networks", *CReAM Discussion Paper* No. 14/11.

Edin, P.-A., P. Fredriksson and O. Aslund (2003), "Ethnic Enclaves and the Economic Success of Immigrants - Evidence from a Natural Experiment", *Quarterly Journal of Economics* 118 (1), 329–57.

Glitz, A. (2013), "Coworker Networks in the Labour Market", *IZA Discussion Paper* No. 7392.

Glitz, A. and R. Vejlin (2015), "Information Flow in Networks of Former Coworkers", *mimeo.*

Granovetter, M. (1974), *Getting a Job: A Study of Contacts and Careers*, Harvard University Press, Cambridge, Massachusetts.

Heath, R. (2013), "Why Do Firms Hire Using Referrals? Evidence from Bangladeshi Garment Factories", *Manuscript*, University of Washington.

Hensvik, L. and O. N. Skans (2013), "Social Networks, Employee Selection and Labor Market Outcomes", *IFAU Working Paper* no. 15.

Ioannides, Y. M. and L. D. Loury (2004), "Job Information Networks, Neighborhood Effects, and Inequality", *Journal of Economic Literature* 42 (4), 1056–93.

Kramarz, F. and O. N. Skans (2015), "When Strong Ties Are Strong: Networks and Youth Labor Market Entry", *Review of Economic Studies,* in press.

Montgomery, J. D. (1991), "Social Networks and Labor-market Outcomes: Toward an Economic Analysis", *American Economic Review* 81 (5), 1407–18.

Munshi, K. (2003), "Networks in the Modern Economy: Mexican Migrants in the U.S. labor Market", *Quarterly Journal of Economics* 118 (2), 549–99.

Saygin, P. O., A. Weber and M. A. Weynandt (2014), "Coworkers, Networks, and Job Search Outcomes", *IZA Discussion Paper* No. 8174.

Schmutte, I. M. (2015), "Job Referral Networks and the Determination of Earnings in Local Labor Markets", *Journal of Labor Economics,* in press.

Simon, C. J. and J. T. Warner (1992), "Matchmaker, Matchmaker: The Effect of Old Boy Networks on Job Match Quality, Earnings, and Tenure", *Journal of Labor Economics* 10 (3), 306–30.

Topa, G. (2011), "Labor Markets and Referrals", in J. Benhabib, A. Bisin and M. O. Jackson, eds., *Handbook of Social Economics*, Chapter 22, North-Holland Press, San Diego, California and Amsterdam, 1193–221.