

Els Opuscles del CREI

num. **17**

January 2006

Social Networks and Labour Market Outcomes

Antoni Calvó-Armengol



CENTRE DE RECERCA
EN ECONOMIA INTERNACIONAL
GENERALITAT DE CATALUNYA
UNIVERSITAT POMPEU FABRA

The Centre de Recerca en Economia Internacional (CREI) is a research centre sponsored by the Universitat Pompeu Fabra (UPF) and the Generalitat de Catalunya. It is located at the campus of the UPF in Barcelona.

CREI's objective is to promote research in international economics and macroeconomics subject to the highest standards of academic scholarship.

Els Opuscles del CREI are the instruments for the diffusion of CREI's research beyond the academic world. Each *Opuscle*, aimed at a general public, gathers together conclusions and observations of specialized papers published, or about to be published, in professional journals. The opinions stated in *Els Opuscles del CREI* are the sole responsibility of the authors.

Published by: CREI
Universitat Pompeu Fabra.
Ramon Trias Fargas, 25-27 08005 Barcelona
Tel. 93 542 24 98
© CREI, 2006
© of this edition: Antoni Calvó-Armengol
ISSN: 1137 - 7844
Design: Fons Gràfic
Printed by: Masanas Gràfiques
Legal register: B-43612-2005

Social Networks and Labour Market Outcomes¹

Antoni Calvó-Armengol

1. Social networks and job search

Labour markets are, essentially, matching markets where both sides search for the best partner. Standard models of the labour market focus on individual characteristics, such as human capital levels, but usually disregard group characteristics, such as the pattern of social connections between individuals in the market. Yet, numerous empirical accounts show that the matching of individuals to jobs is widely influenced by networks of interpersonal ties. Roughly one half of all jobs are filled through contacts both in Europe and in the US (see Table 1). There is thus more to labour markets than a mere anonymous confrontation between demand and supply.

Although a half of all workers have heard about their current job through a friend or a relative, job seekers only devote around a sixth of their search effort to activating their social circles. Job search through friends and relatives is thus a very productive search channel, three times as productive as the average search mode. This results from the combination of two features. First, social contacts convey lots of job

Table 1**Employment search channels**

Search Channel (%)	US	France	Italy	Spain
Social Contacts	52.0	33.1	24.1	39.4
Direct Applications	–	25.9	33.2	35.0
Newspapers	9.4	9.8	3.9	5.3
Employment Agencies	5.8	9.8	7.4	7.4
Other	32.8	21.4	31.4	12.9

European data from ECHP1999 and US data from PSID1980. Note that search modes partially overlap (for instance, direct applications may follow recommendation by some acquaintance, etc.). Note, also, that these numbers vary slightly across occupations, skill levels and socio-economic backgrounds.

opportunities for a low search cost. Second, offers gathered through social contacts display a higher acceptance rate than alternative search modes.

The use of social networks is pervasive both for employer recruiting and for worker job-seeking. Social contacts help workers to find jobs, and employers to find employees. Social networks convey rich and trustworthy information, that they spread broadly and quickly across both sides of the labour market. As such, they constitute a cost-effective search channel both to broaden the number of information items available to firms and workers, and to deepen the quality of each item.

The exact role of social networks in the job search process, however, is not straightforward. The range of its effects varies with the type of information they convey. It also has different implications for the demand or the supply side of the labour market. Isolating these effects is important to enhance our understanding of the labour market functioning and improve its regulation.

Economists have long understood the qualitative role of social networks in the hiring market for firms, the demand-side of the labour market. Only recent findings have improved our understanding of the supply-side effects of social networks on workers. We first briefly describe the former. The latter then constitutes the core of this essay.

2. Firms

Recruiting is a critical activity for most firms. Given a position to be filled, recruiting amounts to searching for the best possible match, a new employee most compatible with the requirements of the job slot. This entails two types of costs for firms, search and screening costs. To recruit, firms first need to identify a pool of applicants or prospective employees, which results in search costs (advertisement, head-hunter services, etc.). Firms then need to screen out these applicants carefully according to their foreseen performance with regards to the characteristics of the position to be filled. This entails screening costs (applicants interviews and other screening procedures).

Economists and sociologists have long understood why and how network-mediated information in the hiring market alleviates both types of costs for firms.

Suppose first that a firm resorts to its incumbent employees and constitutes a pool of applicants out of their circle of acquaintances. By doing so, the firm saves on search costs, as it need not spend any extra resource to look for prospective employees.

Besides, and very importantly, it also saves on screening costs. Indeed, empirical accounts all show that referred applicants' likely performance

is, on average, higher than that of external applicants. Two reasons can be invoked. On one hand, because their reputation within the firm is at stake, incumbent employees are likely to refer only trusted acquaintances. Referrals thus tend to display a better fit for the job, including observable skills, but also traits that are harder to observe, such as loyalty or attachment to the job. On the other hand, the long-term nature of the relationship between incumbent employees and the firm provides the latter with superior information on the incumbents' traits. Based on the presumption that individuals tend to socialize with individuals like themselves,² the firm can then finely untangle different referred applicants depending on the characteristics of the refereeing incumbent, which are known to the firm.

Referrals also display another important advantage for the firm. Because of their social connections with the firms' current workforce, referrals are generally better informed about all aspects of the working environment provided by the firm, including, again, features that are difficult to observe, such as non-pecuniary characteristics attached to the job, or the firm corporate culture. They also reduce search costs for job seekers, who can get information about job opportunities from their friends and acquaintances passively, without actively devoting resources to explicit search activities.

Altogether, referrals reduce employer uncertainty about prospective worker's productivity. They also help applicants to adjust their expectations about the working conditions they will find in the firm. Both applicants and employers thus reap informational returns from referrals, and always for a low search cost. For both reasons, referrals generally yield better matches than alternative recruiting channels.

Consistent with these arguments, empirical accounts show that referrals display a lower turnover rate and stay longer in the firm. Some companies even give bonuses to their current employees to induce them to propose applicants. Recent estimates evaluate the savings on hiring, screening and training costs associated to hiring referrals to be equivalent to a 66.6% rate of return to referral bonuses.

The use of referral networks, however, has also a very important dark side. Because referral recruiting is circumscribed to the network boundaries of the current employees' friends, it excludes or, at least, penalizes, applicants from outside these boundaries. This, of course, entails equity costs for the workers and often, in addition, efficiency costs for the firm that may be disregarding superior candidates. For this reason, referral networks are not perfect substitutes for other formal channels. Rather, they simply provide a complementary search mode to more standard recruiting processes.

3. Workers

The supply-side effects of social networks in the labour market amount to reducing the informational asymmetries faced by the firms –and, to some extent, by the applicants– during the recruiting process.

The demand-side effects cover a broader set of issues, some of which are not straightforward. In what follows, we highlight the specific implications that the inner details of the network structure have on the anatomy and dynamics of employment, including turnover, wages, and reemployment probabilities, both at the individual and at the group level. We also relate labour

market withdrawals to network connections. By doing so, we argue that persistent inequality across groups with seemingly identical economic characteristics can be amenable to differences in their network composition and in the work history of the group.

We proceed in two steps. First, and to begin with, we analyze how a fixed network shapes employment outcomes. This corresponds to the sociologist viewpoint, where social resources impose constraints on individuals and, ultimately, influence their outcomes. We then turn to the case where individuals try to manipulate the network of connections to their advantage. This is the economist viewpoint, where individuals invest in social resources as a response to their needs.

The sociological approach relates network structure to labour outcomes, and helps understanding a number of seemingly disparate empirical regularities of the labour market. Building on this connection, the economic approach then relates network structure (and subsequent labour outcomes) to participation rates and human capital investment, and sheds light on inequality issues. We conclude by deriving the main public policy implications from the previous findings.

4. The sociology of labour market networks

Consider a network of social contacts among workers. Following the sociologist approach, we view social connections as part of the agents' endowment in the economy. The network of social contacts does not depend on agents' decisions it is a given.

The labour market is subject to some turnover, and currently employed workers sometimes lose their job as a result of exogenous shocks to the economy. When this happens, unemployed workers may try and search for a new job through formal search methods, such as employment agencies, newspapers advertisements, union hiring halls, etc. They may, as well, and as evidence corroborates, activate their social connections to try to gather as much information as possible about potential labour opportunities. Social connections are thus a valuable resource in the labour market which, in due course, can help to provide access to job information. Social contacts are complementary to own search effort, and provide a partial insurance against the labour market uncertainty associated to the economy shocks.

The standard economic view of the labour market focuses on standard search methods. The emphasis of search-theoretic models of the labour market is placed on the understanding of the processes of job creation and job destruction, and on the mechanisms of wage formation. The search and matching approach to labour markets constitutes a flexible and applicable paradigm, that sheds light on a number of issues such as unemployment rate, unemployment spells, labour market turnover, and wage distribution. Nonetheless, given the high reliance on social connections by job seekers in real markets, this approach can, at best, provide only a partial understanding of the labour market functioning. More importantly, labour market policies based solely on the individual search paradigm are bound to have a limited effect, insofar as they are not purposefully intended to impinge also on the main recruitment mode, social networks.

In what follows, we concentrate on the role of social networks in matching job seekers with vacancies. We thus envision the labour market in a way that complements the already available tools. In particular, although we allow for own search effort to get as broad as possible a picture, in contrast to the standard view, we leave the determinants of such individual search activities largely unspecified. Because we track throughout on how word-of-mouth communication helps workers to find jobs, and its consequences for the market as a whole, our analysis of the labour market can also be understood as an attempt to open the black-box of the standard matching paradigm.

More precisely, consider an economy populated with both employed and unemployed workers, with unemployed workers actively searching for a job. We allow for on-the-job search, and currently employed workers can seek to improve their current status. Unlike the standard search model, however, but consistent with empirical observations, we allow acquainted agents to communicate job opportunities to each other through word-of-mouth casual communication. We thus explicitly recognize the role of the social network as a means by which job opportunities are disseminated in the economy. We analyze its effects, and derive consequences for public policy.

In this economy, unemployed agents can rely on two very different channels of job information: own search effort in the labour market, and information exchange within their social circle. As a sociologist would do, we view information exchange between acquainted agents as part of the standard reciprocal behaviour to be expected in a network of ongoing relationships. We do not explicitly analyze information exchange through the prism of individual incentives.

Communication within each pair of agents can display a variety of real-life features. Of course, and most likely, we should only expect information flows to arise from employed to unemployed agents, or from high to low-wage workers, as otherwise the initial recipient could obtain a wage raise by concealing the information item instead of displaying it. Beyond that, the information may be relayed to distant recipients through some intermediate agents, possibly with some decay due to these intermediate relays. It can also be dispatched according to some a priori ranking on the choice of the recipient by the source, where the ranking reflects various intensities of social attachment. Else, this relational heterogeneity can be accompanied with some volatility in the pattern of socialization, as existing relationships may only be activated randomly, when agents meet each other.

In a nutshell, the bilateral communication pattern consistent with our analysis is very general, and encompasses a huge variety of situations.

Within a given economy, the collection of bilateral communication flows forms a network. As just stated, this network can have weighted, directed, or even stochastic links. More importantly, the precise manner in which jobs flow in the economy is tightly related to the network geometry and its inner intricacies. For instance, consider n agents connected through a network. Then, any information item can travel between any two agents in the population in at most two steps if the network is a star encompassing them all, while up to $n-1$ steps are needed if, instead, the network is a line. Note, incidentally, that both the star and the line have exactly the same total number of links ($n-1$), so that an outside observer cannot distinguish

between them by only focusing on network macroscopic characteristics; instead, it needs to know precisely who is linked with whom to discriminate between both networks. When interactions are not anonymous, networks are the best way to keep track of them all.

The importance of the network details in shaping information flows is true at the group level. It is also true at the individual level, where the ego-centred network that spans a given agents' direct and indirect contacts, and the structural differences across such networks, is an important source of heterogeneity across agents. So, for instance, more connections give potential access to more information. But this is not without qualification, as the structure of friendship ties among one's friends, and one's friends of friends, etc. also affects the intensity and the geometry of information flows.

Sociologists roughly distinguish between two types of network connections, *strong* and *weak* ties. On the one hand, strong ties –with close friends– constitute cohesive subgroups whose members are densely tied to each other and form network cliques. Strong links tend to be transitive: if *ian* and *john* know each other, and so do *john* and *kate*, then most likely *ian* knows *kate* when these ties are strong. On the other hand, weak ties –with acquaintances– constitute chains of connections and form far-reaching and open-ended networks that penetrate the social boundaries impermeable to stronger ties. Weak links tend to be intransitive: even when *ian* and *john*, and *john* and *kate* know each other, it is very unlikely that *ian* knows *kate* when these ties are weak.

Weak and strong ties play different roles in communication or information diffusion processes.

The performance of a single connection as an information bridge, both for the two agents located at its end-points, and for the economy as a whole, is largely determined by the tie strength of this connection. Introverted close relationships –strong ties– vehicle redundant information, while extroverted friendship links –weak ties– transmit more diversified information. For this reason, weak ties constitute richer information sources than strong ties. As Granovetter (1973) claims “whatever is to be diffused can reach a larger number of people, and traverse greater social distance (...), when passed through weak ties rather than strong ties” (p. 1366). In substance, weak ties foster widespread diffusion, while strong ties breed local sharing.

Altogether, this discussion suggests that the information flows associated with two chains of contacts of identical length, but different levels of introvertedness along this chain of contacts, are generally different. For this reason, the role of networks in shaping economic outcomes is often subtle, and we resort in the sequel to a number of examples to stress the different effects that may arise.

Note that the superiority of weak ties –with respect to strong ties– to disseminate job information in the market is not at odds with the (relative) superiority of referrals by close friends –with respect to anonymous applicants– to alleviate informational asymmetries in the firm. Indeed, search for information can have either an intensive or an extensive margin. Search at the intensive margin consists in getting additional information about an offer already received, while search at the extensive margin consists in broadening the access to more information. Referrals from close friends increase the intensive margin of information flows from the supply-side

(workers) to the demand-side (firms) of the labour market. Instead, weak ties increase the extensive margin of information flows within the supply-side of the market –between workers.

In what follows, we examine in turn the role of social networks on individual and group employment dynamics. We show how a number of apparently disparate empirical features of the labour market can be accounted for by the functioning of labour market networks. Altogether, we believe that the following analysis qualifies the social setting as a new and plausible driver for these empirical observations.

5. Employment patterns: close friends

To begin with, consider a very stylized economy populated with only two agents. Wages are a given, and agents are either employed at this fixed wage level, or unemployed. For ease of exposition, most of our discussion will borrow from stylized models of this sort, but the main intuitions that we present carry over to more general and realistic set-ups.

In the labour market, jobs are permanently created and destroyed. To fix ideas, we suppose that job breakdown occurs, on average, every five quarters –sixty-seven weeks in the calculations. To counter this outflow from employment, the economy also generates new job opportunities every once and a while. For the sake of the example, we assume that all agents, either employed or unemployed, hear about a new available job slot, on average, once every ten weeks.

This stylized economy shares some real-life features. In particular, the uncertainty built into

the labour market creates some volatility in employment outcomes. For this reason, the work histories of individual agents consist on a sequence of employment and unemployment periods. The frequency and spell length of these different periods depend on the relative values of the breakdown and job acquisition rates.

In general, the job acquisition rate depends on the rate of job creation, the cost of job search, and the economy-wide wage distribution. This, at least, is the perspective taken by standard search theory. Here, given that wages are fixed, unemployed workers always have strict incentives to accept any job offer that comes along. It never pays to turn down a given job offer, which would only delay re-employment. In our over-stylized economy, the job acquisition rate then coincides exactly with the frequency with which job information becomes available to an unemployed worker. This frequency itself increases with the number of valuable sources of information available to each unemployed worker.

Luckily enough, all the intuitions that we present in the sequel still hold true when we enrich our framework by allowing unemployed agents to delay strategically re-employment, exactly as in a standard search model. For ease of presentation, we stick to variations of this over-stylized economy throughout.

Suppose first that the two agents do not know each other. Then, an unemployed worker only recovers a job when he himself hears directly from the labour market of an available job offer, probably as a result of his own search effort. The labour market is the only information source for an isolated agent. Therefore, in our example, unemployed workers find a job, on average, after ten weeks being jobless. Given the value for the

breakdown rate, we can conclude that every worker is unemployed 13.2% of his time, on average. Abusing slightly, we say that the economy unemployment rate is 13.2%. (See Table 2a).

Suppose now that the two agents know each other, maybe because they live in the same neighbourhood, or because they attended the same school, or are members of the same clubs, etc. Then, altruistic concerns attached to this intimate relationship induce agents to share with each other valuable information about the job market they may possess, at least when this is harmless. In particular, when one worker is employed while the other is unemployed, the jobless agent can now rely on two different sources of job information. First, he may still obtain a job offer directly, from his own search effort. Second, his employed friend can also become aware of a job opportunity, maybe from conversations at his own workplace. Then, given their ongoing friendship tie, and insofar as the new job opportunity does not yield to any wage raise for the employed agent, the employed and informed agent relays the information about the available job opportunity to his unemployed acquaintance.

Table 2a
Unemployment rate per network geometry

Network	Unemployment
1 • • 2	13.2%
1 —• 2	8.3%

The upper panel corresponds to two isolated agents; the lower panel depicts a dyad. For the simulations, we assume that each worker receives one job offer every 67 weeks, and loses his current job every 10 weeks, on average.

Friendship ties thus potentially widen the sources of job information for unemployed agents. For this reason, the length of unemployment spells in an economy where the two agents know each other decreases, compared to the case where agents are isolated. Here, we find that acquainted agents are unemployed for 8.3% of their work history, so that friendship ties reduce the unemployment rate by more than one third.

In a volatile economy, friendship ties constitute a costless job information channel, complementary to the costly individual search effort in the labour market. The returns to friendship ties as information providers, however, are themselves highly volatile. Indeed, at every given point in time, the economy is populated with both employed and unemployed workers, and every social contact of any given agent can experience either labour status. In other words, friendship ties have their own work history. Given that only currently employed acquaintances are (potentially) valuable information sources, the informational value of friendship ties from a given agents' perspective changes through time with the collective work history of his set of friends. When most direct acquaintances are employed, a given unemployed agent has access to a rich variety of information sources, and quickly moves to employment. When most indirect acquaintances are themselves unemployed, however, less information sources are available, and job recovery takes longer.

The length and occurrence of one's unemployment spell is thus intimately related to the length and occurrence of the unemployment spells of one's friends, and work histories of acquainted agents go hand-in-hand. For this reason, we observe co-movements in individual outcomes within close knits of direct acquaintances. These

co-movements correspond to concomitant individual employment and unemployment spells for group members, so that the group as a whole stays, on average, either close to full employment, or close to full unemployment.

The concomitance in individual work histories also reduces the length of unemployment spells for each group member, so that full unemployment periods last shorter with social ties than without. The reason is as follows. A single individual transition from unemployment to employment within the group turns on a new information channel for this group. This, in turn, triggers more and quicker transitions to employment for group members which, again, ignites additional information channels for the rest of the group. Every individual job recovery thus creates a positive feed-back for the group members that are still unemployed. Subsequent individual transitions to employment then occur at a faster rate. For this reason, a networked group snaps back to full employment more quickly when pushed away from it, than a group without inner social contacts.

The co-movements in individual work histories for directly acquainted agents, and the corresponding relative stickiness of full employment for the group as a whole, explain the smaller likelihood of being unemployed with a friendship tie than without one, in our stylized economy with only two agents.

6. Employment patterns: the whole network

Co-movements in work histories between directly linked agents are not surprising. Within pairs, the interests of the two directly connected

parties are trivially aligned. Each worker prefers an employed friend, who may help him find a new job if needed, than a jobless tie in need. The information about an available job slot, however, is of a rival nature. The more jobless agents happen to know about it, the lower the likelihood that any given agent will end up filling this particular job slot. Two-link away indirect contacts in a network are thus potential competitors for the job information held by any common friend.

Given this rival nature of job information, and the corresponding competition exerted by distant connections, the work histories of distant (and competing) agents in a network may, a priori, follow very asynchronous trends. This intuition, however, is wrong. Co-movements in work histories always arise between any pair of agents in a network, even when these are only indirectly connected. In fact, in the long-run, two-links away, and even further away, indirect contacts help improving the employment status of common friends. This is beneficial, as employed common friends are potential information sources. As a matter of fact, this long-run beneficial effect of distant connections balances the short-run detrimental effect of competition for rival information. The intensity of these co-movements for distantly connected agents, however, quite naturally dampens with the number of intermediate agents between them.

To illustrate this fact, consider an economy populated with four agents. Suppose first that agents are connected two by two, so that we have two separate dyads. As in the case of a single dyad, the overall unemployment rate is equal to 8.3%.

Suppose now that each agent is involved in two different dyadic relationships, so that we

have four overlapping dyads that form a wheel. The unemployment rate then falls to 6.3%, which represents a reduction by one quarter with respect to the dyad economy. The reason for this decrease is twofold. First, each agent has now at his disposal two direct links, instead of just one, which broadens the sources of job information. Second, each agent has now one indirect acquaintance. Together with the two direct links, this gives as much as three partners with whom labour outcomes co-evolve. Broader channels and strengthened outcomes are responsible for the unemployment decrease. (See Table 2b).

In this example, we can measure exactly the strength of the work histories' linkage both for direct and indirect contacts. We simply compute the (conditional) individual probability to be unemployed when one's (direct or indirect) link is also unemployed, and compare it to the (unconditional) unemployment probability, which coincides with the unemployment rate. We find that the individual unemployment probability when one's dyad mate is also unemployed is equal to 6.5%, which represents a 3% increase with respect to the unconditional individual unemployment probability, equal to 6.3%. Now, we look at agents

Table 2b
Unemployment rate per network geometry.

Network	Unemployment
	6.3%
	6.0%

The upper panel corresponds to two four-agents wheels; the lower panel depicts an eight-agents wheels.

diametrically opposed in the wheel which are two-links away from each other. Despite being apparent competitors for any job information held by their common friends, their work histories go, in fact, hand-in-hand. For two-link away agents, the conditional unemployment probability is equal to 6.4%, so that co-movements in work histories for indirect connections are responsible for a 2% increase in unemployment when the prospects of the indirect contacts are bad.

Such spatial correlation of employment is directly observable in census track data for Chicago, where the impact on individual employment probability of a one standard deviation increase on average unemployment in one's neighbouring census tracts is 1.5 times higher than the impact of a one standard deviation increase on educational level. Also, evidence from plants lay-offs in Italy indicates that increasing the network size from the first to the second quartile reduces unemployment duration by as much as -7%, while a 10% increase of the share of employed contacts reduces unemployment duration by -14%, roughly.

7. Network morphology and labour outcomes

The fact that work histories go hand-in-hand for every given pair of agents in a network, however distantly connected, has very important implications. In particular, knowing exactly who passes information to whom is critical in determining a person's long-term employment prospects. Unraveling all the possible chains of information transmission in a network becomes a pre-requisite to formulate accurate predictions for the economy as a whole.

In what follows, we show by means of examples how increasingly detailed features of the network morphology play a role in shaping economic outcomes. Lacking fine enough information on the network embeddedness of economic actors, the economist is deprived of a key explanatory variable to account properly for empirical observations on the labour market.

Consider for instance an economy with eight agents. This economy is populated with dyads, that is, all the agents in this economy are connected two by two, and no worker is wandering around in isolation.

Suppose, to begin with, that every single agent is engaged in only one dyadic relationship, so that the economy consists of four different dyads. The society is fragmented into four groups containing two directly connected agents each. This is just a replica of the single dyad economy, and we can predict an unemployment rate equal to 8.3%.

Suppose now that each agent is involved in two different dyadic relationships. A number of geometric configurations for the network of social contacts is compatible with this fact. Furthermore, in the absence of additional information to single out one specific network, any such configuration can arise a priori. For the sake of illustration, we consider only two possible network geometries, and obtain the corresponding values for the unemployment rate.

In the first network geometry, the overlap across dyads yields to two disconnected wheels with four different agents each. The society is thus fragmented into two connected and unrelated components. Within each component –the four-agents wheel– each agent has exactly

two direct dyad partners, and one indirect two-links away connection. This is a replica of an economy already analyzed above. The unemployment rate is equal to 6.3%, significantly lower than the four disconnected dyads case.

Alternatively, the overlap across dyads can yield a single connected wheel involving the eight agents. (See Table 2b). This is still a network where, as in the two-wheels network, each agent has exactly two direct contacts. Contrarily to the two-wheels configuration, however, the society is now more connected, and information can travel from any agent to every other agent through chains of intermediate dyad partners. The set of indirect contacts is broader than in the two-wheels fragmented network, where each agent has only one two-links away indirect contact. Here, agents have five indirect contacts each, consisting of two two-links away, two three-links away, and one four-links away connections. Given the positive co-movements in labour outcomes for every pair of agents in a network, however distantly connected, and the higher relative number of indirectly paired agents in this society, the unemployment rate is lower than that of the two-wheels economy. We get a 6% unemployment rate for this single-wheel economy.

The previous examples show that, other things being equal, the unemployment rate in an economy decreases when both the number of direct and indirect contacts among agents in the economy increases. More precisely, doubling the number of direct dyad partners reduces unemployment by around –24%, from 8.3% to 6.3%. Increasing the number of indirect contacts from one to four then further decreases the unemployment rate by almost –5%, from 6.3% to 6%.

These examples suggest that reducing network connections to a simple dyadic relationship, or even to a collection of disconnected dyads, seldom provides a rich enough set up to capture adequately the whole range of network effects that may arise. To map the social context into economic outcomes accurately, one needs to introduce indirect connections and to keep track explicitly of their effects on the aggregate behaviour. In a networked economy, two is never a good representation for many.

Beyond the number of direct and indirect contacts, the distribution of indirect contacts in terms of path length also affects the overall behaviour of the economy. In particular, and as illustrated below, aggregate unemployment increases with the average path length in a network.

We still deal with an economy with eight agents. Now, each agent is involved in three different dyadic relationships. We consider two polar network configurations consistent with this fact. We start with the eight-agents single wheel economy, where everybody takes part in two dyadic relationships. To increase the total number of dyad partners for each agent by one, we add four additional links to this network. We distinguish two different ways to add those links. (See Table 2c).

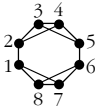
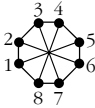
A first possibility is to add four diameters to the wheel. The new links connect agents who are diametrically opposed in the original wheel. We get a highly distributed network, where agents' locations are interchangeable. In this network, each agent has three direct contacts and four two-links away indirect contacts. Indirect contacts are thus very close, and can all be reached through only one intermediate dyad partner. As a result, in the wheel with diameters, the span of social

contacts is highly spread and the average path length is low, equal to 1.6 links. The unemployment rate is equal to 4.8%.

A second possibility is to add the four new links so that to create two local four-agents clusters in the wheel. In this case, agents have more distant indirect contacts. More precisely, agents in the centre of the cluster have two two-links away and two three-links away indirect contacts, while agents in the periphery of the cluster have three two-links away and one three-links away indirect contacts.

Compared with the wheel with diameters, this clustered network is more closely knitted. The cluster forms a relatively introverted pattern of relationships –strong ties–, where agents' friends tend to be friends with each other. The breadth of social contacts is narrower than in the wheel with diameters, specially for agents at the centre of the cluster, which are located further away from information sources outside their social group. The resulting average path length is longer, equal to 1.8 links, which represents a 13% increase with

Table 2c
Unemployment rate per network geometry

Network	Unemployment
	4.9%
	4.8%

The upper panel corresponds to an eight-agents wheel with four-agents clusters 2, 3, 4, 5 and 1, 8, 7, 6; the lower panel depicts an eight-agents wheel with diameters 1-5, 2-6, 3-7, and 4-8.

respect to the wheel with diameters. Because intragroup information sources tend to convey redundant information, the unemployment rate in this clustered network is also higher, and equal to 4.9%. The corresponding 2% increase in unemployment with respect to the wheel with diameters is solely amenable to a relative longer length in indirect contacts, as their number is the same across both types of networks.

The previous example establishes a positive relationship between average path length an unemployment rate, thus wedging a bridge between macro structural characteristics of the network, and macro economic variables of the labour market. On the one hand, the average path length decreases when the span of network contacts spreads, that is, when network relationships get less introverted. On the other hand, the average unemployment decreases when the breadth of available social ties gets wider, and the sources of job information more diversified. This relationship, however, is only a side effect of the micro features of information diffusion in the network, and can not be understood without exploring the micro structural properties of the network pattern of dyadic relationships.

In substance, the micro details of the network structure matter in shaping outcomes. The different values for the unemployment rates computed above can not be accounted for adequately with only some macroscopic knowledge of the pattern of socialization, such as the total number of links. Both the double four-wheels economy and the single eight-wheel economy have eight links each. Yet, their unemployment rates differ by as much as 5%. Even finer information on this pattern of contacts, such as knowing the exact number of direct and indirect contacts each agent has, is not enough. A

path length increase of 13% accounts for as much as a 2% difference in unemployment when we compare the wheels with diameters to the clustered wheel, where agents have exactly the same number of both direct and indirect contacts.

Detailed structural differences between networks cannot be ignored to explain the economy's behaviour. In a volatile environment, social contacts diversify the sources of job information. With more social contacts, job inflows increase relative to job outflows, and unemployment decreases. The higher the number of information sources, the bigger the relative inflow increase, and the steeper the unemployment decrease. This explains the role played by the number of direct contacts. However, any structural difference across two networked economies, even with identical number of direct contacts, and however tiny, may mould information flows that can differ in length, intensity, and seeded pattern, thus yielding to different unemployment rates. Beyond standard explanatory variables, such as labour market regulation, skill composition mix in the economy, capital and technological endowment, etc. social networks and structural differences in socialization patterns across social groups, regions or countries, are also responsible for differences in observed labour market outcomes.

8. Duration dependence

So far, we have mainly explored the cross-sectional implications of the co-movements in work histories that arise naturally across networked agents who share information in the labour market. In particular, our previous discussion highlights the role of the structural pattern of socialization in shaping outcomes, and

explains the likely presence of employment and unemployment clusters of relatively close-knitted agents at different spots in the economy. We now explore the time series effects of labour market networks.

We have already mentioned the emergence of a relative persistence in aggregate employment dynamics, that takes the form of a combined stickiness and attraction effect. The economy tends to oscillate between (nearly) full employment and (approximately) full unemployment: the closer the economy gets to one extreme, the greater the pull is from this extreme. Besides, different parts of the network may display asynchronous behaviour, with some parts experiencing booms and others busts at the same time.

We show that labour market networks lead to a serial correlation at the level of individual work histories. This serial correlation corresponds to negative duration dependence of exit rates from unemployment. In other words, the individual probability to move from unemployment to employment is not stationary in a networked economy. Rather, it displays a clear time trend, namely, the longer the unemployment spell for a given agent, the smaller the probability for this agent to recover a job.

Consider first an economy with no network ties. Agents wander in isolation. Job outflows are regulated by the exogenous breakdown shocks to the economy, while job inflows depend on individual's own search effort. Consider an unemployed agent at a given point in time. Given our assumptions on own search effort, the probability of becoming employed for this unemployed agent is stationary. It does not depend on the particular date, and is

independent of the observed work history of the worker. Exit rates from unemployment do not display any duration dependence.

Consider now an economy with networked agents. We explore the dark side of co-movements in work histories, and identify a source of duration dependence.

When a given agent is unemployed, and while he remains so, he does not constitute a valuable source of information for his acquaintances. No acquaintance can expect to gather any job opportunity from him for the time being, as the agent himself is in need for this information. When every such acquaintance becomes unemployed, therefore, transitions from unemployment to employment take longer than if the agent was employed –less information sources are available. Now, the longer the agent stays unemployed, the more it is likely that any of his acquaintances faces some adverse contingency during this (long) unemployment spell and loses his current job.

When this happens, the newly unemployed acquaintance finds himself in a pretty bad shape, as he can not rely on the (relative) long-term unemployed agent to recover a job quickly. The newly unemployed acquaintance, therefore, is likely to face a long unemployment spell himself. But this, in turn, is bad news for our long-term unemployed agent, who has just lost a valuable job information channel (a direct acquaintance has just turned unemployed) and, besides, who cannot expect that this valuable source will be restored soon (the newly unemployed acquaintance is likely to stay unemployed for a long time). Altogether, the unemployed agents' own job prospects are harmed.

In a networked economy, therefore, the exit rate from unemployment is not stationary. In essence, this stems from the fact that co-movements in work histories consist of reciprocal influences. First, because of such co-movements, bad outcomes for a given agent impoverish the agent's own social environment. Second, by the very same co-movements, this poorer environment harms the agent's prospects in return. Through this reciprocal influence that operates back and forth, one's own past work history ends up influencing one's own future employment prospects. The expected probability of obtaining a job decreases in the length of the agent's unemployment spell. We get duration dependence.

That exit rates from unemployment tend to be lower when the duration of the spell goes up is also a well-documented fact, even though the accuracy of the estimates is still subject to some debate. Economic explanations for duration dependence, however, are a bit scarcer. They either appeal to some sort of exogenous capital accumulation and depletion mechanism, or to some stigma effect.

The first type of explanation considers that workers accumulate some sort of human capital during employment, and that this stock is exogenously depleted during unemployment, which complicates re-employment. The second type of explanation posits that long-term unemployed are stigmatized, and systematically relegated to the bottom ladder of potential employees' ranking by hiring firms.³

Econometricians also invoke another plausible explanation, unobserved heterogeneity. In other words, agents may possess an idiosyncratic characteristic that affects their attractiveness as employees, which is unobservable to the

econometrician but observed by the recruiting firm. The pool of long-term unemployed would then be overwhelmingly composed of unattractive workers, and the more so, the longer the unemployment spell. Because of this progressive degradation of the pool composition, long-term unemployed are less and less likely to be re-employed, and the hazard rate exhibits duration dependence.

Our explanation for duration dependence is orthogonal to these proposals. With labour market networks, agents who have been unemployed longer have increasingly more ties with unemployed agents. If we interpret employed acquaintances as some sort of relational capital that is valuable for job search, we can view long spells of unemployment as progressively depleting the stock in relational capital. As agents get more and more isolated in the workplace, their stock of capital decreases further. Contrary to previous explanations, the stock of relational capital in a networked economy evolves endogenously with the labour market turnover, and synchronously across directly and indirectly connected agents.

The network explanation for duration dependence can also be disentangled from the common unobserved heterogeneity argument the following way. The network channel states that duration dependence can be accounted for if we can control for the state of the whole network, whereas unobserved heterogeneity does not allow for this. Rich enough data may allow to implement a direct test. Else, provided that networks correlate with, say, geographic location, the network channel could be identified insofar as location interacts with other workers' variables to account for employment outcomes, something that is not predicted by other models.

9. Dropping out

So far, we have established that individuals confronted with different social settings, that is, endowed with different social assets, tend to experience different economic outcomes, insofar as those social assets provide access to economic resources –here, relationships provide information about jobs. More precisely, we have worked out how differentials in labour outcomes can be amenable to the pattern of socialization among economic agents and the details of its inner structure. The particular location a given agent holds in a network of social contacts accounts for part of his economic achievements.

To a large extent, the nature and pattern of social contacts depend on the agent's behaviour and decisions. The choice of a residential location, or of some leisure activity, for instance, constitutes also a choice of a social environment. More generally, anyone has some discretion on the choice of his own social circle, and on the strength of each social relationship. Given the economic value of social ties identified above, we can thus expect agents to try to manipulate the network of social contacts within which they are embedded at their advantage.

This is the economic viewpoint. Agents invest in their social network capital while pursuing rationally their own self-interest. Here, the decision to establish social connections is conceived similarly to, for instance, the investment decisions on human capital. Of course, actually, network formation involves large and complex social processes, and the incentives approach to network formation only constitutes a crude –yet, very informative– approximation to such real-life mechanisms.

In what follows, we analyze the economic consequences of the spontaneous formation of social networks as a response to their individual members' incentives. This analysis provides a basis for understanding the observed higher drop-out rates in some social groups versus others –such as blacks versus whites in the US, or female versus male in the most OECD countries–, and the sustained wage gap and inequality of employment rates across such groups.

The white-black wage gap is of the order of 25-40% in the US, while the gender wage gap is of the order of 10-35% in OECD countries. This wage gap is partly explained by differences in factors such as the quality of education and skill levels, but nonetheless, there is still a significant residual that is unexplained. Moreover, once one accounts for the fact that blacks (resp. women) drop out of the labour force at a higher rate than whites (resp. men), the true gap in wages actually increases. Indeed, drop-out rates are 2.5 to 3 times higher for blacks compared to whites in the US, and of the same order for women compared to men in many countries. Even if one believes any inequality in wages between social groups to be entirely explainable by differences in factors such as education, skills, and drop-out rates, one is then left to explain why those should differ.

An analysis of social networks provides a basis for observing both higher drop-out rates in one social group versus another and sustained inequality in wages and employment rates even among those remaining in the labour force. The key is to let the network emerge as a response to individual incentives.

There are a number of ways to let the agents decide at their own convenience on the network geometry of the set of bilateral links that

surrounds them. The most natural model, at least at first sight, allows agents to consciously pick each and all of their dyad partners, selecting them one by one. This approach, however, poses a number of major technical hurdles that complicate the analysis heavily and obscure the interpretation of the results. A close approximation to real-life processes seldom comes without a cost in terms of model tractability which, ultimately, may ruin a proper understanding of the mechanisms at play in the model. This truism is specially salient in networked economies, where the object of study—the pattern of inter-twinned interactions—is particularly complex to handle.

We take a very simple alternative route that happens to provide rich and insightful predictions. Starting from a given network, we simply let each agent decide whether he keeps the current location he holds in this network, or he withdraws from it. The collection of these binary decisions to drop in or out results in a network of agents participating in the labour force, and a group of agents who have dropped out from the labour market. For those who have dropped in, the structure of the network connecting them shapes their work histories.

The binary decision to drop in or out results from an individual cost-benefit analysis.

First, staying in the network is costly. Think, for instance, of an alumni network. The cost then summarizes all the efforts exerted in college. For the case of a club's network, instead, the cost could reflect the opportunity cost of time to attend the club's activities. The examples abound. We assume that agents are identical and interchangeable ex ante, that is, they all face the same chance to be endowed with any possible cost value. Ex post, however, each agent is

assigned a particular realization of the cost value, which reflects some population heterogeneity in, for example, social or cognitive skills.

Second, staying in the network is beneficial. Here, the network benefits take the form of expected labour earnings. Given the network embeddedness of individual work histories, the benefits accruing to each agent from his particular location in the network vary across such locations with their structural characteristics. The social structure, however, does not affect the value of individual socialization costs. Therefore, gross labour earnings carry alone the whole burden of network effects.

We know that every single agent work history co-varies with the work histories of any other direct or indirect mates in the network, however distantly connected. Therefore, from a given agent's perspective, having more agents participate in the network is better news, as this increases his number of direct or indirect contacts which implies better employment prospects for the future, that is, longer and more frequent expected employment spells. Every single agent's benefit of entering the labour market network thus depends positively on the others' decisions to do so.

This positive interdependence in network labour earnings generates some bandwagon or contagion effect in the decisions to drop in or out. When one agent drops out, the prospects worsen for all those who are still in. If any agent in the market was already on the verge of dropping out, he may now make up his mind and effectively drop out. This, in turn, increases the chances that some of the agents still in the market also drop out, and so forth. The contagion effect also works in reverse, and drop-ins call for

additional drop-ins which then propagate through the network as prospects improve for those in.

Consider, for instance, our simple dyad economy. In the absence of a network link, the individual employment probability is 86.8%. With a network link, instead, the individual employment probability of each dyad partner increases to 91.7%. Suppose that the cost to wage ratio takes any possible value between 0 and 100 for all agents, and that each value is equally likely. Then, in the economy populated with single agents, only those for which the ex post cost to wage ratio is higher than 86.8 drop out from the market. For all the others, the benefit to stay in compensates for its cost. So, before individual costs are assigned, we can expect a drop out rate of 13.2%.⁴

Now, let's analyze the dyad economy. Suppose, to begin with, that each agent takes the other drop in decision as a given. Then, mimicking our previous reasoning, we can conclude that only agents for which the ex post cost to wage ratio is higher than 91.7 drop out from the market, which yields a drop out rate equal to 6.3%. Suppose, now, that one of the agent's cost value is between 86.8 and 91.7 while the other agent's cost is above 91.7. We know that, in this case, the second agent decides to drop out. The first agent, then, is alone in the market. As a result of his partner dropping out, his employment prospects worsen, and the agent's employment probability decreases from 91.7% to 86.8%. We are back to the single agent case. Given that his cost to wage ratio is above 86.8, the first agent also decides to drop out. Note, however, that his cost to wage ratio is below 91.7. Therefore, if his partner stayed in the market, he would rationally decide to stay in. His drop-out


decision is only the consequence of his partner's drop-out decision, and we have contagion. When contagion is taken into account, the drop-out rate increases from 6.3% to 6.6%⁵. Contagion, here, is responsible for a 4.8% increase in drop-outs.

More generally, the individual drop-out decisions and the population pattern of contagion effects for these decisions depend on the inner geometric details of the network connecting agents, and on the employment outcome of each agent in the population at the time the drop-out decision is taken. (See Table 3).

10. Inequality persistence

The network contagion effect has important implications for understanding the persistent inequality in employment outcomes across social groups along many divide lines, such as race, gender, or country of origin.

Table 3
Drop-out rate in a bridge network



Agent	1	2, 3, 4, 5	6	7, 8, 9, 10
Drop-out Rate	47%	42%	91%	93%

Initially, agents 1, 2, 3, 4 and 5 are all employed, while agents 6, 7, 8, 9 and 10 are all unemployed. For the simulations, the ratio of the entry cost to the wage lies in [80, 100] for each value being equally likely. A drop-out rate for agent 1 equal to 47% means that agent 1 drops out in 47% of all the possible population cost/wage combinations in this range.

Consider two interchangeable groups with identical inner network structure, cost values and employment outcomes at a given date. Both groups, therefore, display the same economic and sociological fundamentals.

Now, modify one group slightly, by cutting out some connections, or by increasing slightly its number of unemployed, or both. The other benchmark group remains unaltered. In the modified group, because either the network or labour market outcomes –or both– are poorer, employment prospects worsen for everyone. Then, a handful of agents may decide to drop out –whereas they were all eager to stay in the market in the initial configuration of the group. But, if some agents drop out, this modifies the group structure further, worsening prospects further, and more drop-outs follow, and so on. Until the cascade of drop-outs stops, maybe because all have dropped out, maybe because those who stay in have low enough costs such that it is still worth for them to stay despite the bad prospects they face.

The group we obtain at the end of this contagion process displays a much lower labour force participation and a much poorer inner network structure –for those who are in the labour market– than that of the unmodified benchmark group. To begin with, initial conditions were only slightly different across the two groups. Eventually, though, differences in drop-out rates are large. Employment outcomes, which depend tightly on the prevailing network structure, will now differ persistently across both groups.

Differences in social endowments, however small, are heavily amplified through contagion and lead to drastic differences in connection

patterns which, in turn, breed persistent disparities in economic outcomes. Social endowments constrain choices in a non-trivial and dramatic fashion. In particular, collective employment history of a group and its social structure arise as two possible sources for persistent inequality in the labour market. The use of social networks in the labour market generates equity costs borne by those excluded from the wider community, irrespective of the exclusion mechanism at play –for instance, segregation or assortative mating. This has obvious consequences for the economics of immigration. It also identifies the social setting as a plausible source for the relatively viscous social mobility patterns measured in most countries, specially for the two tails of the earnings distribution.

The contagion mechanism described above is also reminiscent of the conformity effects often advocated by sociologists to account for the (sometimes) striking observed differences in, for example, high school drop-outs across various social groups with otherwise very similar characteristics. Here, preferences for conformity are not imposed as an exogenous characteristic of the individuals who, for some reason, would comply with a gregarious behavior. Rather, the observed cascade in decisions arises endogenously across ego-centered individuals from their use of social connections in the labour market.

11. Lessons for public policy

Our recent contributions clarify the economic consequences of information gathering through social networks by job-seekers, largely documented by sociologists. We show how the

use of social contacts has a big overall impact on labour market outcomes (wages, hazard rates, drop-outs, inequality), and argue that the economic analysis of the labour market cannot ignore the social embeddedness of economic actors without sacrificing explanatory power.

Ex ante identical individuals connected through a network can end up with very different outcomes. This ex post heterogeneity is mainly due to differences in geometry across local ego-centred networks. Indeed, the network-driven externality between connected individuals varies across them with the local characteristics of the network geometry. In a sense, background social variables act as a long-term credit constraint with economic significance. The economic roles of social networks identifies a new type of market failure that opens the scope for public intervention.

The sort of public policy that is called for in this case needs to be tailored to the explicit role of the network. In particular, the planner can alter the allocation of network externalities across individuals, and implement optimal redistributive schemes, by suitably manipulating the network locally.

Interventions to improve and sustain the employment status of a given agent also improve the outlook for the social acquaintances of the agents targeted by the intervention –a contagion effect in reverse. In a networked society, education subsidies and other labour market regulation policies display local increasing returns for the social topology. Targeted interventions magnify the initial effect of the policy, and the more so, the more the subsidies and programs are circumscribed to tightly clustered agents in a network. For this reason, targeting is more

efficient than spreading resources more broadly, and the more concentrated the interventions, the higher the efficiency gains.

Network effects call for highly localized interventions. In this context, optimal policies need to respond to two different concerns. On the one hand, optimal policies require the identification of candidate optimal group targets –this is the pure normative concern. On the other hand, optimal policies require to operate a selection of the actually targeted group among those candidates –this the political economy concern, a NIMBY⁶ problem in reverse.

Notes

(1) This Opuscle borrows extensively from our joint work with Matt Jackson, to whom I am most grateful; all the views presented here are inexorably tied to this joint work. I also thank Antonio Cabrales for his very useful comments, and Yves Zenou for our joint network projects.

(2) Sociologists refer to the mechanism according to which “birds of a feather flock together” as assortative mating or inbreeding social bias.

(3) Note that the stigma effect could be reinterpreted as some exogenous and progressive loss of some sort of reputational capital in the labour market, that affects hiring.

(4) Only one agent drops out with probability 0.132×0.868 while two of them drop out with probability 0.132×0.132 . So, the drop-out rate is $2 \times (0.132 \times 0.868 + 0.132 \times 0.132) / 2 = 0.132$, that is, 13.2%.

(5) The contagious drop-out decision arises when one value of the cost to wage ratio is higher than 91.7 while the other is between 86.8 and 91.7, an event with probability $0.063 \times (0.917 - 0.868) = 0.003$. So, the drop-out rate accounting for contagion is 6.6%.

(6) An acronym for “Not In My BackYard”.

References

- Calvó-Armengol, Antoni (2004) “Job Contact Networks” *Journal of Economic Theory*, 115(1), pp. 191-206.
- Calvó-Armengol, Antoni and Matthew O. Jackson (2004) “The Effects of Social Networks on Employment and Inequality” *American Economic Review*, 94(3), pp. 426-454.
- Calvó-Armengol, Antoni and Matthew O. Jackson (2005a) “Networks in Labor Markets: Wage and Employment Dynamics and Inequality” *Journal of Economic Theory* (in press)
- Calvó-Armengol, Antoni and Matthew O. Jackson (2005b) “Like Father, Like Son: Social Networks, Human Capital Investment and Social Mobility” CalTech and ICREA
- Calvó-Armengol, Antoni and Yves Zenou (2005) “Job Matching, Social Networks and Word-of-Mouth Communication” *Journal of Urban Economics*, 57, pp. 500-522
- Calvó-Armengol, Antoni and Yannis Ioannides (2005) “Social Networks in Labor Markets,” in *The New Palgrave*, L. Blume and S. Durlauf (eds), London: MacMillan Press (in press)
- Granovetter, Mark (1973) “The Strength of Weak Ties” *American Journal of Sociology*, 78(6), pp. 1360-1380.
- Ioannides, Yannis M. and Linda Datcher Loury (2004) “Job Information Networks, Neighborhood Effects, and Inequality” *Journal of Economic Literature*, 92, pp. 1056-1093.

Opuscles already published

- 1. Reconsidering Spanish Unemployment**
Ramon Marimon (June 97)
 - 2. Reducing Unemployment. At Any Cost?**
Fabrizio Zilibotti (December 97)
 - 3. Capital and Labor Taxes, Macroeconomic Activity, and Redistribution**
Albert Marcet (November 98)
 - 4. The Lender of Last Resort in Today's Financial Environment**
Xavier Freixas (November 99)
 - 5. Why does the Public Sector Grow? The Role of Economic Development, Trade and Democracy**
Carles Boix (November 99)
 - 6. Gerontocracy and Social Security**
Xavier Sala-i-Martin (July 2000)
 - 7. The Political Viability of Labour Market Reform**
Gilles Saint-Paul (December 2000)
 - 8. Are EU Policies Fostering Growth and Reducing Regional Inequalities?**
Fabio Canova (May 2001)
 - 9. Agglomeration Effects in Europe and the USA**
Antonio Ciccone (September 2001)
 - 10. Economic Polarization in the Mediterranean Basin**
Joan Esteban (May 2002)
 - 11. How do Households Invest their Wealth?**
Miquel Faig (October 2002)
 - 12. Macroeconomic and Distributional Effects of Social Security**
Luisa Fuster (April 2003)
 - 13. Educating Intuition: A Challenge for the 21st Century**
Robin M. Hogarth (September 2003)
 - 14. Capital Controls in Post-War Europe**
Hans-Joachim Voth (April 2004)
 - 15. Taxation of Financial Intermediaries**
Ramon Caminal (September 2004)
 - 16. Ready to Take Risks? Experimental Evidence on Risk Aversion and Attraction**
Antoni Bosch-Domènech / Joaquim Silvestre i Benach (November 2005)
 - 17. Social Networks and Labour Market Outcomes**
Antoni Calvó-Armengol (January 2006)
-



Antoni Calvó-Armengol

Antoni Calvó-Armengol earned a joint PhD in economics from École Nationale des Ponts et Chaussées, Paris and Universitat Pompeu Fabra, Barcelona in 2000. He had previously graduated from École Polytechnique in Paris, from École Nationale des Ponts et Chaussées in Paris, and from Escuela Superior de Caminos, Canales y Puertos at Universidad Politécnica de Madrid. He is currently ICREA Research Professor at Universitat Autònoma de Barcelona, CEPR Research Affiliate, and Associate Editor for the *Journal of Economic Theory* and the *Journal of the European Economic Association*. He has previously taught at Université de Toulouse, Universidad Carlos III de Madrid and Universitat Pompeu Fabra. His research on the economics of social networks and applied microeconomics has been published, among others, in *American Economic Review*, *Journal of Political Economy*, *Journal of Economic Theory*, *International Economic Review*, and *Journal of Urban Economics*. He has received the Young Economist Award of the European Economic Association.

CENTRE DE RECERCA EN ECONOMIA INTERNACIONAL

GENERALITAT DE CATALUNYA

UNIVERSITAT POMPEU FABRA

Ramon Trias Fargas, 25-27 - 08005 Barcelona

Tel: 93 542 24 98 - Fax: 93 542 18 60

E-mail: crei@upf.edu

<http://www.creiweb.org>

€ 6.00