# The Effects of Social Networks on Employment and Inequality

By Antoni Calvó-Armengol and Matthew O. Jackson
*American Economic Review* 94.3 (2004)

James J. Heckman

THE UNIVERSITY OF CHICAGO

Econ 350, Winter 2021

- The importance of social networks in labor markets is pervasive and well documented.
- Mark Granovetter (1973, 1995) found in a survey of residents of a Massachusetts town that over 50 percent of jobs were obtained through social contacts.
- Earlier work by Albert Rees (1966) found numbers of over 60 percent in a similar study.
- Exploration in a large number of studies documents similar figures for a variety of occupations, skill levels, and socioeconomic backgrounds

- In this paper, we take the role of social networks as a manner of obtaining information about job opportunities as a given and explore its implications for the dynamics of employment.
- In particular, we examine a simple model of the transmission of job information through a network of social contacts.
- Each agent is connected to others through a network. Information about jobs arrives randomly to agents.
- Agents who are unemployed and directly hear of a job use the information to obtain a job.
- Agents who are already employed, depending on whether the job is more attractive than their current job, might keep the job or else might pass along information to one (or more) of their direct connections in the network.

- Before proceeding to the model, let us also mention a fourth feature that is also exhibited by the model. Unemployment exhibits duration dependence and persistence.

- That is, when conditioning on a history of unemployment, the expected probability of obtaining a job decreases in the length of time that an agent has been unemployed.

- Such duration dependence is well documented in the empirical literature, for example, in studies by Stuart O Schweitzer and Ralph E. Smith (1974), Heckman and George Borjas (1980), Christopher Flinn and Heckman (1982), and Lisa M. Lynch (1989).

- To get a feeling for the magnitude, Lynch (1989) finds average probabilities of finding employment on the order of 0.30 after one week of unemployment, 0.08 after eight weeks of unemployment, and 0.02 after a year of unemployment.

# I. A Simple Network Model

- The model we consider here is one where all jobs are identical.
- We refer the reader to a companion paper, Calvo´-Armengol and Jackson (2003), for a more general model that nests this model, and also looks at wage dynamics, and allows for heterogeneity in jobs, decisions as to whether to switch jobs, repeated and selective passing of information, competing offers for employment, and other extensions of the model presented here.
- In short, the results presented here extend to wage inequality as well, and are quite robust to the formulation.
- There are $n$ agents.
- Time evolves in discrete periods indexed by $t$.
- The vector $s_t$ describes the employment status of the agents at time $t$.
- If agent $i$ is employed at the end of period $t$, then $s_{it} = 1$ and if $i$ is unemployed then $s_{it} = 0$.

- Any two people either know each other or do not, and in this model information only flows between agents who know each other.

- A graph g summarizes the links of all agents, where $g_{ij} = 1$ indicates that $i$ and $j$ know each other, and $g_{ij} = 0$ indicates that they do not know each other.

- It is assumed that $g_{ij} = g_{ji}$, meaning that the acquaintance relationship is a reciprocal one.

- If an agent hears about a job and is already employed, then this agent randomly picks an unemployed acquaintance to pass the job information to.

- If all of an agent's acquaintance are already employed, then the job information is simply lost.

- The probability of the joint event that agent $i$ learns about a job and this job ends up in agent $j$'s hands, is described by $p_{ij}(s)$, where

$$p_{ij}(\mathbf{s}) = \begin{cases} a & \text{if } s_i = 0 \text{ and } i = j, \\ \dfrac{a}{\sum\limits_{k: s_k = 0} g_{ik}} & \text{if } s_i = 1,\ s_j = 0, \text{ and } g_{ij} = 1; \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

and where the vector $s$ describes the employment status of all the agents at the beginning of the period.

# II. The Dynamics and Patterns of Employment

- The relationship between the one-period ahead employment status of an agent and his pattern of connections, as described by the $p_{ij}(s)$'s above, is clear.

- Having links to employed agents improves $i$'s prospects for hearing about a job if $i$ is unemployed.

- In addition, decreasing the competition for information from two-link-away connections is helpful.

- That is, if friends of my friends are employed rather than unemployed, then I have a higher chance of being the one that my friends will pass information to.

- Further indirect relationships (more than two-links away) do not enter the calculation for the one-period-ahead employment status of an agent.

# Figure 1: Negative Correlation in Conditional Employment

*Example 1 (Negative Conditional Correlation):*

- Consider Figure 1, a network with three agents, and suppose the employment from the end of the last period is $s_{t-1} = (0,1,0)$.
- In the picture, a darkened node represents an employed agent (agent 2), while unemployed agents (1 and 3) are represented by empty nodes.
- A line between two nodes indicates that those two agents are linked.
- Conditional on this state $s_{t-1}$, the employment states ${s_1}_t$ and ${s_3}_t$ are negatively correlated.
- This is due to the fact that agents 1 and 3 are "competitors" for any job news that is first heard by agent 2.
- Despite this negative (conditional) correlation in the shorter run, agent 1 can benefit from 3's presence in the longer run. Indeed, agent 3's presence helps improve agent 2's employment status.

*Example 2 (Correlation and Network Structure):*

- Consider an example with $n = 4$ agents and let $a$=0.100 and $b$=0.015.

- If we think about these numbers from the perspective of a time period being a week, then an agent loses a job roughly on average once in every 67 weeks, and hears (directly) about a job on average once in every ten weeks.

- Figure 2 shows unemployment probabilities and correlations between agents' employment statuses under the long-run steady state distribution.8

- If there is no network relationship at all, then we see an average unemployment rate of 13.2 percent.

- Even moving to just a single link ($q_{12} = g_{21} = 1$) substantially decreases the probability (for the linked agents) of being unemployed, as it drops by more than a third, to 8.3 percent.

- The resulting unemployment rate aggregated over the four agents is 10.75 percent.

# Figure 2: Correlation and Network Structure I

| $g$ | $\text{Prob}(s_1 = 0)$ | $\text{Corr}(s_1, s_2)$ | $\text{Corr}(s_1, s_3)$ |
|---|---|---|---|
| 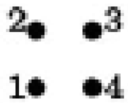 | 0.132 | – | – |
| 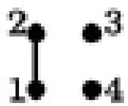 | 0.083 | 0.041 | – |
| 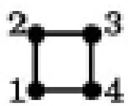 | 0.063 | 0.025 | 0.019 |
| 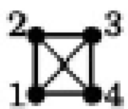 | 0.050 | 0.025 | 0.025 |

- As we see from Figure 2, adding more links further decreases the unemployment rate, but with a decreasing marginal impact.
- This makes sense, as the value to having an additional link comes only in providing job information when all of the existing avenues of information fail to provide any.
- The probability of this is decreasing in the number of connections.
- The correlation between two agents' employment is (weakly) decreasing in the number of links that each an agent has, and the correlation between agents' employment is higher for direct compared to indirect connections.
- The decrease as a function of the number of links is due to the decreased importance of any single link if an agent has many links.
- The difference between direct and indirect connections in terms of correlation is due to the fact that direct connections provide information, while indirect connections only help by indirect provision of information that keeps friends, friends of friends, etc., employed.

- Next, Figure 3 examines some eight-person networks, with the same information arrival and job breakup rates, $a$ =0.100 and $b$ =0.015.10

- Here, again, the probability of unemployment falls with the number of links, and the correlation between two employed agents decreases with the distance of the shortest path of links (geodesic) between them.

- Also, we can see some comparisons to the four-person networks: an agent has a lower unemployment rate in a complete four-person network than in an eight-person circle.

- In this example, the direct connection is worth more than a number of indirect ones.

- More generally, the trade-off between direct connections and indirect ones will depend on the network architecture and the arrival and breakup rates.

# Figure 3: Correlation and Network Structure II

| $g$ | $\text{Prob}(s_1 = 0)$ | $\text{Corr}(s_1, s_2)$ | $\text{Corr}(s_1, s_3)$ | $\text{Corr}(s_1, s_4)$ | $\text{Corr}(s_1, s_5)$ |
|---|---|---|---|---|---|
| 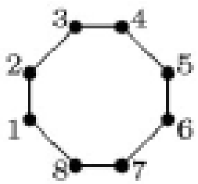 | 0.060 | 0.023 | 0.003 | 0.001 | – |
| 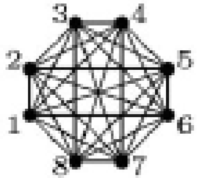 | 0.030 | 0.014 | 0.014 | 0.014 | 0.014 |

*Example 3 (Bridges and Asymmetries):*

- Consider the network in Figure 4. Again we calculate employment from simulations using the same arrival and breakup rates as in the previous examples.
- In this network the steady-state unemployment probabilities are 4.7 percent for agents 1 and 6, 4.8 percent for agents, 2, 5, 7, and 10, and 5.0 percent for the rest.
- While these are fairly close, simple differences of an agent's position in the network affects his or her unemployment rate, even though all agents all have the same number of connections.
- Here agents 1 and 6 have lower unemployment rates than the others, and 3, 4, 8, and 9 are the worst off.
- If we compare agent 3 to agent 1, we note the following: the average geodesic (minimum path) distance between any two agents who are directly connected to agent 3 is 4/3.

## **Figure 4:** A Network with a Bridge

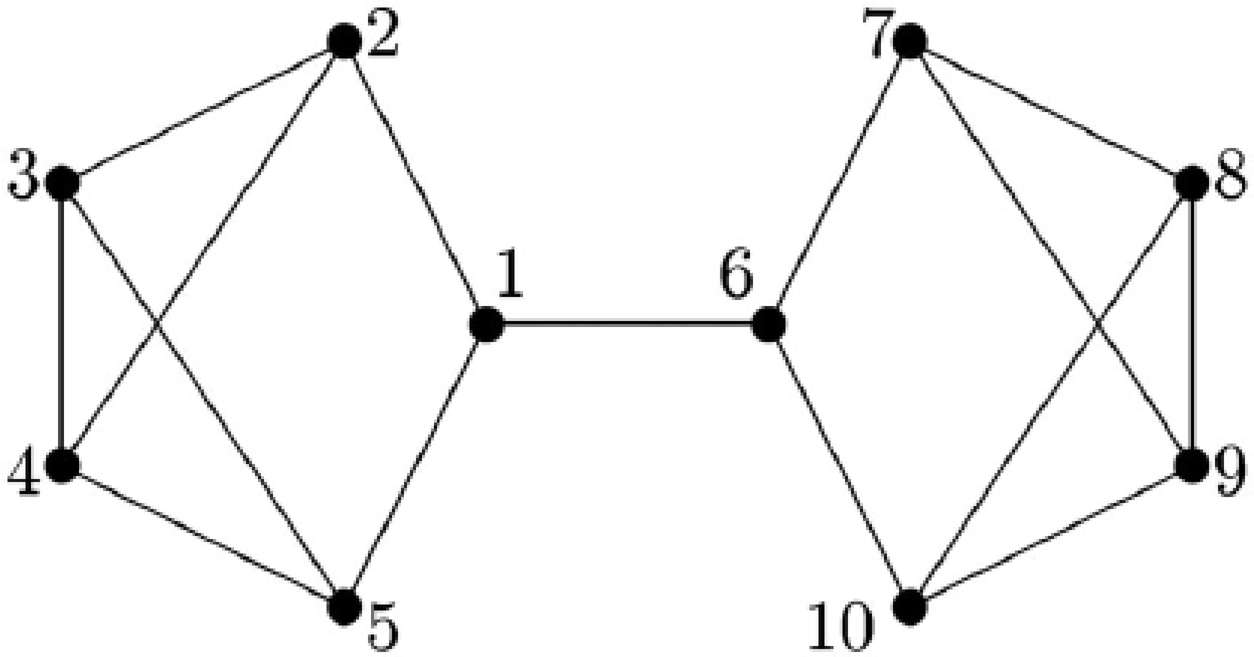

*Example 4 (Structure Matters: Densely Versus Closely Knit Networks):*

- The model can also show how other details of the network structure matter. Compare the long-run average unemployment rates on two eight-person networks with 12 links each. In both networks, all agents have exactly three links.

- But, the average length of the paths connecting agents is different across networks. Again, we run simulations with $a$ =0.100 and $b$ =0.015; see Figure 5.

- The average path length is lower for the circle with diameters than for the circle with local four-agents clusters, meaning that the latter is more closely knit than the former.

- Indeed, the average path length decreases when the span of network contacts spreads; that is, when the relationships get less introverted or less closely knit.

- The average unemployment increases with closed-knittedness, reflecting the fact that the wider the breadth of current social ties, the more diversified are the sources of information.

# **Figure 5:** Path Length Matters

| $g$ | average path length | average unemployment |
|---|---|---|
| 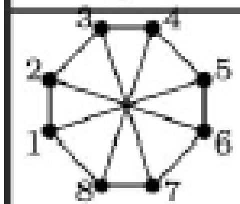 | 1.571 | 0.048 |
| 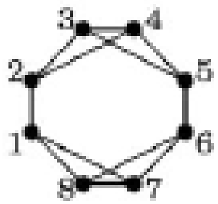 | 1.786 | 0.049 |

**PROPOSITION 1:** *Under fine enough subdivisions of periods, the unique steady-state longrun distribution on employment is such that the employment statuses of any path-connected agents are positively correlated.*

- The proposition shows that despite the shortrun conditional negative correlation between the employment of competitors for jobs and information, in the longer run any interconnected agents' employment is positively correlated.
- This implies that there is a clustering of agents by employment status, and employed workers tend to be connected with employed workers, and vice versa.
- This is consistent with the sort of clustering observed by Topa (2001).
- The intuition is clear: conditional on knowing that some set of agents are employed, it is more likely that their neighbors will end up receiving information about jobs, and so on.

**PROPOSITION 2:** *Under fine enough subdivisions of periods, starting under the steady-state distribution, the employment statuses of any two path-connected agents are positively correlated across arbitrary periods*

# III. Duration Dependence and Persistence in Unemployment

- As mentioned in the introduction, there are some other patterns of unemployment that have been observed in the data and are exhibited by a networked model.

- To see this, let us examine some of the serial patterns of employment that emerge.

- Again, consider job arrival and breakup rates of $a$ =0.100 and $b$ =0.015.15

- Ask the following question: suppose that a person has been unemployed for at least each of the last X periods.

- What is the probability that he or she will be employed at the end of this period?

- We examine the answer to this question in Figure 6 as we vary the number of periods of observed past unemployment and the network.

## **Figure 6:** Duration Dependence

| $g$ | 1 period | 2 periods | 10 periods | limit |
|---|---|---|---|---|
| 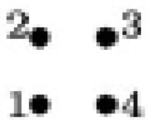 | 0.099 | 0.099 | 0.099 | 0.099 |
| 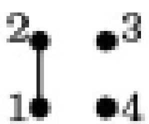 | 0.176 | 0.175 | 0.170 | 0.099 |
| 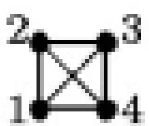 | 0.305 | 0.300 | 0.278 | 0.099 |

**PROPOSITION 3:** *Under fine enough subdivisions of periods and starting under the steadystate distribution, the conditional probability that an individual will become employed in a given period is decreasing with the length of their observed (individual) unemployment spell.*

- Indeed, longer past unemployment histories lead to worse inferences regarding the state of one's connections and the overall state of the network.
- This leads to worse inferences regarding the probability that an agent will hear indirect news about a job.
- That is, the longer $i$ has been unemployed, the higher the expectation that $i$'s connections and path connections are themselves also unemployed.
- This makes it more likely that $i$'s connections will take any information they hear of directly, and less likely that they will pass it on to $i$.
- In other words, a longer individual unemployment spell makes it more likely that the state of one's social environment is poor, which in turn leads to low forecasts of future employment prospects.

# *A. Comments on Stickiness in the Dynamics of Employment*

- The duration dependence for individuals is reflective of a more general persistence in employment dynamics.

- This persistence can be understood by first noting a simple feature of our model.

- When aggregate employment is relatively high, unemployed agents have relatively more of their connections employed and face relatively less competition for job information, and are more likely to hear about jobs.

- Conversely, when aggregate employment is relatively low, unemployed agents are relatively less likely to hear about jobs.

- To illustrate this point, consider the bridge network in Figure 4 that connects ten agents.

- We calculate the (average) individual probability that an unemployed agent finds a job within the current period, conditional on the total number of employed agents in the network.

- When there is no network connecting agents, the probability that an unemployed agent finds a job is simply the arrival rate a.

- In contrast, when agents are connected through a network (here, the bridge network of Figure 4), the probability of finding a job varies with the employment state.

- This conditional probability is a when everybody is unemployed, but then increases with the number of employed agents in the network, as shown in Table 1.

## Table 1: Probability of Finding Employment for Agents in the Bridge Network

| Number of employed | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
|---|---|---|---|---|---|---|---|---|---|---|
| $a = 0.100$; $b = 0.015$ | 10.0 | 10.4 | 12.0 | 14.5 | 17.9 | 20.7 | 25.4 | 25.7 | 28.7 | 34.4 |
| $a = 0.050$; $b = 0.050$ | 5.0 | 5.9 | 6.2 | 6.9 | 8.6 | 9.3 | 11.3 | 12.2 | 15.0 | 18.5 |

- This state dependence of the probability of hearing about a job, then implies a persistence in aggregate employment dynamics.

- As a network gets closer to full employment, unemployed agents become even more likely to become employed.

- Symmetrically, the lower the employment rate, the lower the probability that a given unemployed agent hears about a job.

- Although the process oscillates between full employment and unemployment, it exhibits a stickiness and attraction so that the closer it gets to one extreme (high employment or high unemployment) the greater the pull is from that extreme.

- This leads to a sort of boom and bust effect, as illustrated in Figure 7.

# **Figure 7:** Time Series of Employment for Networked Versus Disconnected Agents

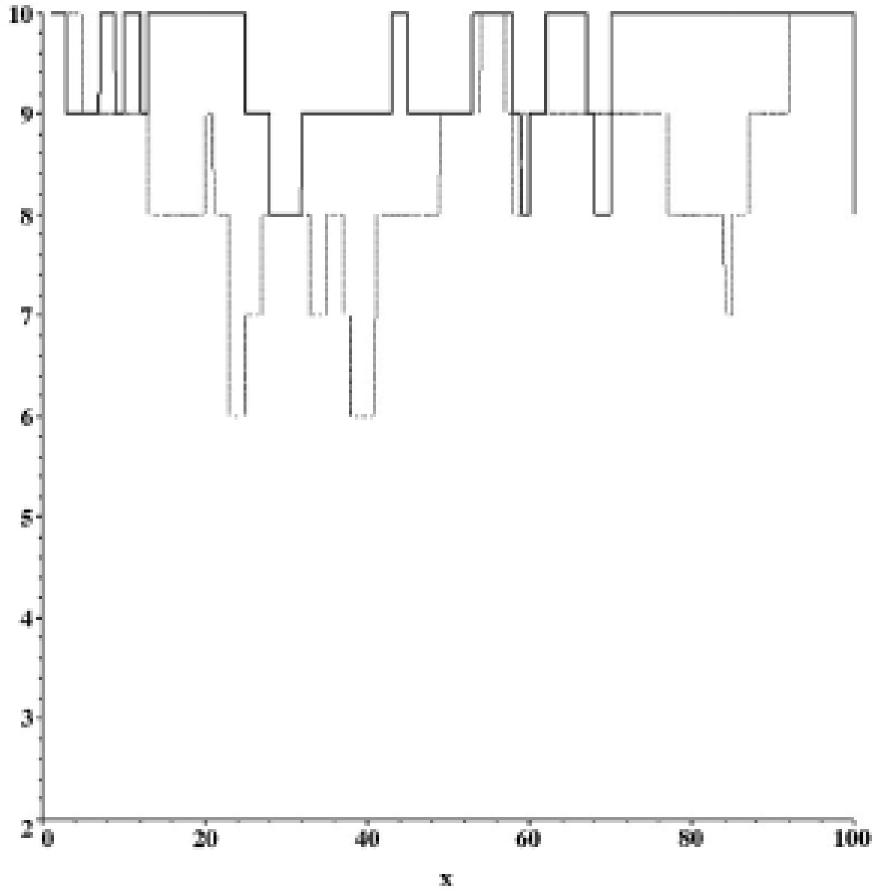


Aggregate employment over time for $a$=0.100 and $b$=0.015

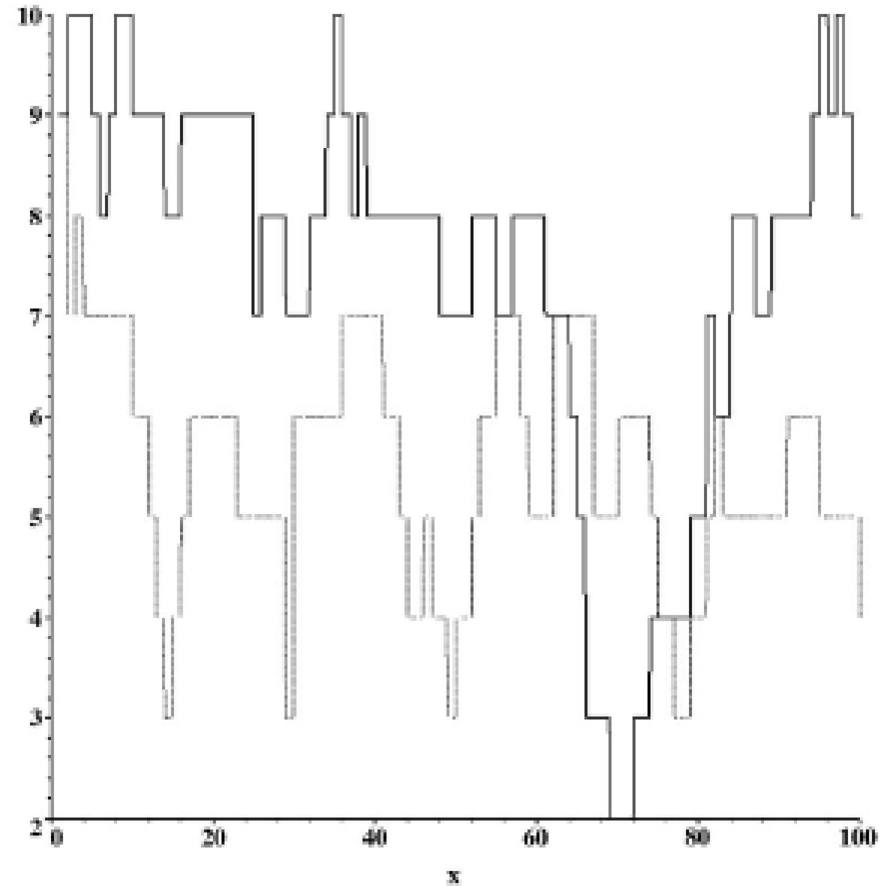


Aggregate employment over time for $a$=0.050 and $b$=0.050

**Figure 8:** Asynchronous Patterns of Employment Across Network Sections

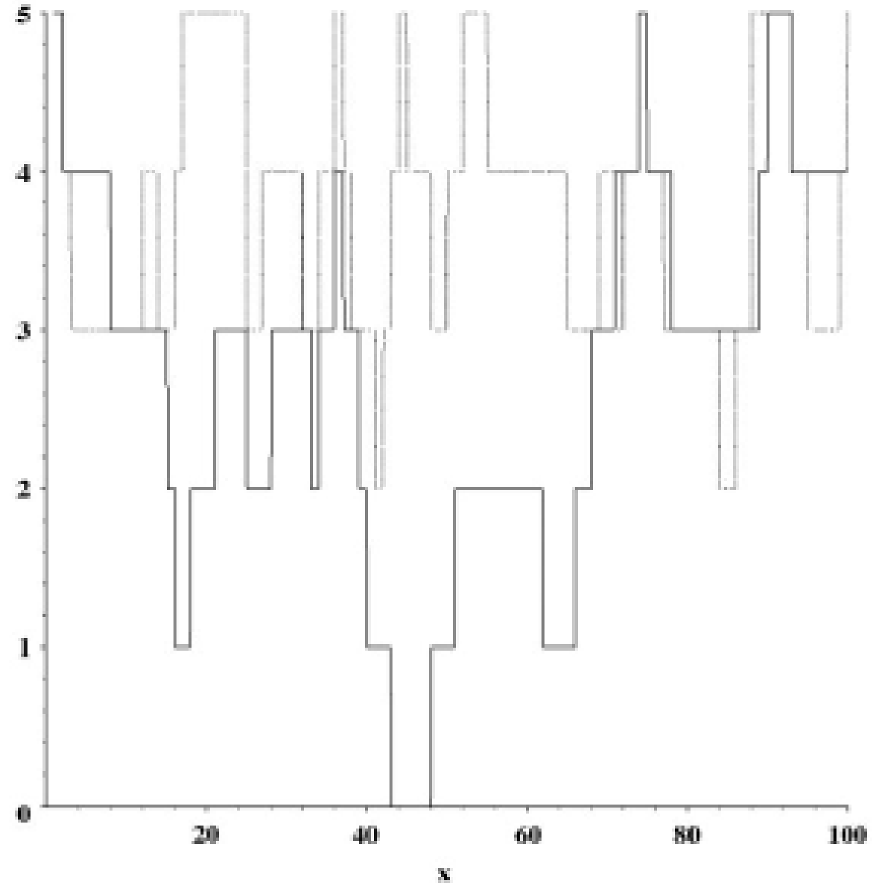


Aggregate employment of agents 1 to 5 versus agents 6 to 10 over time for $a$=0.050 and $b$=0.050

- We also point out that employment need not be evenly spread on the network, especially in a network such as the bridge network from Figure 4.
- As a result temporal patterns may be asynchronous across different parts of the network, with some parts experiencing booms and other parts experiencing recessions at the same time.
- This asynchronous behavior is illustrated in Figure 8, which plots separately over 100 periods the aggregate employment of agents 1 to 5 (the dotted line) and that of agents 6 to 10 (the plain line) in the bridge network from Figure 4 from a simulation with a=0.050 and b=0.050.

# IV. Dropping Out and Inequality in Employment

- We now turn to showing how the network model has important implications for inequality across agents, and how that inequality can persist.

- Our results so far show that an agent's employment status will depend in important ways on the status of those agents who are path connected with the agent.

- This leads to some heterogeneity across agents, as their networks and the local conditions in their networks will vary.

- Note, however, that in the absence of some structural heterogeneity across agents, their long-run prospects will look similar.

- That is, if the horizon is long enough, then the importance of the starting state will disappear.

*Example 5 (Initial Conditions, Dropouts, and Contagion):*

- To measure the contagion effect, we first ask how many people would drop out without any equilibrium effect, that is, if they each did the calculation supposing that everyone else was going to stay in.
- Then we can calculate how many people drop out in equilibrium, and any extra people dropping out are due to somebody else dropping out, which is what we attribute to the contagion effect.
- For these calculations, we take the cost of staying in the network, ci , to be uniformly distributed on [0.8, 1] and fix the per period wage to be w  1.
- We do the calculations with complete networks, where each participating agent is directly linked to every other agent.

- For Tables 2 and 3, the calculations are don for a discount rate of 0.9, where we simplify things by assuming that an agent starts in the initial state, and then jumps to the steady state in the next "period."

- This just gives us a rough calculation, but enough to see the effects.

- So, a agent who stays in gets a net payoff of $0.9_{p_i} - c_i$, where $p_i$ is the agent's steady-stat employment probability in the maximal equilibrium.

- We again set a=0.100 and b=0.015.21.

- So, for instance, in Table 3, when $n$ =16 and everybody is initially unemployed, we have 68 percent of the people dropping out on average.

- This means that we expect about 11 people to drop out on average and about 5 people to stay in.

## **Table 2:** Dropouts and Contagion—Starting Employed

| $s_0 = (1, \dots, 1)$ | $n = 1$ | $n = 2$ | $n = 4$ | $n = 8$ | $n = 16$ | $n = 32$ | $n \to \infty$ |
|---|---|---|---|---|---|---|---|
| Drop-out percentage | 58.3 | 44.5 | 26.2 | 14.7 | 9.7 | 7.8 | 6.8 |
| Percentage due to contagion | 0 | 8.8 | 5.0 | 1.4 | 0.4 | 0.2 | 0 |

## **Table 3:** Dropouts and Contagion—Starting Unemployed

| $s_0 = (0, \dots, 0)$ | $n = 1$ | $n = 2$ | $n = 4$ | $n = 8$ | $n = 16$ | $n = 32$ | $n \to \infty$ |
|---|---|---|---|---|---|---|---|
| Drop-out percentage | 100 | 97.8 | 92.9 | 82.2 | 68.0 | 60.6 | 56.8 |
| Percentage due to contagion | 0 | 12.1 | 21.7 | 18.9 | 8.7 | 3.0 | 0 |

**PROPOSITION 4:** *Consider two social groups with identical network structures. If the starting state person-by-person is higher for one group than the other, then the set of agents who drop out of the first group in the maximal equilibrium is a subset of their counterparts in the second group. These differences in drop-out rates generate persistent inequality in probabilities of employment in the steady-state distributions, with the first group having weakly better employment probabilities than their counterparts. There is a strict difference in employment probabilities for all agents in any component of the network for which the equilibrium drop-out decisions differ across the two groups.*

*Example 6 (Connected Social Groups and Dropouts):*

- Consider the network structure from Example 3; see Figure 9.
- Agents 1 to 5 start employed and agents 6 to 10 start unemployed.
- We do drop-out calculations as in Example 5.
- We take the ci to be uniformly distributed on [0.8, 1], fix $w = 1$, use a discount rate of 0.9, and have agents who stay in get a net payoff of $0.1_{s_i} + 0.9_{p_i} - c_i$, where $p_i$ is the agent's steady-state employment probability in the maximal equilibrium of the drop-out game, and $s_i$ is their starting employment state.
- The drop-out probabilities for the different agents are illustrated in Table 4.

## **Table 4:** Drop-outs Rates in the Bridge Network with Asymmetric Starting States

| Agent | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| Drop-out rate | 0.47 | 0.42 | 0.42 | 0.42 | 0.42 | 0.91 | 0.93 | 0.93 | 0.93 | 0.93 |

# V. A Look at Policy Implications

- Another lesson is that there is a positive externality between the status of connected individuals.
- So, for instance, if we consider improving the status of some number of individuals who are scattered around the network, or some group that are more tightly clustered, there will be two sorts of advantages to concentrating the improvements in tighter clusters.
- The first is that this will improve the transition probabilities of those directly involved, but the second is that this will improve the transition probabilities of those connected with these individuals.
- Moreover, concentrated improvements lead to a greater improvement of the status of connections than dispersed improvements.
- This will then propagate through the network.

## *A. Concentration of Subsidies*

- The experiment we perform here is the following.
- In each case we subsidize two agents to stay in the market—simply by lowering their cost ci to 0.28.
- The question is which two agents we subsidize.
- In the network, each agent has four connections.
- The network structure is as follows. Each agent has three links—two immediate neighbors and one that is slightly further away. This is pictured in Figure 10.
- Table 5 provides the percentage of agents who stay in the network as a function of who is subsidized (two agents in each case) and what the range of costs (randomly drawn) are.
- There are some interesting things to note.

## **Figure 10:** The Starting Network Structure

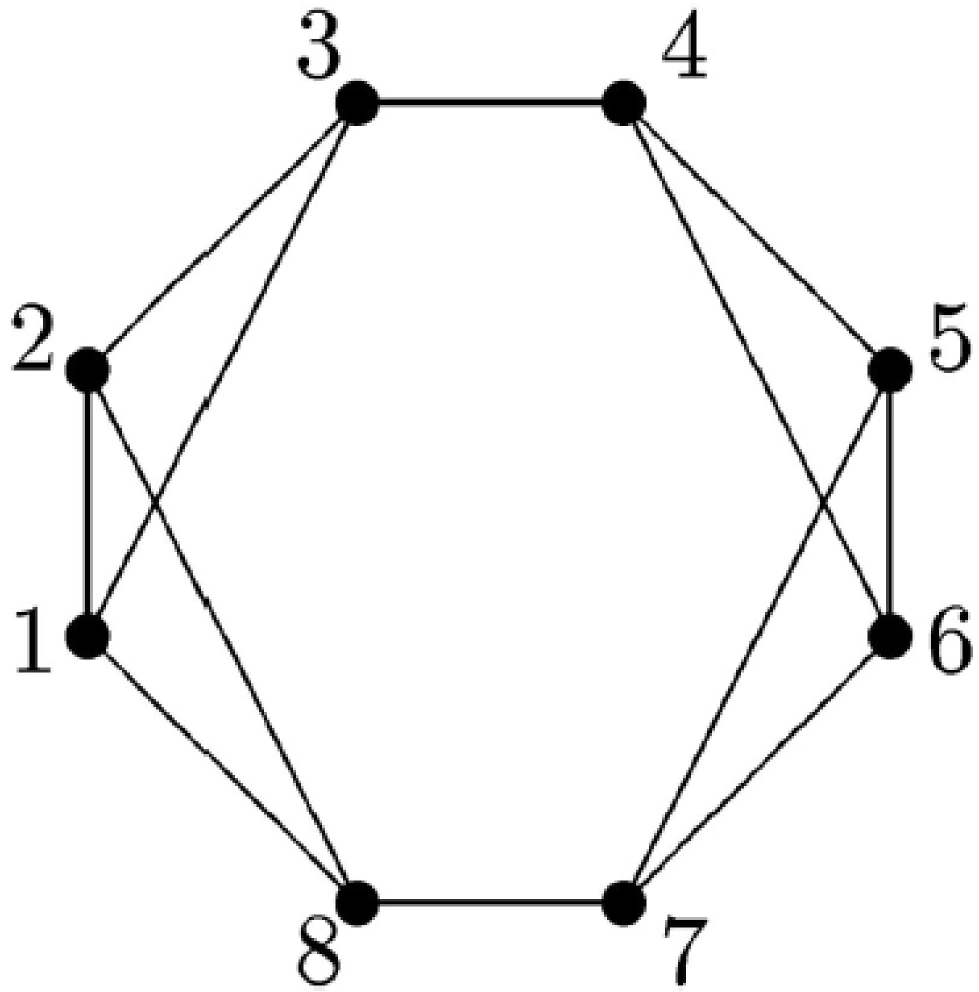

## **Table 4:** Subsidization Structure and Percentage of Agents Who Stay In

| | Cost range | | | |
|---|---|---|---|---|
| Agents subsidized | 0.80 to 1 | 0.82 to 1 | 0.84 to 1 | 0.86 to 1 |
| 1 and 2 | 52.9 | 39.4 | 27.8 | 25.0 |
| 1 and 3 | 53.6 | 39.4 | 27.1 | 25.0 |
| 1 and 4 | 57.2 | 43.4 | 27.9 | 25.0 |
| 1 and 5 | **57.9** | 43.8 | 27.0 | 25.0 |
| 1 and 6 | 57.9 | **44.0** | 27.0 | 25.0 |
| 1 and 7 | 57.1 | 43.4 | 27.8 | 25.0 |
| 1 and 8 | 53.5 | 39.4 | 27.1 | 25.0 |
| 3 and 4 | 54.5 | 39.3 | 26.1 | 25.0 |
| 3 and 7 | 57.7 | 43.6 | 27.4 | 25.0 |
| 3 and 8 | 56.2 | 42.9 | **29.1** | 25.0 |

## **Table 5:** Subsidization Structure and Percentage of Agents Who Stay In

| | Cost range | | | |
|---|---|---|---|---|
| Agents subsidized | 0.80 to 1 | 0.82 to 1 | 0.84 to 1 | 0.86 to 1 |
| 1 and 2 | 52.9 | 39.4 | 27.8 | 25.0 |
| 1 and 3 | 53.6 | 39.4 | 27.1 | 25.0 |
| 1 and 4 | 57.2 | 43.4 | 27.9 | 25.0 |
| 1 and 5 | **57.9** | 43.8 | 27.0 | 25.0 |
| 1 and 6 | 57.9 | **44.0** | 27.0 | 25.0 |
| 1 and 7 | 57.1 | 43.4 | 27.8 | 25.0 |
| 1 and 8 | 53.5 | 39.4 | 27.1 | 25.0 |
| 3 and 4 | 54.5 | 39.3 | 26.1 | 25.0 |
| 3 and 7 | 57.7 | 43.6 | 27.4 | 25.0 |
| 3 and 8 | 56.2 | 42.9 | **29.1** | 25.0 |

# VI. Possible Empirical Tests

- While as we have discussed, the model generates patterns of employment and dropouts that are consistent with the stylized facts from a number of studies, one might want to look at some additional and more direct tests of the model's predictions.

- Note that drop-out rates and contagion effects depend both on the costs ranges and on the values for the arrival rate and breakup rate.

- Some comparative statics are quite obvious: (1) as the expected cost increases (relative to wages), the drop-out rate increases; (2) as the breakup rate increases, the drop-out rate increases; and (3) as the arrival rate increases, the drop-out rate decreases. However, there are also some more subtle comparisons that can be made.

- For instance, let us examine what happens as job turnover increases.
- Here, as the arrival and breakup rates are both scaled up by the same factor, we can see the effects on the drop-out rates.
- Note that such a change leaves the base employment rate (that of an isolated agent) unchanged, and so the differences are attributable entirely to the network effects.
- Table 6 pulls out various rescalings of the arrival and breakup rates for the two cost ranges when $n = 4$ and agents are related through a complete network.
- As before, the first figure is the drop-out rate and the second is the amount attributable to contagion effects.

## **Table 6:** Dependence of Dropouts and Contagion on Arrival and Breakup Rates

| Scaled by $a$ and $b$ | 1 | 3 | 5 | 7 | 9 |
|---|---|---|---|---|---|
| | 0.05, 0.015 | 0.15, 0.045 | 0.25, 0.075 | 0.35, 0.105 | 0.45, 0.135 |
| $c_i \sim [0.8, 1]$ | 69:27 | 76:27 | 83:26 | 88:24 | 96:20 |
| $c_i \sim [0.6, 1]$ | 24:3 | 28:3 | 34:5 | 37:5 | 42:5 |

# VII. Concluding Discussion

- As we have mentioned several times, we treat the network structure as given, except that we consider drop-out decisions.

- Of course, people have more specific control over whom they socialize with both in direct choice of their friendships, as well as through more indirect means such as education and career choices that affect whom they meet and fraternize with on a regular basis.

- Examining the network formation and evolution process in more detail could provide a fuller picture of how the labor market and the social structure co-evolve by mutually influencing each other: network connections shape the labor market outcomes and, in turn, are shaped by them.