Emerging Small-World Referral Networks in Evolutionary Labor Markets

Troy Tassier, Filippo Menczer

Abstract— **We model a labor market that includes referral networks using an agent based simulation. Agents maximize their employment satisfaction by allocating resources to build friendship networks and to adjust search intensity. We use a local selection evolutionary algorithm, which maintains a diverse population of strategies, to study the adaptive graph topologies resulting from the model. The evolved networks display mixtures of regularity and randomness, as in small-world networks. A second characteristic emerges in our model as time progresses; the population loses efficiency due to over-competition for job referral contacts in a way similar to social dilemmas such as the tragedy of the commons. Analysis reveals that the loss of global fitness is driven by an increase in individual robustness, which allows agents to live longer by surviving job losses. The behavior of our model suggests predictions for a number of policies.**

Keywords— **Labor markets, referral networks, local selection, small-world, social dilemmas, affirmative action.**

I. Introduction

We use an evolutionary model to study a robust finding in US labor markets: approximately 50% of workers in the US economy first hear about their job through a friend, relative, or other social contact [1], [2], [3], [4], [5], [6].[1] In particular, we use an evolutionary model based on *local selection* algorithms [7], [8]. Such an approach differs from traditional genetic algorithms in two main ways. First, individual fitness is not determined by comparison to the fitness of other individuals in the population. Instead agents survive and reproduce if they are able to gather enough energy (in fixed supply) to meet a survival requirement. An agent survives if he finds a suitable strategy niche. This difference allows a more diverse population to evolve. Second, since fitness is endogenous (there is no predetermined fitness function), the population size varies depending on the collective strategies of the existing agents. These two characteristics make local selection algorithms a natural choice to study diverse dynamic systems such as economies and markets.

Previous research suggests that reducing the uncertainty of a new hire is the primary reason for hiring by referral [9], [10], [6]. Referral based hiring reduces uncertainty because an employer can be more sure of the quality of a worker if someone she knows refers the worker. Here, we concentrate on another idea. In order for referral based hiring to be effective, social networks must efficiently transfer job information. How are social networks able to do this? People do not usually pick friends in order to get a job. People might pick friends for example because of similar interests or geographic location. Therefore friends tend to exist in clustered groups where the friends of a person tend to know each other. If Jane is a friend of Ken and a friend of Judy, it is likely that Ken and Judy are friends too.

Very clustered (locally structured) social networks inhibit the transfer of information. If everybody whom you know also knows the same people, it will be difficult for information to transfer across social groups. For this reason Granovetter [2] argues that friends outside of the usual social circle of a person are especially important in locating jobs. These friends will hear about jobs that you and your usual circle of friends will not. Therefore it would seem that a network with a smaller amount of local structure would yield a better transfer of job information than a social network.

In this paper we propose to model a simple labor market with an agent based simulation. Agents attempt to maximize individual employment outcomes by choosing a mix of social networks and direct job search. We then view the characteristics of the social networks that develop to establish whether they transfer information efficiently using measures existing in the social network literature.

Our model yields two main results. First, the evolved social networks resemble small-world networks [11]. There is a large amount of local structure in the evolved networks. But the local structure does not significantly affect the ability of the network to transfer information between members of the population. The result indicates that the local structure we observe in real social networks does not preclude efficient information transfer.

Second, as evolution progresses agents over-compete with regard to the amount of the search strategy employed. Agents evolve toward expending more energy (on either social networks or direct search, depending on the relative costs) than is efficient for the population. However, the strategy chosen appears individually efficient; agents live longer on average than at the point of maximum population efficiency. This result may be characterized as a social dilemma because cooperation between agents would serve to benefit the population as a whole but doing so is not optimal for the individual agents. The second result helps us to understand the difference between global (population) and local (individual) efficiency in our evolutionary model.

A detailed description of our model and the evolutionary algorithm we employ is given in the next section. We

`http://dollar.biz.uiowa.edu/~fil/Papers/ieee-tec-ace.pdf`

Troy Tassier is with the Economics Department at the University of Iowa, W243 PBB, Iowa City, IA 52242, USA. Phone: 319-335-1403. Fax: 319-335-1956. Email: `troy-tassier@uiowa.edu`

Filippo Menczer is with the Management Sciences Department at the University of Iowa, S320 PBB, Iowa City, IA 52242, USA. Phone: 319-335-0884. Fax: 319-335-0297. Email: `filippo-menczer@uiowa.edu`

[1] Granovetter [3] provides an excellent review of the history of referral based hiring research through 1995.

then characterize our results and place them in the context of previously existing research in Section III. How can programs that limit the importance of referral based hiring (such as affirmative action) be used to curb over-networking in the form of "old-boys networks"? This and other policy considerations based on our findings are discussed in Section IV.

II. The Model

The setup of our model has two components: jobs and agents. Jobs in the economy are fixed in supply, with randomly assigned, uniformly distributed wages. Agents search for open jobs in the economy using two methods: first, each agent is able to directly search for jobs. Expending effort on direct search increases the likelihood an agent finds a job on her own. We model direct search as an abstraction of methods by which a person may improve the likelihood of finding a job. For instance many individuals look for employment by reading classified advertisements in a newspaper, by applying to various firms directly, or by improving education. All of these methods are considered direct search in our model.

Additionally, agents may choose to expend effort by making friends. Friends are valuable because agents tell their friends about the jobs they find through direct search. We use the term "friend" to refer to all activities by which a person may learn about a job from another person. Examples include: social events, telephone calls, electronic mail, or personal conversations.

Agents will be subject to selection pressure in the economy. Each search method has an energy cost for the agent. Agents gain energy by earning wages. The agents who run out of energy die and the agents who gain in excess of a threshold reproduce by cloning a duplicate agent with similar search strategies and social networks.

Specifically, there are N agents in the economy indexed as i. Note that N will change as the economy evolves. Let $F_i = [F_{i1}, F_{i2}, ..., F_{if_i}]$ be the vector of friends of agent i where each element refers to an agent currently alive in the population. The size of F_i is f_i, the number of friends of agent i. We define a social graph or referral network by defining each agent i as a vertex and creating an unweighted directed edge from each element of F_i to i since every element of F_i is another vertex in the graph. Define an edge connecting a specific friend, j, of agent i as edge (i, j). Agents define an evolving social graph through their choices of friends.

Let $s_i \in (0, 1)$ be the amount of effort used in direct search by agent i; and let e_i be the energy held by agent i. Let $W = [w_1, w_2 \ldots w_{N_0}]$ be the wages associated with each job in the economy. Note that the number of jobs N_0 is a fixed parameter. If agent i holds a job his wage is w_i, which corresponds to one of the wages in W. Each job can only be occupied by one agent or a job may be vacant.

In each period t, each agent currently alive is charged unit costs c_s and c_f for the amount of effort spent searching and for the number of friends, respectively. The agent receives replacement energy from his wage, w_i. There is also a fixed energy cost for each period, c_{min}. Therefore:

$$e_i(t+1) = e_i(t) - c_{min} - s_i(t)c_s - f_i(t)c_f + w_i(t). \quad (1)$$

Now define a reproduction threshold as θ. If $e_i(t+1) > \theta$, agent i reproduces and a new agent j is formed. If $e_i(t+1) < 0$, agent i dies. The evolutionary dynamic, the distribution of jobs and wages, the costs of friends and searching, and the network topology determine a carrying capacity for the population in the economy. If agents use energy efficiently (through their choices of search and friends) the size of the population increases.

In each period, the level of s_i and f_i is subject to a shock or mutation for every agent. Specifically, a draw from a uniform distribution $[-\delta_s, +\delta_s]$ is added to the agent's level of search intensity. And, with probability δ_f, the agent gains or loses a friend. These mutations provide the necessary variation for the members of the population to evolve toward an optimal network and search strategy.

A. The Simulation Algorithm

The simulation runs according to a local selection algorithm. The algorithm's connection with ecological models is discussed at length elsewhere [7], [12]. Here we focus on the characteristics of the algorithm that are employed in our job market model. The fitness of an agent is not derived by ranking the members of the population. Instead, an agent survives if she gathers enough energy to exceed her costs. If an agent finds a strategy niche that sustains her, she survives even if other members of the population do better. Thereby local selection maintains diversity in a population. Additionally, because fitness is not predetermined, the systems defines an endogenous carrying capacity. A system subject to local selection carries as many agents as are able to find successful strategies. Because of the ability to maintain diversity in strategies and the ability to evolve a population size, local selection algorithms provide a natural choice for modeling a labor market.

We provide an outline of the algorithm in Figure 1. In period 0 of the simulation N_0 agents are created with each agent i having an initial level of search intensity, $s_i = s_0$, and an initial number of friends, $f_i = f_0$, for all i. We set the initial levels of each as a parameter of the model. All agents have the same strategy in the initial period. For each agent i, f_i directed non-duplicated edges are created from random agents in the population to i. The set of these edges across all agents defines an initial social network for the population. Next, for each of N_0 jobs a wage is chosen from a uniform distribution in $(0, 1)$. We randomly assign each of these jobs to one member of the initial population. Therefore there is full employment at period 0.

For each period following period 0, all agents go through three loops: job search, energy allocation, and mutation. In the job search loop agents are chosen in random order. When an agent i is chosen she is allowed to attempt to upgrade her job. To do so we maintain a list of the open jobs in the economy. For each open job j in the economy, each agent i in the population is notified about j with probability s_i. Agent i may learn of job j on her own or one of

```
Initialize N_0 agents with s_0 search
Initialize social graph with f_0 random edges per agent
Initialize N_0 jobs with random wages
Assign each job to an agent in the population
for Z periods
for each agent i
Notify i of available jobs based on s_i
Send information through the social network
Let i pick highest wage job
endfor
for each agent i
Collect costs c_s s_i + c_f f_i + c_min
Pay wage w_i to i if employed
if (e_i > θ)
Let i reproduce
else if (e_i < 0)
Let i die
endif
Update social graph
endfor
for each agent i
Fire i with probability α
Mutate s_i and f_i
Update social graph based on new f_i
endfor
Measure statistics of the evolved social graph
endfor
```

Fig. 1. Psuedocode of the local selection evolutionary algorithm used to simulate our job market.

her f_i friends may tell her about a job. Information about a job only travels one step along the social network. If an agent tells a friend about a job, the friend is not allowed to tell an additional friend. We chose to limit the number of steps information travels across the network to model the notion that referral based hiring adds certainty to the hiring process. If information travels many steps the additional certainty an employer gains about a worker through referral decreases because the referee is not likely to know the potential employee well. Agent i picks the job with the highest wage among the open jobs she hears about, or keeps her current job if it has a higher wage. We continue the job search loop until every agent has had an opportunity to search for a job.

Following the job search loop, we allocate energy. Each agent pays costs c_f and c_s for each of the friends he maintains and the effort he expends in searching for jobs. He also pays an additional energy cost c_{min} in each period. If agent i is employed he is paid a wage w_i that corresponds to the job he holds. Each period agents collect the difference between costs paid and wages received in an energy reserve, e_i, that is updated according to Equation 1.

If e_i is greater than the threshold θ the agent reproduces by producing a clone with identical strategies and social network. If e_i is less than 0 the agent dies. When agent i replicates a new agent j, she replicates her current strategies and social network as well: $s_j(t) = s_i(t)$ and $F_j(t) = F_i(t)$. The energy of agent i is split equally between i and j: $e_i(t+1) = e_j(t+1) = e_i(t)/2$. When an agent dies all incoming edges are deleted. All outgoing edges are replaced as follows: if agent i dies, and agent i has an outgoing edge to agent k, a new edge is created from a randomly chosen agent h to agent k that replaces edge (k, i). Replacing these edges keeps the number of incoming edges constant for all surviving members of the population. If this was not done there would be a downward bias in the number of edges for the population.

The final loop of each period of the simulation introduces the variations and mutations necessary for the population to evolve. We fire each agent with probability α in every period. Firing agents provides an incentive for them to maintain search even if employed. If agents were not occasionally fired, the agents occupying jobs paying wages greater than c_{min} would have no incentive to search for a better job; they would live forever without further effort.

There are two sources of mutation that allow strategies of individual agents to evolve. First, with probability δ_f each agent gains or loses a friend in each period. If an agent i loses a friend, a random incoming edge to agent i is deleted. If agent i gains a friend, a random agent j is selected (as long as j is not already a friend of i) and an edge (i, j) is created. Second, a mutation to the search intensity of each agent occurs. A random draw from the uniform distribution $[-\delta_s, +\delta_s]$ is added to the current search intensity s_i for each agent i. The job search, wage, and mutation loops are repeated for Z periods in each run of the simulation.

B. *Measuring the Networks*

We characterize our results using measures introduced in the small-world networks literature [11], [13], [14]. We measure the global structure of a graph using the *diameter* or *path length*, L. The path length is the average of the shortest path between all pairs of vertices in the graph, i.e. the average number of edges that must be traveled along the shortest path between all pairs of vertices; $L = \frac{1}{P}\sum_{p=1}^{P} L_p$ where $P = N(N-1)$ (excluding paths from a node to itself) is the number of pairs of nodes in the graph; p indexes each individual pair and L_p measures the distance between each individual pair p. Our algorithm does not guarantee that an evolved graph will be fully connected. Therefore there may be some pairs of nodes for which no path exits, i.e., $L_p = \infty$. To accommodate for a non-connected graph we actually use the following alternative definition of diameter [15]:

$$L = \left(\frac{1}{P} \sum_{p=1}^{P} L_p^{-1} \right)^{-1}. \qquad (2)$$

The measured path length of a social network will help us quantify the ability to transfer information across the network. A short path length implies that information flows easily (efficiently) in the network. The longer the path length, the more difficult it is for agents to find out about jobs.

To quantify the local structure of a graph, define a *clustering coefficient*, C, as follows: if agent i has f_i adjacent neighbors in the social graph, then this neighborhood defines a subgraph with at most $f_i(f_i - 1)$ edges. C_i is the fraction of these edges that exist in the neighborhood of i and C is the average C_i for all i:

$$C = \frac{1}{N} \sum_{i=1}^{N} C_i. \qquad (3)$$

Essentially C is the likelihood that agents j and k are connected given that each is connected to i. It therefore measures the social "cliquishness" of the graph. We measure average clustering, C, average search intensity, number of friends, and population size each period. We measure average path length, L, at the end of every run.[2]

We normalize our results using random graphs. Random graphs provide us with an approximation to the theoretical lower bound for minimal path length (see [13] p. 501). In addition, random graphs contain minimal structure at the local level. Since all edges are equally likely to be connected, a random graph provides a lower bound on the clustering of a graph. Therefore, random graphs represent the natural choice as a benchmark. If the statistics from the evolved graph are similar to the random graph we infer that the evolved graph is random in nature. Note that all edges, except for those created through selective reproduction, are created at random. Therefore any deviation from randomness in the evolved graph indicates an evolutionary selection toward the characteristics of the graph observed.

When each simulation is complete we create a set of G random graphs as follows: we create a graph g with the same number of nodes and edges as in the evolved graph. We randomly choose two nodes, i, j, for each edge and connect them thereby creating edge (i, j). When we have done this for every edge we measure average path length, L_g^{rand}, and clustering, C_g^{rand}, for the randomly created graph according to Equations 2 and 3. We average these quantities over the G random graphs and then normalize each evolved graph by the corresponding random diameter and clustering coefficient. We repeat this process for each run r of the simulation. Finally, we average the ratios over R runs using a given set of parameters, to obtain the normalized ratios:[3]

$$L_{avg} = \frac{1}{R}\sum_{r=1}^{R}\frac{1}{G}\sum_{g=1}^{G}\frac{L_r}{L_{g,r}^{rand}} \tag{4}$$

$$C_{avg} = \frac{1}{R}\sum_{r=1}^{R}\frac{1}{G}\sum_{g=1}^{G}\frac{C_r}{C_{g,r}^{rand}}. \tag{5}$$

III. Results

We obtain the results from a set of 10 simulations with the parameters shown in Table I. We randomly seed each simulation and present results averaged over the simulations.

TABLE I
Parameters for Base Simulations

Parameter	Value
δ_f	0.001
δ_s	0.001
α	0.02
s_0	0.01
f_0	2
c_s	1
c_f	1/50
c_{min}	0.30
N_0	50
θ	10
Z	200

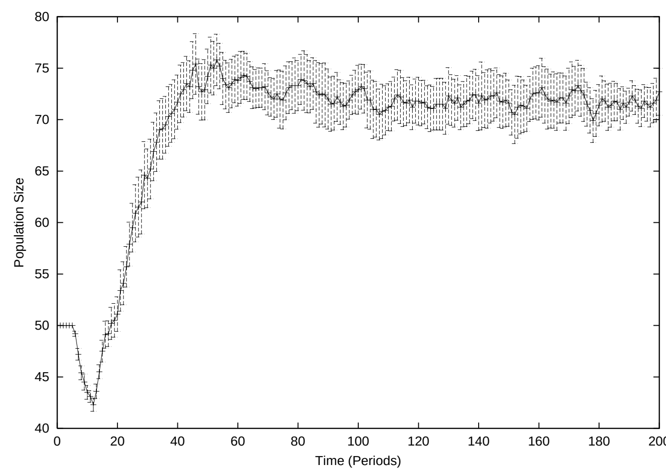

Fig. 2. Evolution of population size

A. Small-World Effect

In a local selection evolutionary algorithm, population size measures the efficiency with which agents use energy. In our model if agents are efficient at finding jobs and transferring information between individuals the population will be high. Figure 2 shows the average population size over the 10 runs. Note the large increase in population between the first and last periods. The population begins with 50 agents but the initial randomly assigned search strategies are poor and the population drops off slightly to around 43 agents after about 12 periods. The agents then begin developing better job search strategies and the population quickly increases to a high of approximately 75 agents before settling to a stable population just above 70 agents.

Regarding the strategies used to accomplish the population increase, the agents have increased both friends and direct search with respect to their initial values. In the initial period each agent had 2 friends. The agents now have an average of 2.22 friends.[4] Direct search has increased from .0100 in the first period to an average of .0114 at the end of the simulation, but this change is not statistically significant for most runs. Interestingly, the large population increase has occurred for relatively small changes in number of friends and intensity of search. A larger change, however, has taken place in the topology of the evolved networks. Let us first take a qualitative look at such change.

[2] Measuring L is very expensive in terms of computation time. Measuring L for a graph with 50 nodes and an average of 5 edges per node takes approximately 2 minutes on a 400MHz PII.

[3] For the results reported below $G = 10$ and $R = 10$. Ten observations on each of these parameters was sufficient to provide statistically significant results concerning the graph statistics.

[4] All runs are significantly different from the starting value of 2 at the 99% confidence level.

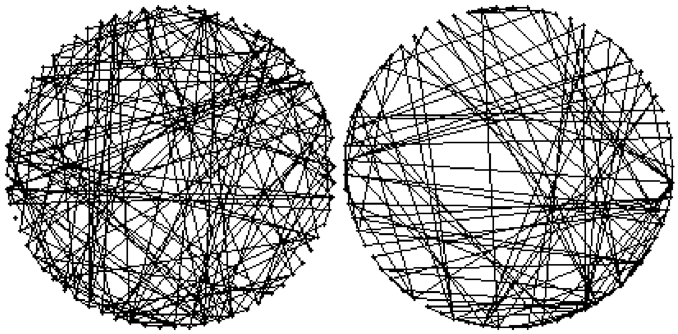

Fig. 3. Topology of one of the 10 evolved networks (right) and of a random network with the same number of nodes and edges (left). To draw these graphs we start with a random agent and iteratively pick a neighbor with a greedy procedure. For each node i we find the agent with the most similar social network by maximizing a score, over agents j not previously placed on the graph, based on the overlap between friends of i and friends, and friends of friends, of j. We repeat the process until all nodes are placed on the circle. If an agent selected has no friends in common with any remaining agent, a new random agent is selected.

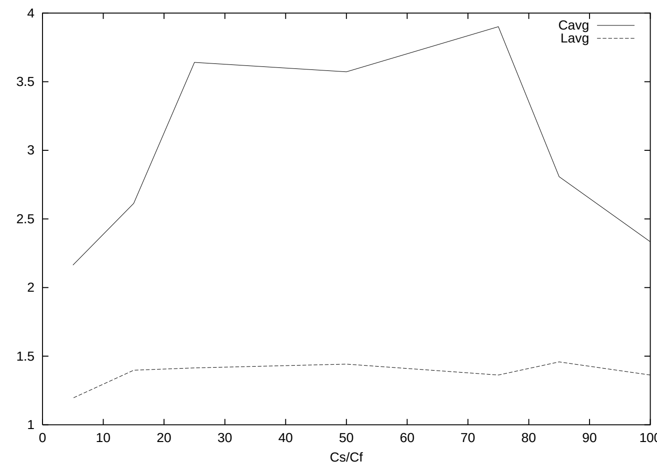

Fig. 4. Small-world effect. A t-test reveals that while L_{avg} is insensitive to the strategy cost ratio, C_{avg} satisfies the small-world definition only in the middle region. A t-test on C_{avg} at the middle and end points of the region reveals statistically significant differences between the sets of simulations (t-statistic=2.065 between $c_s/c_f = 50$ and $c_s/c_f = 5$, and t-statistic=2.022 between $c_s/c_f = 50$ and $c_s/c_f = 100$).

Figure 3 illustrates the clustering of friends produced by the adaptive process. The problem of placing nodes on a plane to minimize edge crossing is NP-complete [16]. Therefore, for more clustered graphs, the greedy procedure used to draw the graphs will result in fewer "long" edges (across the middle of the network) and in denser connections at the periphery. Upon comparing the two graphs, the difference in local structure is evident. Many more of the agents in the evolved graph are friends of their neighbors and also friends with the friends of their friends. The agents have organized a highly clustered population.

As discussed above, clustered networks generally are poor at transferring information because the average path length between nodes tends to be long. However, Watts and Strogatz [11], [13] have quantified how clustered networks can transfer information efficiently. Starting with a "regular" graph (say, a grid) they replace existing edges of the graph with random edges and find that the characteristic path length of the graph decreases to that of a random graph quickly. The clustering of the graph falls as well, but not nearly as quickly as the path length. A small-world graph is then roughly defined as a graph which exhibits a path length close to that of a random network but with clustering close to that of a regular network [13]. Graphs of film appearances of actors ("Kevin Bacon numbers"), power grids, and neural networks are used as real world examples to show that small-world networks exist in reality [13], [17].

By gauging the diameter and clustering measures of our evolved network with respect to random networks, we are able to quantitatively compare the evolved structure to the definition of a small-world network using C_{avg} and L_{avg}. Results from the ten simulations with the parameters above yield $C_{avg} = 3.57$ and $L_{avg} = 1.44$ after 200 periods; therefore our graphs resemble a small-world graph ($C_{avg} \gg 1$ and $L_{avg} \approx 1$.) From this result we infer that clustered social networks still transfer job information efficiently. The local structure observed in social networks does not preclude the efficient transfer of information.

As a sensitivity analysis of the small-world effect we examine a range of strategy costs. Leaving $c_s = 1$ we vary c_f. The values for C_{avg} and L_{avg} at the end end of the runs are plotted vs c_s/c_f in Figure 4. The amount of local structure depends on the relative costs of friends and direct job search. For low c_f, many friends are easy to maintain and therefore the structure does not matter as much. For really high c_f, friends are too expensive to maintain and the local structure disappears. But in the critical region of intermediate cost ratios we observe a consistent small-world phenomenon due to emergent local structure.

B. *Second Order Evolution*

In the previous subsection we observed that agents are able to produce efficient networks in a short period of time. As a further test of the robustness of the efficiency found above we increase the length of the simulations to 5,000 periods. Again we run 10 simulations and obtain results from averages over the runs. Figure 5 displays the average population size. After reaching a maximum around the 150th period, the population size steadily decreases until reaching an equilibrium level of about 63 agents around period 3,500. As we stated above, population size measures how efficiently a system uses energy. Therefore it appears that the agents in the population have lost some of the efficiency gained early in the simulation.

Figure 6 displays the average number of friends over the same period. Average friends increase dramatically and the increase corresponds strongly to the decrease in population and efficiency loss. The correlation between average friends and population size is -0.896.

To quantify the degree of efficiency loss we compare the evolution of the population to that of a population with

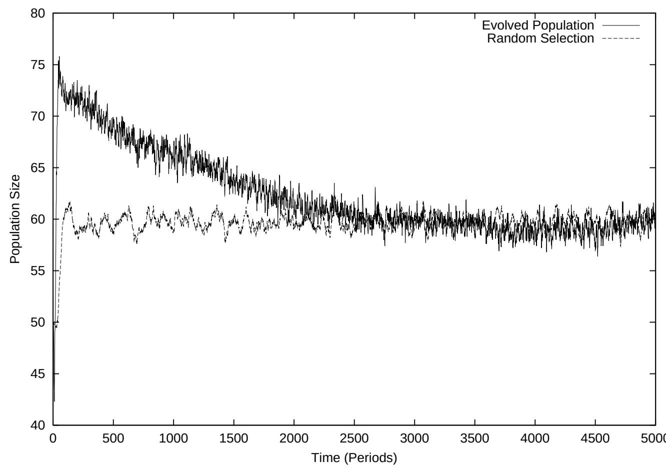

Fig. 5. Second-order evolution of population size. The population dynamics of a random population is also plotted for comparison. A t-test confirms that the two population sizes are significantly different at the height of the evolved population (99% confidence level, t-statistic=4.386).

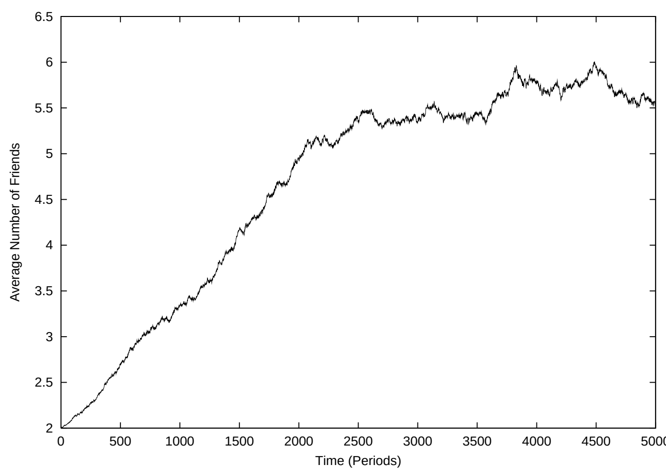

Fig. 6. Evolution of the number of friends in the population. The increase in the number of friends between periods 0 and 5000 is statistically significant (t-statistic=9.639).

random selection. Random selection is accomplished using the algorithm of Figure 1 with two variations: first, wages are assigned randomly; each period every agent is assigned the wage of a randomly chosen agent from the population regardless of his job status. Second, costs are fixed for all agents regardless of strategy choice. For comparison to the original simulation we choose the fixed cost to be the average (total) cost paid by the agents in the last period of the evolved simulation. Figure 5 also plots the random selection population size over 5,000 periods.

As can be seen in the figure, the size of the evolved population is much larger than the size of the population with random selection at the height of its early peak. The average population at period 200 is approximately 72.4 (standard error 2.49) for the evolved agents and the average population for the randomly selected agents is 58.6 (1.93). But the evolved population size steadily decreases over time until converging to a population level close to that of random selection. The decrease to the level of random selection implies that the population has lost all of the efficiency gain from the structure of the network observed early in its evolution.

In order to better understand the loss of efficiency we explore the individual fitness of agents. Following ecological models we can measure individual fitness of agents by the age of agents in the population [7]. The average age of the agents in the 200th and 5,000th periods of the simulation are 57.4 and 69.9 periods, respectively.[5] The average agent in the 5,000th period is twenty-percent older than the average agent in the 200th period. While the evolved strategy of the agents in the 200th period is more efficient for the population (highest observed population size), individual agents can do better (live longer) by increasing the number of friends maintained.

Increasing the number of friends serves two purposes: first, an employed agent with many friends finds out about better jobs with a higher likelihood. In this sense friends provide an opportunity for upward mobility. Second, if an agent becomes unemployed his expected duration of unemployment is lower if he has more friends. Friends provide insurance in the case of job loss. But each of these advantages is offset by the higher maintenance cost that must be paid to keep many friends.

In addition the search strategy chosen depends on the strategies chosen by other members of the population. For instance having many friends does little good if everyone else in the population never searches. In this case the friends of an agent never learn about any jobs and thereby cannot provide her with any job information. Also, since the number of jobs in the economy is fixed and only one agent can occupy each job, agents must compete with each other over the jobs. If agent j finds out about a job before agent i does, agent i cannot get that job until agent j vacates it. Agents must choose a combination of search strategies that do well given the search strategies of other population members. Therefore it may be advantageous to try to have more friends than the status quo members of the population, as long as it is not too costly. Other members of the population may then respond to your strategy by increasing their number of friends. The dynamic results in an "arms race" for social networks where everyone picks a strategy that is best for the individual at the expense of the population.

Therefore the observed phenomenon is similar to a social dilemma in the spirit of the prisoner's dilemma and the tragedy of the commons. Here we have a resource in fixed supply, jobs, that provides the population with energy. Agents do not choose the efficient strategy as a collective because it is not a stable equilibrium. Instead they maximize individual fitness; a strategy that benefits the individual but hurts global efficiency.

To confirm the above interpretation we perform two final experiments. In the first, we decrease the number of friends

[5] A t-test confirms that the average age at the 200th period is significantly different from the average age at the 5,000th period (99% confidence, t-statistic=3.135).

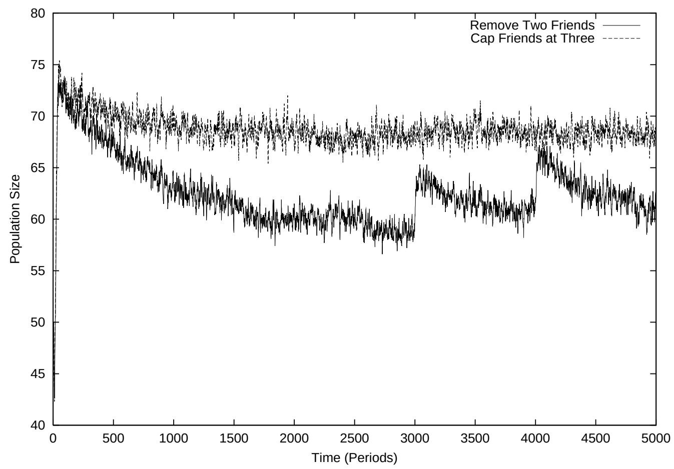

Fig. 7. Population dynamics in two control experiments. In one experiment, two friends are taken from each agent at period 3,000 and 4,000. The population size increases following each period. In the second experiment, the agents are only allowed a maximum of three friends. Capping the number of friends increases the long run population size and cures the social dilemma of over-competition observed previously.

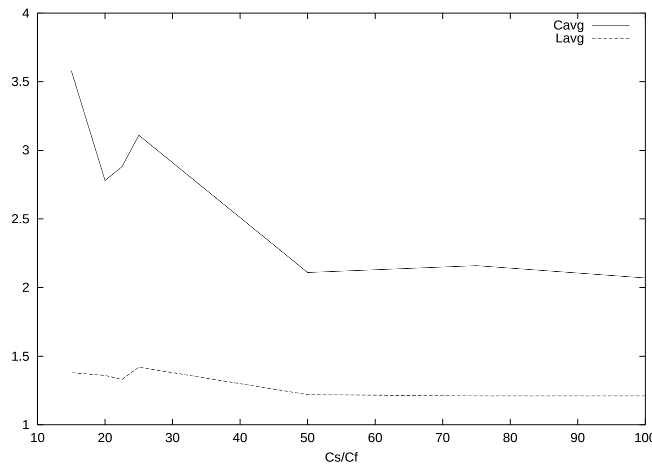

Fig. 8. Small-world effect at period 5,000. As above we compare the difference in C_{avg} across the cost ratio levels by performing a t-test. While the difference between $c_s/c_f = 50$ and $c_s/c_f = 15$ is statistically significant (t-statistic=2.630), the difference between $c_s/c_f = 50$ and $c_s/c_f = 100$ is not (t-statistic=0.211).

held by every agent in the population by two at the 3,000th and 4,000th periods. As can be seen in Figure 7 the average population size increases by about 5 agents directly following the removal of friends. However, shortly after the new peak, the population size follows the previously observed decrease.

In the second experiment we place a restriction on agents that does not allow them to have more than three friends. As seen before in Figure 2, the population quickly increases to about 75 agents. Following the peak the population size decreases somewhat, as shown in Figure 7, but reaches a steady state around 70 agents. In contrast to previous simulations the population size does not decrease further and remains around 70 agents for the duration of the simulation. Limiting the over competition (over networking) cures the social dilemma.

A final note on second order evolution can be made upon viewing Figure 8, which plots C_{avg} and L_{avg} against c_s/c_f for the networks observed at the 5,000th period. Observe that the structure of the resulting networks remains similar to that of the shorter runs. As observed previously in Figure 4, the path length ratio remains constant between 1.3 and 1.5 for all cost ratios. The clustering ratio increases with the cost of friends; however, it takes higher c_f to generate the same magnitude increase. Since the agents have increased the number of friends they hold, each individual friend is less important. Therefore it takes a larger decrease in the number of friends (caused by more expensive friends) to generate the small-world effect.

IV. Policy Discussion

One of the leading topics in the discussion of referral based hiring has been the effect on disadvantaged social groups [18], [5], [3], [1], [19]. The research makes two main points about inequality in regard to referral based hiring: first, social segregation may lead to wage inequality if referral based hiring is prevalent. When many jobs are filled through referral, individuals in poor social networks may have fewer opportunities to apply for jobs if firms do not also use additional search methods such as advertising or other forms of job posting. In this fashion referral based hiring limits the job opportunities available to individuals outside of the "old-boys network."

Second, an individual with sufficiently better connections may expect more job offers than someone with worse connections. The increase in the number of offers has the potential to increase the wages of a better connected individual by improving his bargaining position. For well-connected individuals "it's not what you know but who you know" [5].

One policy approach that affects the potential inequality created by referral based hiring is affirmative action. With regard to our model, affirmative action could help create equality in two ways. First, it may be thought of as a way to limit over-competition for friendship networks and to cap the social dilemma problem. As observed at the end of the previous section, limiting the amount of networking increases the efficiency of the job market.

A second way that affirmative action may be useful is by limiting unintentional discrimination on the part of employers. For instance, if the percent of white workers in a firm is high and social networks are stratified by race, applicants to this firm who are referred are likely to be white as well.[6] If employers commonly hire by referral and get few minority applicants, the likelihood a minority is hired decreases even if the employer treats minorities equally with regard to the hiring decision. In this way employers may unintentionally discriminate against minorities

[6] For example Mouw (see [19] p. 89) finds that firms who employ few black workers (.05%-10% of the workforce being black) are 75% less likely to hire a black worker if they hire through referral instead of using a newspaper advertisement.

without any ill-will. Affirmative action can remedy the problem by encouraging socially segregated firms to hire workers by means other than referral and thereby increase the diversity of applicants.

Another issue raised in our model is in the relative costs of the various means to search for jobs. The advent of the Internet has certainly changed the ways that individuals maintain friendships. Electronic mail, for instance, has made it cheaper to keep in contact with friends, especially periphery friends that one does not see often. What implication does this have for our results above?

As friendships get cheaper to maintain, agents spend a larger portion of their time networking. Depending on the degree of integration in a society, our model predicts that poorly networked people may fair better or worse. If a population remains highly segregated as the use of networking increases, poorly-connected people are shut out from an increasingly important method to access jobs. But, in our model we also observe that as networking increases the structure of the social networks becomes less local. The social networks are wider and the amount of clustering relative to a random graph decreases. Therefore it may be easier for a poorly connected person to break into the "job loop."

For real labor markets the question becomes "which effect will dominate?" Will the less local structure of social networks allow poorly connected individuals to expand into well connected social networks? Or will inexpensive friendship maintenance cause extreme over-competition with respect to individuals in well connected networks and thereby increase inequality? If networks become so cheap that everyone knows everyone else (the network is sufficiently dense), disadvantaged individuals will prosper along with advantaged individuals. But there is a problem: if the real reason for hiring by referral is decreasing uncertainty about the quality of a new hire, it is not likely that an outsider to the "old-boys network" will be able to use a connection to a distant person not known well. And, an employer only gains confidence through referral if the referee knows the applicant well. With this in mind, individuals newly connected to the "old-boys network" may not be able to profit from a better network position until they become well established in the new network. Therefore an increase in the ease of maintaining friendships has unclear implications for creating equality.

As a final policy consideration we may think of a third job search strategy that may be added to our model: education. While education is not directly a search strategy it is important (or required) for getting many jobs. If we incorporate education as a choice variable that helps agents get jobs it is likely that we will still get over-networking given a sufficiently low cost of maintaining friends. If this is true the high dependence on friends for attaining jobs may reduce the value of education in a society — potentially an even worse social dilemma.

Another possibility with respect to education more directly concerns poorly connected people. If social networks become increasingly important in finding jobs even a well educated person may have difficulty finding employment if he is not well connected. In turn this may affect his education choice in two ways: first, if he is not well connected his value of education may be lower than another person with better connections. This is because he may not be able to capitalize on the education he holds. Second, if networks are important in finding jobs, a person has an incentive to choose the occupation for which he has the best network. If all of his friends have low education jobs it may be optimal to choose the same low education profession because it may maximize the likelihood of becoming employed. If so, individuals may encounter poverty-traps where low education choices are reinforced by the other members of a population group.

V. Conclusion

We used an evolutionary model with endogenous fitness to study the efficiency of information transfer in a labor market setting. In our model small-world referral networks emerge as an efficient means to transfer information in the labor market. We observe a significant amount of local structure in the evolved networks but the local structure does not preclude the efficient transfer of information as measured by the average path length. As evolution progresses we observe a second characteristic: the population loses efficiency over time as indicated by a decrease in population size. Investigation reveals that over competition for social networks cause the efficiency loss. The result is reminiscent of traditional social dilemmas such as the prisoner's dilemma, in that agents maximize individual (local) efficiency at the expense of population (global) efficiency.

The predictions of our model allow us to make some considerations on the inequality issues related to the prevalence of referral based hiring methods and the structure of social networks. Additional research directly related to these policy areas is possible within our model with minor modifications. Of particular interest may be the effect of referral based hiring on education choices. By adding education as a choice variable to our model we can view how the education a person chooses is affected by her referral network. This approach would add to considerations of peer effects in the attainment of education.

Another area we hope to consider is the effect of referral based hiring on social segregation. Individuals may have incentives to open or close a social network for the purposes of gaining jobs. As an example, if a person has many unemployed friends he would like to meet a group of employed people. But, since his friends are unemployed he is of less value than a better connected person as a job contact. People in poorly connected groups would like to join a better connected group but the people in this group would like to remain exclusive to lessen competition for jobs. Studying this dynamic would add to our understanding of the effect that labor markets and financial prosperity have on efforts to better integrate society.

Acknowledgments

We are grateful to Mary Beth Combs for her comments on a previous draft of this paper. The first author acknowledges support from the MacArthur Network on Social Interactions and Economic Inequality.

References

[1] Kenneth Arrow, "What has economics to say about racial discrimination?," *Journal of Economic Perspectives*, vol. 12, pp. 91–100, 1998.
[2] Mark Granovetter, "The strength of weak ties," *American Journal of Sociology*, vol. 78, pp. 1360–1380, 1972.
[3] Mark Granovetter, *Getting a Job: A Study of Contacts and Careers*, Harvard University Press, Cambridge, MA, second edition, 1995.
[4] Harry Holzer, "Search method use by unemployed youth," *Journal of Labor Economics*, vol. 6, pp. 1–20, 1988.
[5] James Montgomery, "Social networks and labor-market outcomes: Toward an economic analysis," *The American Economic Review*, vol. 81, pp. 1408–1418, 1991.
[6] Curtis Simon and John Warner, "Matchmaker, matchmaker: The effect of old boy networks on job match quality, earnings and tenure," *Journal of Labor Economics*, vol. 10, pp. 306–329, 1992.
[7] F Menczer and RK Belew, "Latent energy environments," in *Adaptive Individuals in Evolving Populations: Models and Algorithms*, RK Belew and M Mitchell, Eds., Santa Fe Institute Studies in the Sciences of Complexity. Addison Wesley, Reading, MA, 1996.
[8] F Menczer, M Degeratu, and WN Street, "Efficient and scalable pareto optimization by evolutionary local selection algorithms," *Evolutionary Computation*, vol. 8, no. 2, pp. 223–247, 2000.
[9] Albert Rees, "Information networks in labor markets," *American Economic Review*, vol. 56, pp. 559–566, 1966.
[10] Albert Rees and George Schultz, *Workers in an Urban Labor Market*, University of Chicago Press, Chicago, 1970.
[11] Duncan J Watts and Steven Strogatz, "Collective dynamics of "small-world" networks," *Nature*, vol. 393, pp. 440–442, 1998.
[12] F Menczer and RK Belew, "From complex environments to complex behaviors," *Adaptive Behavior*, vol. 4, pp. 317–363, 1996.
[13] Duncan Watts, *Small World Networks: The Dynamics of Networks between Order and Randomness*, Princeton University Press, Princeton, 1999.
[14] Mark Newman, "Models of the small world: A review," *J. Stat. Phys.*, Forthcoming.
[15] Mark Newman, ," 2000, Personal communication.
[16] M Garey and D Johnson, "Crossing number is np-complete," *SIAM Journal of Algebraic and Discrete Methods*, vol. 4, no. 3, pp. 312–316, 1983.
[17] James Collins and Carson Chow, "It's a small world," *Nature*, vol. 393, pp. 393–410, 1998.
[18] William Julius Wilson, *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy*, University of Chicago Press, Chicago, 1987.
[19] Ted Mouw, *Four Essays on the Social Structure of Urban Labor Markets*, Ph.D. thesis, University of Michigan, Department of Sociology, 1999.