

Contribution for
Innovations in Financial
and Economic Networks
Anna Nagurney, editor

Mark Pingle
Department of Economics
University of Nevada
Reno, NV 89557-0016
<http://unr.edu/homepage/pingle/>
pingle@unr.edu

Leigh Tesfatsion
Department of Economics
Iowa State University
Ames, IA 50011-1070
<http://www.econ.iastate.edu/tesfatsi/>
tesfatsi@iastate.edu

April 11, 2003

1 Evolution of Worker-Employer Networks and Behaviors Under Alternative Non-Employment Benefits: An Agent-Based Computational Study

Mark Pingle and Leigh Tesfatsion

1.1 Introduction

Determining the effects of labor institutions on macroeconomic performance is a central concern of economic policymakers. Differences in labor institutions have been conjectured to be a key explanation for observed cross-country differences in the level and persistence of unemployment, in the distribution of income and wealth, and in growth rates for labor productivity and GDP.

For example, as discussed by Blau and Kahn (1999), Ljungqvist and Sargent (1998), and Nickell and Layard (1999), European OECD countries over the past twenty years have tended to rely on administered wages and legislated job protection and have experienced sluggish job growth and persistently high unemployment. In contrast, the United States has had a relatively

more flexible, less regulated labor market and has achieved much greater job growth and relatively lower unemployment rates. This has led many European policymakers to argue the need for reforms in their labor institutions.

Unfortunately, as discussed by Acemoglu and Shimer (2000), Blau and Kahn (1999), and Freeman (1998), it is difficult to obtain conclusive empirical results regarding how labor institutions affect economic performance. Regression methods relating changes in labor institutions to economic outcomes quickly tax degrees of freedom. This problem is compounded if one institution's impact depends on the presence or absence of other institutions. Also, labor institutions are inherently endogenous. For example, governments continually revise labor institutions in response to economic and political pressures. This endogeneity makes it difficult to interpret the validity of empirical investigations.

In recognition of these difficulties, Freeman (1998, pp. 19-20) suggests that agent-based computational modeling might offer a promising additional way to study the impact of labor institutions, particularly from a market design perspective. Tesfatsion (1998,2001,2002a) reports some preliminary work along these lines. An *agent-based computational economics (ACE)* framework is used to study path dependence, market power, and market efficiency outcomes for a labor market under systematically varied concentration and capacity conditions. ACE is the computational study of economies modeled as evolving systems of autonomous interacting agents (Tesfatsion,2002b,2003).¹

In this study we conduct an ACE labor market experiment to test the sensitivity of labor market outcomes to changes in the level of a non-employment payment. Our ultimate objective is to understand how the basic features of real-world unemployment benefit programs affect labor market performance.² However, as will be clarified below, a human subject experiment was run in parallel to this computational experiment as a check on the reliability of our findings and the adequacy of the learning representations for our computational workers and employers. To facilitate this initial benchmark check, a deliberately simplified experimental design is used.

Specifically, we consider a balanced labor market with equal numbers of workers and employers. In each trade cycle (work period), every worker has one work offer to make and every employer has one job opening to fill. An employer can reject a work offer received from a worker on two possible grounds: unacceptable past work history; or capacity limitations.

¹See <http://www.econ.iastate.edu/tesfatsi/ace.htm> for extensive resources related to ACE, including surveys, an annotated syllabus of readings, research area sites, software, teaching materials, and pointers to individual researchers and research groups.

²For example, in the U.S., unemployment benefits are *financed by taxes on employers* and are intended to provide *temporary* financial assistance to workers who become unemployed *through no fault of their own* and who *continue to seek work*. Understanding the separate and combined impacts of these and other unemployment benefit program features on labor market performance over time is an extremely challenging problem. A detailed discussion of theoretical and empirical labor market studies focusing on unemployment benefits and related issues can be found in Pingle and Tesfatsion (2001).

The workers repeatedly submit their work offers to preferred employers until either they succeed in being hired or they become discouraged by rejections and exit the job market. A worker must pay a small transaction cost each time he submits a work offer to an employer. As in MacLeod and Malcomson (1998), each matched worker and employer individually chooses to shirk or cooperate on the work-site, and these choices are made simultaneously so that neither has a strategic informational advantage. Any worker or employer who does not enter an employment relationship during the trade cycle in question receives an exogenously specified non-employment payment. Workers and employers evolve their work-site behaviors over time on the basis of past experiences in an attempt to increase their earnings.

In this labor market, then, full employment with no job vacancies is possible. Nevertheless, the unemployment rate and the particular set of workers and employers in employment relationships endogenously evolve over the course of successive trade cycles. Two interdependent choices made repeatedly by the workers and the employers shape this evolutionary process: namely, their choices of work-site partners; and their behavioral choices in interactions with these partners.

Three non-employment payment (NEP) treatments are experimentally studied: a zero NEP; a low NEP; and a high NEP. As reported in Table 1.1, one main finding is that the average utility levels attained by workers and employers do not substantially change as the NEP is increased from zero to high. As will be clarified in Section 1.4, although the increase in the NEP increases the worker unemployment rate and the employer vacancy rate, this greater loss of productive activity is offset in part by the higher NEP and in part by the increased levels of mutual cooperation exhibited by the workers and employers who manage to match.

Another main finding reported in Table 1.1 is that a somewhat higher average utility level for workers and employers is attained with a low NEP than with either a zero or a high NEP in the short and intermediate runs (generations 1 through 50). As will be explained more carefully in Section 1.4, a zero NEP encourages shirking on the work-site (low risk of quits or firings in response to defections) while a high NEP results in a high risk of lost earnings due to coordination failure (high risk of quits or firings in response to defections). Interestingly, however, average utility tends to increase over time under each NEP treatment as the workers and employers become better at sustaining mutual cooperation on the work-site. Moreover, this movement towards higher average utility is strongest under the high NEP treatment. Thus, in the long run (generation 1000), the average utility level attained by workers and employers with a high NEP exceeds the average utility levels attained with a zero or low NEP.

On the other hand, program costs should be taken into account as well as utility benefits in order to obtain a more accurate measure of economic efficiency. Let net earnings be measured by the average (per agent) utility

level attained by workers and employers minus the average NEP paid to these workers and employers. Define efficiency to be the ratio of actual net earnings to maximum possible net earnings. As indicated in Table 1.1, although a high NEP results in a high average utility level, it also results in a significantly lower efficiency level than either the zero or low NEP due to high program costs. Overall, considering the short, intermediate, and long run, the low NEP delivers the highest overall efficiency level. Consequently, evaluated in terms of efficiency, our findings indicate that a low NEP is preferable to either a zero or a high NEP.

The aggregate outcomes reported in Table 1.1, while interesting, are only the tip of the iceberg. A careful study of individual experimental runs indicates that the response of the ACE labor market to changes in the NEP is much more intricately structured than this table suggests. As reported in Figure 1.1, the 20 runs generated for each NEP treatment tend to gravitate towards one of two “attractor states.” The configuration of these two attractor states is similar under the zero and low NEP treatments: the first attractor state is characterized by latched pairs of mutually cooperative workers and employers, while the second attractor state is characterized by latched pairs of workers and employers who intermittently defect and cooperate. In contrast, under the high NEP treatment, one attractor state is characterized by latched pairs of mutually cooperative workers and employers while the other attractor state is a state of economic collapse in which each worker and employer ultimately becomes inactive. This apparent existence of multiple attractor states suggests caution in interpreting the aggregate outcomes reported in Table 1.1, since these outcomes could be based on inappropriately pooled data.

The existence of multiple attractor states for each NEP treatment is due to strong network and learning effects. Starting from the same initial structural conditions, chance differences in the initial interaction patterns among the workers and employers can cause the labor market to evolve towards persistent interaction networks supporting sharply distinct types of expressed behaviors. For example, with a high NEP, the labor market evolves either towards a highly efficient economy in which all workers and employers are in long-run mutually cooperative relationships or towards economic collapse with 100% unemployment. Thus, while a change in the level of the NEP can be expected to have substantial systematic effects on key labor market outcomes such as efficiency and unemployment, our findings suggest that these effects will be in the form of spectral (multiple peaked) distributions with large standard deviations.

These computational experiment findings can be compared to findings reported in Pingle and Tesfatsion (2002) for a human-subject experiment using a similarly structured labor market but with a smaller number of workers and employers participating in a much smaller number of trade cycles per experimental session. In the human-subject experiment, as in the computa-

tional experiment, a higher NEP resulted in higher average unemployment and vacancy rates as well as higher average utility levels among those who successfully matched. In the human-subject experiment, however, most relationships that formed between workers and employers were either short-lived or intermittent, with only modest amounts of behavioral coordination in evidence. In contrast, in the computational experiment almost all workers and employers who succeeded in matching ended up in long-run relationships with one partner in which the behaviors of the partners were highly coordinated.

As detailed more carefully in Pingle and Tesfatsion (2002), this difference in findings raises interesting questions. To what extent are the human-subject and computational experiments capturing the same economic structure but reporting over different time scales, short run versus long run? In particular, could it be that the “shadow of the past” weighs heavily on human subjects over the necessarily shorter human-subject trials, biasing behaviors towards unknown past points of reference? If so, the computational experiment might be providing the more accurate prediction of what would happen in actual labor markets over a longer span of time. Alternatively, the two experiments might differ structurally in some fundamental way so that differences in outcomes would be observed regardless of time scale. In particular, is the representation of agent learning in the computational experiment too inaccurate to permit valid comparisons with human-subject labor market experiments? Are the observed differences in types of network formations due to the different frequencies with which transaction costs are incurred due to scale effects? Careful additional studies, both empirical and experimental, will be needed to resolve these questions.

The ACE labor market model is presented in Section 1.2. Section 1.3 outlines the experimental design of our study, and Section 1.4 provides a detailed report of our experimental findings. Concluding remarks are presented in Section 1.5.

1.2 The ACE Labor Market Model

Overview:

The ACE labor market comprises 12 workers and 12 employers. Each worker can work for at most one employer at any given time, and each employer can employ at most one worker at any given time. The workers and employers repeatedly seek preferred work-site partners using a modified form of a matching mechanism (Gale-Shapley, 1962) that has been observed to evolve in various real-world labor market settings (Roth and Sotomayor, 1992). The workers and employers who successfully match then engage in risky work-site interactions modeled as prisoner’s dilemma games. At regular intervals the workers and employers separately update their work-site rules of behavior on the basis of the past earnings obtained with these rules.

The computational experiment is implemented by means of the *Trade*

Network Game Laboratory (TNG Lab), an agent-based computational laboratory developed by McFadzean, Stewart, and Tesfatsion (2001) for studying the evolution of trade networks via real-time animations, tables, and graphical displays.³ The specific TNG parameter settings used for the experiment at hand are described below. All other TNG parameter settings are the same as in Tesfatsion (2001).

Implementation Details:

As depicted in Table 1.2, ACE labor market activities are divided into a sequence of 1000 generations. Each generation in turn is divided into three parts: (a) a trade cycle loop consisting of successive trade cycles during which work-site interactions take place; (b) an environment step in which each worker and employer assesses their current utility (fitness) level as a function of their accumulated earnings to date; and (c) an evolution step in which the workers and employers separately evolve their work-site rules of behavior on the basis of the past earnings attained with these rules.

Each worker and employer in the initial generation is assigned a work-site rule in the form of a randomly specified pure strategy for playing an iterated prisoner's dilemma game with an arbitrary partner an indefinite number of times. This work-site rule governs the behavior of the agent in his work-site interactions throughout the entire trade cycle loop for the initial generation. Each work-site rule is represented by means of a "finite state automaton"⁴ with 16 internal states. Thus, the set of feasible work-site rules for each worker and employer, while extremely large, is nevertheless finite. Each worker and employer in the initial generation also assigns an initial expected utility assessment U^o to each of his possible work-site partners, where U^o is equal to the mutual cooperation payoff.⁵

The workers and employers in the initial generation then participate in a trade cycle loop consisting of 150 successive trade cycles. In each trade cycle they engage in two main activities: (1) a matching process during which they search for preferred work-site partners on the basis of their current expected utility assessments for these partners; and (2) an employment process during which each matched worker-employer pair engages in one work-site interaction. Throughout these processes the workers and employers update their current expected utility assessments for each other every time they obtain a payoff from an interaction with each other.

³See <http://www.econ.iastate.edu/tesfatsi/tnghome.htm> for source code, executables, user instructions, tutorials, and research related to the TNG Lab.

⁴A *finite state automaton* is a system comprising a finite collection of internal states together with a state transition function that gives the next state of the system as a function of the current state and other current system inputs. For the application at hand, the latter inputs are the actions selected by a worker and employer engaged in a work-site interaction.

⁵This is not an innocuous specification, since it strongly affects the extent to which the workers and employers engage in experimentation with new partners. This issue is further considered in Section 1.3.

Each worker and employer also has an exogenously specified minimum tolerance level, assigned as part of the initial generation configuration process. In the current experiment, these minimum tolerance levels are set equal to the non-employment payment. Thus, entering into a risky work-site interaction is viewed as a tolerable gamble if and only if it is expected to yield at least as high a payoff as would be earned through inactivity. If the expected utility assessment assigned to an employer by a worker ever falls below the minimum tolerance level, the worker will stop directing work offers to this employer. Similarly, if the expected utility assessment assigned to a worker by an employer ever drops below the minimum tolerance level, the employer will stop accepting work offers from this worker.

The manner in which workers direct work offers to employers during the matching process for each trade cycle proceeds as follows. Each worker and employer has a preference ranking over possible partners, determined by his current expected utility assessments. Each worker starts by directing a work offer to a most preferred tolerable employer. Each employer receiving at least one tolerable work offer places his most preferred tolerable work offer on his work offer list and refuses all the rest. Each worker having a work offer refused then redirects this work offer to a next most preferred tolerable employer who has not yet refused him in the current matching process, if any such employer exists. Once employers stop receiving new work offers, they accept the work offers currently on their work offer lists and the matching process comes to a close. Throughout this process, ties are broken by random selection.

Once a worker and employer are matched, they enter into a work-site interaction. This interaction is modeled as a prisoner's dilemma game with cooperation interpreted as meeting all work-site obligations and defection interpreted as shirking with regard to these obligations. As depicted in Table 1.3, one of four possible payoffs can be earned in each work-site interaction: a low payoff $L=10$, earned by an agent who cooperates against a defecting partner; a mutual defection payoff $D=20$; a mutual cooperation payoff $C=40$; or a high payoff $H=60$ earned by an agent who defects against a cooperating partner. Also, a worker incurs an offer cost $OC=1.0$ each time he directs a work offer to an employer, whether or not the work offer is accepted. A worker or employer who is not matched earns a non-employment payment (NEP) for the trade cycle. Each worker and employer records all payoffs he receives during the course of each trade cycle, including work-site payoffs, negative payoffs due to offer costs, and non-employment payments.

Each worker and employer uses a simple reinforcement learning algorithm to update his expected utility assessments for possible partners in response to new payoffs. Recall that each agent (worker or employer) initially assigns an initial expected utility assessment $U^o=C$ to each possible work-site partner. Subsequently, each time an agent v interacts with an agent z , agent v forms an updated expected utility assessment for z by summing U^o together with

all payoffs received to date from interactions with z and dividing this sum by one plus the total number of these interactions. The payoffs included in this summation include work-site payoffs and negative payoffs due to offer costs. Consequently, an updated expected utility assessment for any agent z is the average of all payments received to date in interactions with z , augmented to include U^o as a virtual additional payoff. Under this method, if an agent interacts repeatedly with another agent for a sufficient length of time, his expected utility assessment for z will eventually approach his true average payoff level from interactions with z .⁶

At the end of the initial generation, the workers and employers enter into an environment step in which each agent calculates his utility (fitness) level. This utility level is taken to be the average total net payoffs per trade cycle that the agent earned during the course of the preceding trade cycle loop, i.e., the agent's total net payoffs divided by 150 (the number of trade cycles per loop). The workers and employers then enter into an evolution step in which they use their attained utility levels to evolve (structurally update) their work-site rules via both inductive and social learning. Inductive learning takes the form of experimentation; agents perturb their work-site rules by introducing random modifications. Social learning takes the form of mimicry; agents deliberately modify their work-site rules to more closely resemble the work-site rules used by more successful (higher utility) agents of their own type. Thus, workers imitate other more successful workers, and employers imitate other more successful employers.

Experimentation and mimicry are separately implemented for workers and for employers by means of genetic algorithms involving commonly used elitism, mutation, and recombination operations. Elitism ensures that the most successful work-site rules are retained unchanged from one generation to the next. Mutation ensures that workers and employers continually experiment with new work-site rules (inductive learning). Recombination ensures that workers and employers continually engage in mimicry (social learning).⁷

At the end of the evolution step, each worker and employer has a potentially new work-site rule. The memory of each worker and employer is then wiped clean of all past work-site experiences. In particular, initial expected utility assessments for possible partners are re-set to the mutual cooperation payoff level without regard for past work-site experiences. The workers and employers then enter into a new generation and the whole process repeats, for a total of 1000 generations in all.⁸

⁶See McFadzean and Tesfatsion (1999) for more details. Briefly, this long-run consistency property follows from the finite state automaton representation for work-site rules which ensures that the action pattern between any two agents who repeatedly interact must eventually enter into a cycle as the number of their interactions becomes sufficiently large.

⁷See McFadzean and Tesfatsion (1999) and Tesfatsion (2001) for detailed discussions of this use of genetic algorithms to implement the evolution of work-site rules.

⁸A final technical remark about implementation should also be noted, in case others wish to replicate or extend this experiment. The minimum tolerance level is hardwired

1.3 The Computational Experiment

The computational experiment focuses on only one treatment variable, the non-employment payment (NEP). The three tested treatments for NEP are NEP=0, NEP=15, and NEP=30. These three treatments are referred to as ZeroT, LowT, and HighT, respectively.

The interest in these three alternative treatments is seen by comparing them with the work-site payoffs depicted in Table 1.3. In treatment ZeroT, non-employment during a trade cycle results in the payment NEP=0. This is the worst possible trade cycle payoff, worse even than the sucker payoff L=10 that results from cooperating with a defecting work-site partner. In treatment ZeroT, then, unemployment or vacancy is never an attractive alternative to employment or hiring, and the workers and employers will be willing to put up with defections to avoid unemployment or vacancy.

In contrast, in treatment LowT non-employment during a trade cycle results in the payment NEP=15. This payment is strictly higher than the sucker payoff L=10, meaning agents will prefer non-employment to being suckered. Thus, each agent who defects against a cooperative partner to attain a high payoff now faces a risk of future non-employment if this current partner chooses not to interact with him in the future. Finally, in treatment HighT non-employment results in the payment NEP=30. This payment dominates both the sucker payoff L=10 and the mutual defection payoff D=20. Consequently, agents will tend to be much more sensitive to defections, preferring unemployment or vacancy in preference to defecting back against a defecting partner.

For each NEP treatment, 20 runs were generated using 20 different seeds for the TNG Lab pseudo-random number generator: namely, $\{0, 5, 10, \dots, 95\}$. In the data tables reported in Section 1.4, each run is identified by its corresponding seed value. Each run consists of 1000 generations in total. To investigate evolutionary change, the twenty runs for each treatment are sampled at three different points in time: generation 12, generation 50, and

to zero in the TNG Lab, the software used to implement the computational experiment. Thus, to retain the non-employment payment NEP equal to the minimum tolerance level, experiments were actually run with each work-site payoff normalized by subtraction of NEP. In addition, for better TNG Lab visualization, the work-site payoffs were further normalized by multiplication by 0.10. For example, $C^* = 0.10[C - NEP]$ was used in place of the mutual cooperation payoff C, and similarly for the other work-site payoffs. The corresponding normalized non-employment payment then equaled $NEP^* = 0.10[NEP - NEP] = 0$. Finally, to maintain consistency with this normalization, the offer cost OC was normalized to $OC^* = 0.10$. Note that it would *not* be consistent to subtract NEP from OC, since OC is a cost per work offer. For example, a worker who is refused k times and never hired during a trade cycle receives a total payoff NEP-kOC at the end of the trade cycle, and this is the payoff from which NEP must then be subtracted to implement the payoff normalization. This subtraction occurs automatically when $NE^*=0$ is used in place of NEP. In all data tables presented below, utility levels and market power levels are translated back into non-normalized form prior to reporting, for easier comparison with the human-subject experimental findings reported in Pingle and Tesfatsion (2002).

generation 1000. For each sampled generation, data is collected regarding network formation, market non-participation rates, work-site behaviors, welfare (utility and market power) outcomes, and persistent relationship type counts.

Before reporting our experimental findings in detail, it is important to explain carefully the descriptive statistics that have been constructed to help characterize the one-to-many mapping between treatment and outcomes.

Measurement of Persistent Relationships:

As previously noted (see footnote 4), work-site rules are represented as finite state automata, implying that the actions undertaken by any one agent in repeated work-site interactions with another agent must eventually cycle. Consequently, the actions of any one agent in interactions with another agent during a trade cycle loop can be summarized in the form of a work-site history H:P. The “handshake” H is a (possibly null) string of work-site actions that form a non-repeated pattern, while the “persistent portion” P is a (possibly null) string of work-site actions that are cyclically repeated. For example, letting c denote cooperation and d denote defection, the work-site history ddd:dc for an agent v in interactions with another agent z indicates that v defected against z in his first three work-site interactions with z and thereafter alternated between defection and cooperation.

A worker and employer are said to exhibit a *persistent relationship* during a given trade cycle loop if two conditions hold. First, their work-site histories with each other during the course of this loop each have non-null persistent portions. Second, accepted work offers between the worker and employer do not permanently cease during this loop either by choice (a permanent switch away to a strictly preferred partner) or by refusal (one agent becomes intolerable to the other because of too many defections).

A persistent relationship between a worker and employer in a given trade cycle loop is said to be *latched* if the worker works continually for the employer (i.e., in every successive trade cycle) during the persistent portions of their work-site histories. Otherwise, the persistent relationship is said to be *recurrent*.

Measurement of Market Non-Participation Rates:

A worker or employer who fails to form any persistent relationship during a given trade cycle loop is classified as *persistently non-employed* for that trade cycle loop. The percentage of workers who are persistently non-employed constitutes the *persistent unemployment rate* for that trade cycle loop. Similarly, the percentage of employers who are persistently non-employed constitutes the *persistent vacancy rate* for that trade cycle loop.

Classification of Networks by Competitive Distance:

We will next construct a distance measure that permits the classification of experimentally observed “interaction networks” into alternative types. This distance measure will calculate the distance between an experimentally

observed interaction network and an idealized interaction network capable of supporting a competitive (full employment) market outcome.

Recall from Table 1.2 that each generation G of the ACE labor market model consists of a single trade cycle loop plus an environment step and an evolution step. The *interaction network* $N(G,R)$ for a particular generation G in a particular experimental run R refers to the observed pattern of interactions occurring among workers and employers in the trade cycle loop for that generation and run.

Each interaction network $N(G,R)$ is represented in the form of a directed graph. The vertices V of the graph represent the workers and employers. The edges of the graph (directed arrows) represent work offers directed from workers to employers. Finally, the edge weight on any edge denotes the number of accepted work offers between the worker and employer connected by the edge. The reduced-form network $PN(G,R)$ derived from $N(G,R)$ by eliminating all edges of $N(G,R)$ that correspond to non-persistent relationships is referred to as the *persistent network* corresponding to $N(G,R)$.

In a standard competitive equilibrium situation, workers are indifferent among employers offering the same working conditions and employers are indifferent among workers offering identical labor services. Moreover, workers offering the same labor services have the same ex ante expected employment rate and employers offering the same working conditions have the same ex ante expected vacancy rate.

In the current labor market model, these same market characteristics would tend to prevail if all workers and employers always cooperated. In the latter case, due to indifference, workers would randomly distribute their work offers across all employers and employers would randomly select work offers from among all work offers received. The resulting interaction pattern would therefore tend to be fully recurrent (no latching and no persistent non-employment) with equal ex ante expected employment rates and vacancy rates for workers and employers, respectively. For these reasons, the following interaction pattern among workers and employers is referred to below as a *competitive interaction pattern*: Each worker is recurrently directing work offers to employers, and every worker and employer has at least one persistent relationship.

The *network distance* for any persistent network $PN(G,R)$ is then defined to be the number of vertices (agents) in $PN(G,R)$ whose edges (persistent relationships) fail to conform to the competitive interaction pattern. By construction, then, a distance measure of 0 indicates zero deviation and a distance measure of 24 (the total number of workers and employers) indicates maximum deviation. In particular, a perfectly recurrent persistent network has a network distance of 0, a perfectly latched persistent network has a network distance of 12, and a perfectly disconnected persistent network (no persistent relationships) has a network distance of 24.

Classification of Work-Site Behaviors:

A worker or employer in generation G of a run R is called a *never-provoked defector (NPD)* if he ever defects against another agent that has not previously defected against him. The percentages of workers and employers who are NPDs measure the extent to which these agents behave opportunistically in work-site interactions with partners who are strangers or who so far have been consistently cooperative.

A worker or employer in generation G of a run R is referred to as a *persistent intermittent defector (IntD)* if he establishes at least one persistent relationship for which his persistent portion consists of a non-trivial mix of defections and cooperations. The agent is referred to as a *persistent defector (AllD)* if he establishes at least one persistent relationship and if the persistent portion of each of his persistent relationships consists entirely of defections. Finally, the agent is referred to as a *persistent cooperator (AllC)* if he establishes at least one persistent relationship and if the persistent portion of each of his persistent relationships consists entirely of cooperations. By construction, an agent in generation G of a run R satisfies one and only one of the following four agent-type classifications: persistently non-employed; a persistent intermittent defector; a persistent defector; or a persistent cooperator.

Two important points can be made about this classification of agent types. First, in contrast to standard game theory, the agents coevolve their types over time. This coevolution is in response to past experiences, starting from initially random behavioral specifications. Thus, agent typing is endogenous. Second, agent typing is measured in terms of persistently expressed behaviors, not in terms of work-site rules. An agent may have coevolved into an AllC in terms of expressed behaviors with current work-site partners, based on past work-site experiences with these partners, while still retaining the capability of defecting against a new untried partner. Indeed, work-site rules continually coevolve in the evolution step through mutation and recombination operations even if expressed behaviors appear to have largely stabilized. This ceaseless change in work-site rules makes any apparent stabilization in the distribution of agent types all the more surprising and interesting.

Measurement of Utility and Market Power Outcomes:

The utility level of a worker or employer at the end of generation G in a run R is measured by the average total net payoffs per trade cycle that the agent earns during the course of the trade cycle loop for generation G .

With regard to market power, we adopt the standard industrial organization approach: namely, market power is measured by the degree to which the actual utility levels attained by workers and employers compare against an idealized competitive yardstick. We take as this yardstick a situation in which there is absence of strategic behavior, symmetric treatment of equals, and full employment. Specifically, we define *competitive market conditions* for the ACE labor market to be a situation in which each worker is recurrently directing work offers to employers, and each worker and employer is a

persistent cooperator (AllC).

Ignoring offer costs, the utility level that each worker and employer would attain under these competitive market conditions is simply the mutual cooperation payoff level, C . Therefore, as in Pingle and Tesfatsion (2001,2002), we define the *market power (MPow)* of each worker or employer in generation G of a run R to be the extent to which their attained utility level, U , differs from C : that is, $MPow = (U-C)/C$.

Classification of Persistent Relationship Types:

A persistent relationship between a worker and employer in generation G of a run R is classified in accordance with the persistent behaviors expressed by the two participants in this particular relationship.

If both participants are persistent intermittent defectors (IntDs), the relationship is classified as *mutual intermittent defection (M-IntD)*. If both participants are persistent defectors (AllDs), the relationship is classified as *mutual defection (M-AllD)*. If both participants are persistent cooperators (AllCs), the relationship is classified as *mutual cooperation (M-AllC)*. Note that the relative shirking rates for an M-IntD relationship can be deduced for the participant worker and employer by examining their relative market power levels.

A persistent relationship in which the worker and employer express distinct types of behaviors is indicated in hyphenated form, with the worker's behavior indicated first. For example, a persistent relationship involving a worker who is an IntD and an employer who is an AllC is indicated by the expression IntD-AllC.

1.4 Experimental Findings

Overview:

The results for the computational experiment display a startling degree of regularity. This regularity is visible as early as the twelfth generation and persists through generation 1000.

For each of the three NEP treatments ZeroT, LowT, and HighT, the twenty trial runs tend to cluster into two distinct attractor states. Each attractor state supports a distinct configuration of market non-participation rates, work-site behaviors, utility levels, market power outcomes, and persistent relationship types. These attractor states can be Pareto-ranked, in the sense that the average utility levels attained by workers and by employers are both markedly higher in one of the two attractor states. The exact form of the attractor states varies systematically across the three NEP treatments.

Network Formation:

For each of the twenty runs corresponding to each treatment ZeroT, LowT, and HighT, the form of the persistent network was determined at three

sampling points: generation 12; generation 50; and generation 1000. Using the network distance measure defined in Section 1.3, the distribution of these persistent networks across runs was then plotted, conditional on treatment and sampled generation. Thus, a total of nine network distributions were plotted, three for each of the three treatments.

These nine network distributions are depicted in Figure 1.1. Network distance is measured along the horizontal axes and the number of runs clustered at this network distance is indicated on the vertical axes. Recall that a network distance of 0 corresponds to a perfectly recurrent (“competitive”) persistent network, a network distance of 12 corresponds to a perfectly latched persistent network, and a network distance of 24 corresponds to a perfectly disconnected persistent network (no persistent relationships).

In treatment ZeroT, perfectly latched networks are strongly dominant even by generation 12. For each sampled generation, all but one or two of the twenty runs exhibit persistent networks consisting of perfectly latched worker-employer pairs. This is indicated by the sharp peak in the network distribution at network distance 12.

In treatment LowT, perfectly latched networks are again dominant. Nevertheless, at each sampled generation, the network distribution is less sharply peaked at network distance 12 than it was for treatment ZeroT.

In treatment HighT, a new phenomenon arises. For each sampled generation, seven runs out of twenty lie at network distance 24, indicating that the workers and employers in these runs have failed to form any persistent relationships. At generation 12, the remaining 13 runs are scattered over network distances from 0 to 23. By generation 1000, however, the network distribution displays two sharp peaks, one at network distance 12 (latching) and one at network distance 24 (complete coordination failure).

Figure 1.1 also indicates the behavioral modes supported by each network distribution. For example, consider the 19 runs clustered at network distance 12 for the ZeroT treatment sampled at generation 50. Figure 1.1 indicates that workers and employers generally attained M-IntD (mutual intermittent defection) relationships in 11 of the runs and M-AllC (mutual cooperation) relationships in the remaining 8 runs.

Overall Utility and Efficiency Levels:

As indicated in Table 1.1, the following results are obtained for the average (per agent) utility achieved by workers and employers across treatments. For generations 12 and 50, average utility is highest in treatment LowT. By generation 1000, however, average utility is actually highest in treatment HighT. The latter finding results from the high NEP for inactive agents as well as the previously noted observation that workers and employers who do manage to match under treatment HighT become increasingly more successful over time at coordinating on persistent mutual cooperation (the first attractor state) rather than persistent non-employment (the second attractor state). Although this evolution over time of increased cooperative behavior between

matched workers and employers is observed under all three NEP treatments, it is observed most strongly for treatment HighT.

On the other hand, consider the overall level of economic efficiency attained under each NEP treatment, where efficiency takes into account both utility benefits and program costs. Specifically, for each NEP treatment, let net earnings be measured as the average utility level attained by workers and employers minus the average NEP paid to workers and employers. Ignoring offer costs, the maximum possible net earnings under each NEP treatment is the mutual cooperation payoff 40, attained by mutually cooperative workers and employers in latched pairings. For each NEP treatment, let *efficiency* be measured by actual net earnings as a percentage of 40. As indicated in Table 1.1, efficiency is substantially lower under treatment HighT than under treatments ZeroT or LowT. In particular, the relatively higher average utility level attained under treatment HighT at generation 1000 is more than offset by the higher average NEP to unemployed workers and vacant employers. In contrast, program costs are minimal at all sampled generations under treatments ZeroT and LowT since unemployment and vacancy rates remain close to zero. As Table 1.1 indicates, evaluated in terms of efficiency, the best program option overall turns out to be the low NEP.

Table 1.4 provides a breakdown of the average utility levels reported in Table 1.1 by agent type (workers and employers) and by attractor state for each of the three NEP treatments. The generally higher average utility levels attained by employers reflects in part the structural asymmetry that workers shoulder all of the offer (transaction) costs associated with network formation and maintenance. As detailed below, in some cases this structural asymmetry also appears to provide employers a strategic advantage in their work-site interactions.

Market Non-Participation Rates, Work-Site Behaviors, Utility Levels, and Market Power Levels:

Table 1.5 reports market non-participation rates, work-site behaviors, utility levels, and market power outcomes for the twenty runs constituting treatment ZeroT, each sampled at generation 12. These descriptive statistics are reported separately for each of the twenty individual runs comprising this treatment. More precisely, for each run, the following descriptive statistics are given: Persistent unemployment rate for workers (UnE-w); Persistent vacancy rate for employers (Vac-e); A count of never-provoked defectors for workers (NPD-w) and employers (NPD-e); A count of intermittent defectors for workers (IntD-w) and employers (IntD-e); A count of always-defectors for workers (AllD-w) and employers (AllD-e); A count of always-cooperators for workers (AllC-w) and employers (AllC-e); Mean utility level for workers (Util-w) with standard deviation (Util-w SD); Mean utility level for employers (Util-e) with standard deviation (Util-e SD); Mean market power level attained by workers (MPow-w); Mean market power level attained by employers (MPow-e).

The twenty runs in Table 1.5 are grouped together, first in accordance with their network distance (NetD), and second in accordance with the type of work-site behaviors expressed by the workers and employers. This grouping reveals that the runs are essentially clustered into two distinct *attractor states* comprising 18 runs in total, each run exhibiting a perfectly latched persistent network pattern (NetD=12). The remaining two runs {5, 7} comprise a mix of recurrent and latched relationships and appear to be transition states between the two attractor states. The workers attain very low mean market power levels in the transition-state runs. This is due to the substantial offer costs they accumulate from refused work offers in the course of maintaining their recurrent relationships.

In the first attractor state comprising four runs {30, 50, 90, 65}, 77% of the workers and 71% of the employers are AllCs (persistent cooperators). Despite the prevalence of AllC agent types, the employers attain an average mean market power level (MPow-e = -0.00) that is markedly higher than the corresponding level obtained by the workers (Mpow-w = -0.12). This is due to the offer costs incurred by workers in the process of forming and sustaining the persistent latched networks and to the modestly higher percentages of NPD (non-provoked defection), IntD (persistent intermittent defection), and AllD (persistent defection) exhibited by employers.

In the second attractor state comprising fourteen runs {60, 35, ..., 70}, very high percentages of the workers and the employers are NPDs and IntDs. Interestingly, the workers and employers obtain similar average mean market power levels in this second attractor state (-0.20 for workers and -0.17 for employers). However, these levels are substantially lower than the average mean market power levels they attain in the first attractor state. Thus, in terms of this market power measure, the first attractor state Pareto dominates the second attractor state.

In parallel to Table 1.5, Table 1.6 reports persistent relationship type counts for treatment ZeroT sampled at generation 12. As in Table 1.5, data are reported for the two transient-state runs {5, 7} plus the eighteen remaining runs grouped into the two attractor states.

The most striking aspect of Table 1.6 is the almost complete lack of mixed persistent relationships, i.e., relationships in which the participant worker and employer are expressing distinct types of behaviors. In particular, Table 1.6 reveals that the first attractor state comprising four runs {30, 50, 90, 65} is dominated by mutual cooperation (M-AllC) whereas the second attractor state comprising 14 runs {60, 35, ..., 70} is dominated by mutual intermittent defection (M-IntD). Mutual defection (M-AllD) is almost entirely absent.

The mean market power levels reported in Table 1.6 reveal, however, that the shirking rates expressed by the workers and employers in their M-IntD relationships in the second attractor state are not generally balanced in any given run. Rather, in about half the runs the workers shirk more than the employers, and in the remaining half the employers shirk more than the

workers. Thus, although the average mean market power levels attained by workers and employers in this second attractor state are very close, this hides an underlying volatility in relative shirking rates across runs.

The characteristics reported in Table 1.5 and Table 1.6 for treatment ZeroT sampled at generation 12 are largely maintained in generation 50 and in generation 1000. One interesting observation, however, is that individual runs can traverse from one attractor state to another as time proceeds. For example, run 30 is in the first attractor state in generation 12, appears as a transition state in generation 50, and ends up in the second attractor state by generation 1000. Conversely, run 60 is in the second attractor state in generation 12 but ends up in the first attractor state by generation 1000.

A second interesting observation is that the number of runs lying in each attractor state evens out over time. In generation 12, the cooperative first attractor state comprises only four runs while the second attractor state dominated by intermittent defection comprises fourteen runs. By generation 50, the first attractor state comprises eight runs while the second attractor state comprises 11 runs. By generation 1000, each attractor state comprises exactly nine runs. Thus, on average, agents over time are improving their ability to coordinate on mutual cooperation.

As in treatment ZeroT, the twenty runs comprising treatment LowT, sampled at generation 12, can be clustered into two attractor states together with a collection of transition states. The first attractor state comprises six runs characterized by perfect latching and a high percentage of AllC agent types in M-AllC relationships. The second attractor state comprises eight runs characterized by almost perfect latching and a high percentage of IntD agent types in M-IntD relationships. The six transition-state runs each comprise a mix of latched and recurrent relationships and have a high percentage of IntD agent types in M-IntD relationships.

In contrast to treatment ZeroT, however, the number of transition-state runs is larger (six runs instead of two) for treatment LowT sampled at generation 12. This is consistent with the network distribution data reported in Figure 1.1. The latter data reveal that, for each sampled generation, the peak at network distance 12 (latching) for treatment LowT is less pronounced than the peak at distance 12 for treatment zeroT. This indicates that the workers and employers in treatment LowT take longer on average to coordinate into perfect latched networks than the workers and employers in treatment ZeroT.

Also in contrast to treatment ZeroT, the average mean market power levels attained by workers and employers in treatment LowT, sampled at generation 12, are not balanced in the second attractor. The employers attain a level of -0.08, whereas the workers attain a markedly lower level of -0.16. The second attractor is dominated by latched relationships, indicating that each worker is persistently incurring only one offer cost per trade cycle. Since each offer cost is small relative to trade payoffs, only 1.0, it follows that accumulation of offer costs does not explain this large discrepancy in market

power. Rather, since the second attractor state is dominated by M-IntD relationships, this discrepancy indicates that the employers are managing to shirk at a substantially higher rate than the workers in these M-IntD relationships.

The outcomes for treatment LowT sampled at generation 1000 closely resemble the outcomes reported in Table 1.5 and Table 1.6 for treatment ZeroT sampled at generation 12. The first attractor state comprises nine runs strongly dominated by M-AllC relationships, and the second attractor state comprises seven runs strongly dominated by M-IntD relationships. (Hence, an increase in the size of the first attractor state is observed for treatment LowT in moving from generation 12 to generation 1000.) Mixed types of relationships are almost entirely absent in the two attractor states. In the first attractor state the workers and employers attain average mean market power levels of -0.04 and -0.03, respectively. In the second attractor state the workers and employers attain uniformly lower but balanced average mean market power levels of -0.15. As for treatment ZeroT, this balance hides an underlying volatility in shirking rates across runs.

Table 1.7 reports market non-participation rates, work-site behaviors, utility levels, and market power outcomes for the twenty runs constituting treatment HighT, sampled at generation 12. Table 1.8 reports persistent relationship type counts for these same runs, again sampled at generation 12. As for the previous two treatments, the twenty runs can be clustered into two attractor states together with a scattering of transition states. Moreover, once again the runs in the first attractor state exhibit perfectly (or almost perfectly) latched persistent networks with a high percentage of AllC agent types. Nevertheless, the nature of the second attractor state is dramatically different. Whereas in the previous two treatments the second attractor state was dominated by M-IntD relationships, now the second attractor state corresponds to complete or almost complete coordination failure. More precisely, the network distance for the runs in the second attractor state varies from 22 (only two persistent relationships) to 24 (no persistent relationships). With a high non-employment payment, agents are opting for non-employment rather than choosing to remain in M-IntD relationships.

As also seen for treatments ZeroT and LowT, increased coordination on the first attractor state occurs over time for treatment HighT. In generation 12, the first attractor state comprises five runs, the second attractor state comprises 9 runs, and the six remaining runs are scattered across transition states. Also, in the first attractor state, an average of 9.8 out of the 12 persistent relationships in each run are M-AllC. By generation 1000, however, the first attractor state comprises 11 runs, the second attractor state comprises seven runs, and only two runs are in a transition state. Moreover, in the first attractor state, an average of 10.73 out of the 12 persistent relationships in each run are M-AllC.

Summarizing the relative market power outcomes of workers and employ-

ers in each treatment, the following regularities are observed. For every treatment, in each sampled generation, the employers consistently attain a higher average market power level than workers in the cooperative first attractor state. This difference is attributable to the relatively higher (although small) incidence of NPD, IntD, and AllD behaviors among employers and to the fact that offer costs are borne solely by the worker. Also, for treatments ZeroT and HighT, the workers and employers attain essentially the same average market power levels in the second attractor state in each sampled generation; and the same is true for treatment LowT when sampled in generation 1000. A balanced market power level in the second attractor state indicates either that workers and employers have essentially the same shirking rates on average (treatments ZeroT and LowT) or that all agents are persistently non-employed (treatment HighT).

With regard to market power in the cooperative first attractor state compared across treatments, the workers attain a modestly negative average market power level in each treatment in each sampled generation; the levels range from -0.12 to -0.02. Interestingly, treatments LowT and HighT have a lower average incidence of NPD behavior and a higher average percentage of M-AllC relationships per run than treatment ZeroT in this first attractor state. Nevertheless, these advantages are offset (in market power terms) by the higher average offer costs incurred by workers due to the longer time taken within each generation to establish a persistent network. (For example, as seen in Table 1.7 for treatment HighT sampled at generation 12, only one run in the first attractor state attains a network distance of 12, i.e., a perfectly latched persistent network.) In contrast to the workers, employers do not incur offer costs, hence they attain close to a zero average market power level in each treatment at each sampled generation in the cooperative first attractor state; the levels range from -0.02 to +0.03.

With regard to market power in the second attractor state compared across treatments, in each sampled generation both the workers and the employers attain their lowest average levels in treatment HighT. The second attractor state in treatment HighT is characterized by complete or nearly complete coordination failure.

Never-Provoked Defection:

The importance of stance toward strangers and first impressions for determining subsequent outcomes in sequential interactions has been stressed by Orbell and Dawes (1993) and by Rabin and Schrag (1999). In the present computational experiment, two sharply differentiated attractor states exist for each treatment, the first dominated by persistent mutual cooperation and the second dominated either by persistent intermittent defection or by persistent non-employment. Thus, outcomes are strongly path dependent, and stance towards strangers and first impressions could play a critical role in determining these outcomes. These aspects of agent behavior are captured by counts of never-provoked defection (NPD).

In treatments ZeroT and LowT, NPD is commonly observed in all sampled generations, particularly in the second attractor state dominated by persistent intermittent defection (IntD). For example, as seen in Table 1.5, for treatment ZeroT sampled at generation 12, 33% of workers and 38% of employers engage in NPD in the first attractor state, and these percentages rise to 52% and 83%, respectively, for the second attractor state. It would appear that these high percentages for NPD in the second attractor state might actually be inducing the resulting predominance of IntD as agents engage in retaliatory defections. Because the non-employment payment is lower than the mutual defection payoff in these two treatments, agents tend to defect back against defecting partners rather than simply refusing to interact with them.

Another interesting observation regarding treatments ZeroT and LowT is that the incidence of NPD for each agent type in each attractor state tends to be higher in treatment ZeroT than in treatment LowT. In treatment ZeroT, the non-employment payment 0 lies below all work-site payoffs, including the sucker payoff $L=10$ earned by an agent who cooperates against a defecting partner. Consequently, there is no risk of refusal on the basis of bad behavior alone, but only from unfavorable comparisons with other agents. In contrast, in treatment LowT the non-employment payment 15 lies between the sucker payoff and the mutual defection payoff $D=20$. In this case, then, an opportunistic agent faces a higher risk of refusal since non-employment is preferred to a sucker payoff.

In treatment HighT the non-employment payment 30 lies above the mutual defection payoff for the first time, and the impact of this change in payoff configuration is substantial. For example, as reported in Table 5, only 13% of workers and 7% of employers in generation 12 engage in NPD in the first attractor state characterized by mutual cooperation. In contrast, 100% of workers and 97% of employers engage in NPD in the second attractor state characterized by complete or almost complete coordination failure. The same pattern holds at generation 50 and generation 1000. Agents are now much pickier with regard to their partners; an early defection from a partner drops that partner's expected utility assessment below the non-employment payoff and hence below minimum tolerability.

1.5 Concluding Remarks

As detailed by Roth (2002), recent advances in experimental methods and game theory using both human subjects and computational agents are now permitting economists to study a wide variety of complex phenomena associated with decentralized market economies. Examples include inductive price discovery, imperfect competition, buyer-seller matching, and the open-ended co-evolution of individual behaviors and economic institutions.

One interesting branch of this literature is the attempt to exploit syner-

gies between experiments with human subjects and experiments with computational agents by means of parallel experimental designs. The few parallel experimental studies to date have largely focused on financial market issues.⁹ However, we conjecture that parallel experiments will ultimately prove to be even more valuable when applied to economic processes such as labor markets in which face-to-face personal relationships play a potentially strong role in determining market outcomes.

The preliminary ACE labor market study at hand highlights the need to carefully align parallel experimental designs to ensure valid comparability. For example, transaction costs must be properly scaled across experiments to ensure comparable agent incentives, and horizons need to be aligned to ensure that short-run and long-run effects are properly distinguished.

In future ACE labor market studies, we intend to calibrate our parallel experimental designs to empirical data. Salient aspects of actual unemployment benefit programs will be incorporated, and findings from previous empirical studies of unemployment benefit programs will be used wherever possible. In addition, the recent construction of linked employer-employee (LEE) data sets is an exciting development facilitating the empirical study of outcomes generated by worker-employer interactions; see Hammermesh (1999). LEE data sets complement beautifully the focus of ACE labor market studies on worker-employer interaction patterns. Consequently, LEE data should permit careful empirical testing of computational findings related specifically to interaction effects, such as strong path dependence and the existence of multiple attractors.

Acknowledgments

The authors are grateful for helpful comments received from an anonymous referee, from Peter Orazem, from Marshall Van Alstyne and other participants in the UCLA Computational Social Sciences Conference held in May 2002 at Lake Arrowhead, California, and from Richard Freeman and other participants in the Harvard University Colloquium on Complexity and Social Networks held on April 7, 2003.

References

- D. Acemoglu and R. Shimer, Productivity Gains from Unemployment Insurance, *European Economic Review* 44 (2000), 1195-1223.
- F. D. Blau and L. M. Kahn, Institutional Differences in Male Wage Inequality: Institutions versus Market Forces, O. Ashenfelter and D. Card, editors, **Handbook of Labor Economics**, Volume 3A, Elsevier Science B.V., 1999, 1399-1461.

⁹See <http://www.econ.iastate.edu/tesfatsi/aexper.htm> for pointers to research using parallel experiments.

R. Freeman, War of the Models: Which Labour Market Institutions for the 21st Century?, *Labour Economics* 5 (1998), 1-24.

D. Gale and L. Lloyd Shapley, College Admissions and the Stability of Marriage, *American Mathematical Monthly* 69 (1962), 9-15.

D. Hammermesh, LEEping into the Future of Labor Economics: The Research Potential of Linking Employer and Employee Data, *Labour Economics* 6 (1999), 25-41.

L. Ljungqvist and T. J. Sargent, The Unemployment Dilemma, *Journal of Political Economy* 106 (1998), 514-550.

D. McFadzean, D. Stewart, and L. Tesfatsion, A Computational Laboratory for Evolutionary Trade Networks, *IEEE Transactions on Evolutionary Computation* 5 (2001), 546-560.

D. McFadzean and L. Tesfatsion, A C++ Platform for Evolutionary Trade Networks, *Computational Economics* 14 (1999), 108-134.

W. B. MacLeod and J. Malcomson, Motivation and Markets, *American Economic Review* 88 (1998), 388-411.

S. Nickell and R. Layard, Labor Market Institutions and Economic Performance, O. C. Ashenfelter and D. Card, editors, **Handbook of Labor Economics**, Volume 3C, Elsevier Science B.V., 1999, 3029–3084.

J. Orbell and R. M. Dawes (1993), Social Welfare, Cooperator's Advantage, and the Option of Not Playing the Game, *American Sociological Review* 58 (1993), 787–800.

M. Pingle and L. Tesfatsion, Non-Employment Benefits and the Evolution of Worker-Employer Cooperation: Experiments with Real and Computational Agents, ISU Economic Report No. 55, June 2001.

M. Pingle and L. Tesfatsion, Non-Employment Benefits and the Evolution of Worker-Employer Cooperation: A Human-Subject Experiment, Working Paper, October 2002.

M. Rabin and J. Shrag, First Impressions Matter: A Model of Confirmatory Bias, *Quarterly Journal of Economics* 114 (1999), 37-82.

A. Roth and M. A. Sotomayor, **Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis**, Cambridge University Press, Cambridge, UK, 1992.

A. Roth, The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics, *Econometrica* 70 (2002), 1341-1378.

L. Tesfatsion, A Trade Network Game with Endogenous Partner Selection. H. Amman, B. Rustem, and A. B. Whinston, editors. **Computational Approaches to Economic Problems**, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1997, 249–69.

L. Tesfatsion, Preferential Partner Selection in Evolutionary Labor Markets: A Study in Agent-Based Computational Economics, V. Porto, N. Saravanan, D. Waagan, and A. E. Eiben, editors, **Evolutionary Programming VII**, Springer-Verlag, Berlin, 1998, 15–24.

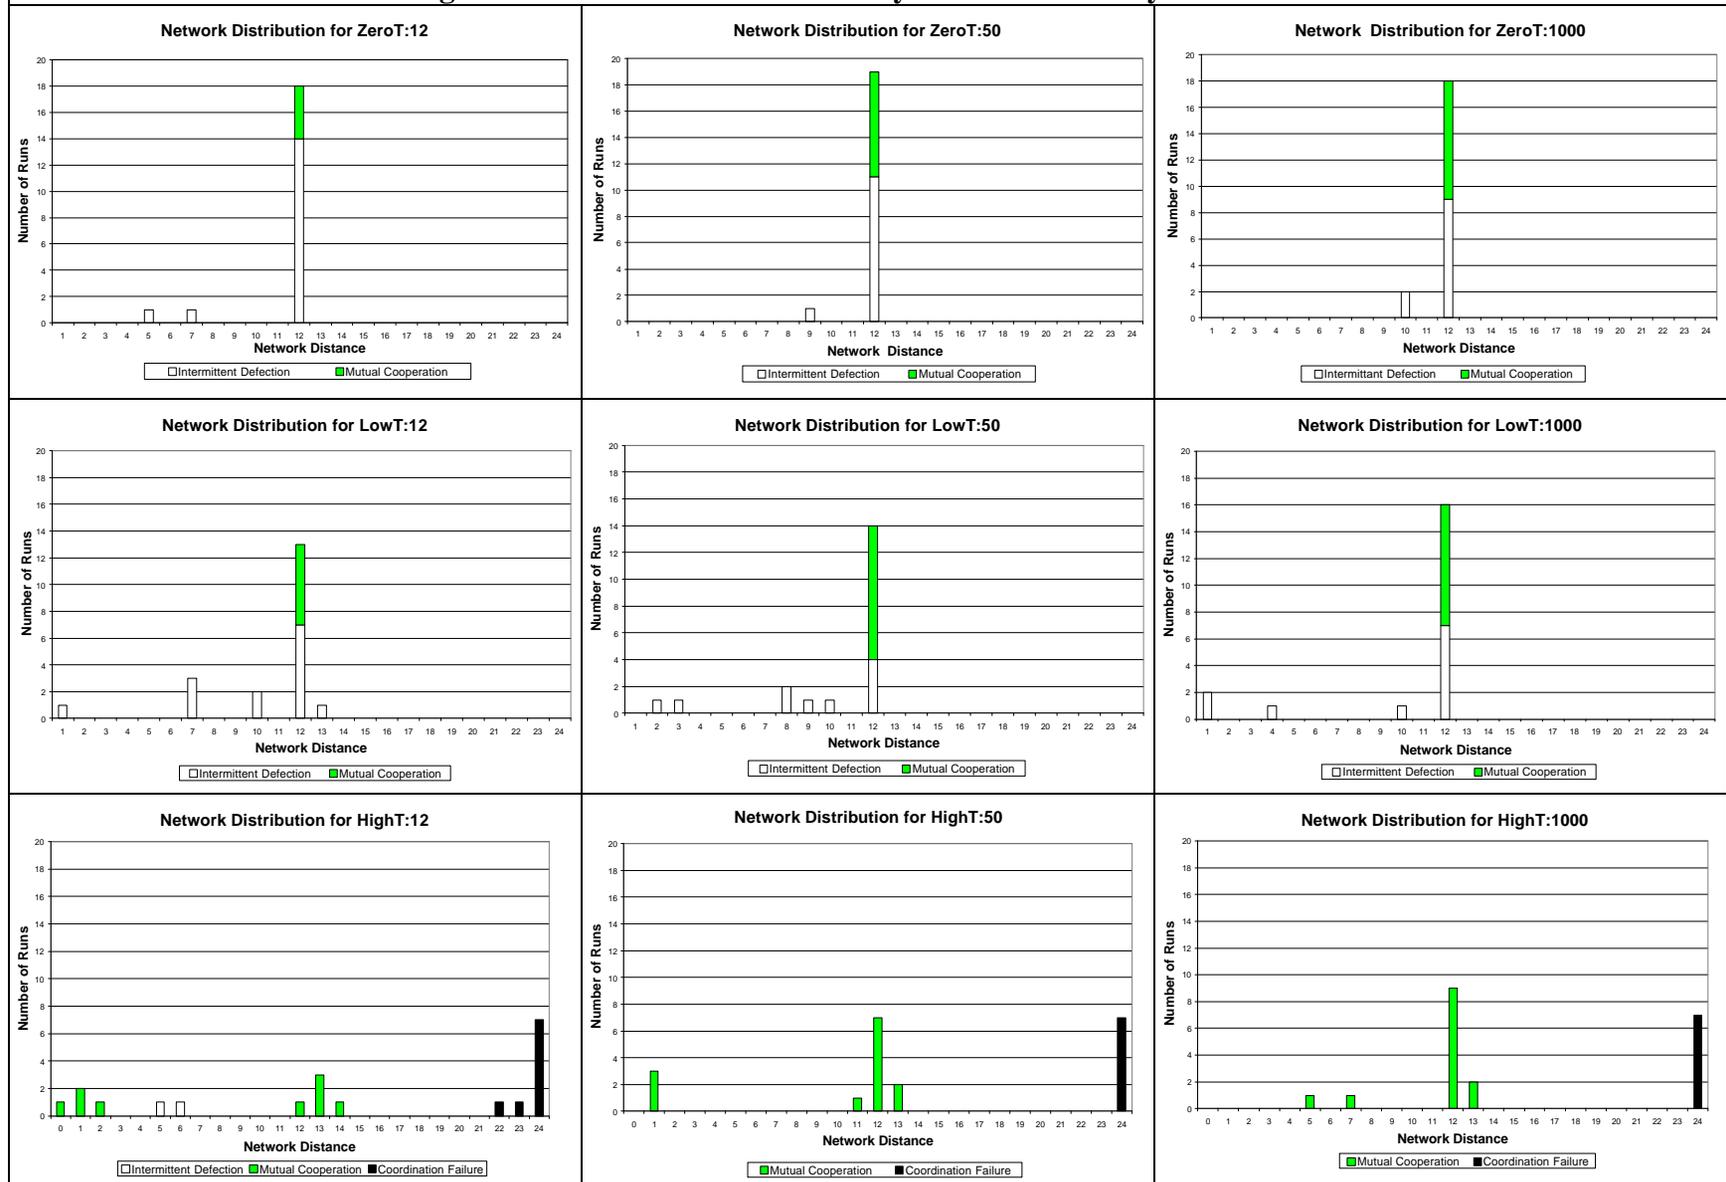
L. Tesfatsion, Structure, Behavior and Market power in an Evolutionary Labor Market with Adaptive Search, *Journal of Economic Dynamics and Control* 25 (2001), 419–457.

L. Tesfatsion, Hysteresis in an Evolutionary Labor Market with Adaptive Search, S.-H. Chen, editor, **Evolutionary Computation in Economics and Finance**, Physica-Verlag Heidelberg, 2002a, 189–210.

L. Tesfatsion, Economic Agents and Markets as Emergent Phenomena, *Proceedings of the National Academy of Sciences U.S.A.*, Volume 99, Supplement 3 (2002b), 7191–7192.

L. Tesfatsion, Agent-Based Computational Economics, Francesco Luna, Alessandro Perrone, and Pietro Terna, editors, **Agent-Based Theories, Languages, and Practices**, Routledge Publishers, 2003, to appear.

Figure 1.1: Network Distribution by Treatment and by Generation



	Generation 12	Generation 50	Generation 1000
ZeroT NEP=0	Utility 33.8 (5.9) Unemp 0% Vacancy 0% Efficiency 85%	Utility 35.2 (2.8) UnEmp 0% Vacancy 0% Efficiency 88%	Utility 35.8 (3.6) UnEmp 0% Vacancy 0% Efficiency 90%
LowT NEP=15	Utility 35.4 (3.6) UnEmp 1% Vacancy 2% Efficiency 88%	Utility 37.2 (2.0) UnEmp 0% Vacancy 1% Efficiency 93%	Utility 35.7 (1.8) UnEmp 2% Vacancy 3% Efficiency 88%
HighT NEP=30	Utility 34.6 (1.0) Unemp 48% Vacancy 50% Efficiency 50%	Utility 35.7 (0.5) Unemp 36% Vacancy 36% Efficiency 62%	Utility 36.4 (0.6) UnEmp 32% Vacancy 33% Efficiency 67%

Table 1.1: Summary of Aggregate Outcomes: Average agent utility level (with standard deviation), worker unemployment rate, employer vacancy rate, and efficiency level for each non-employment payment (NEP) treatment.

```

int main () {

    InitiateEconomy();           // CONSTRUCT initial subpopulations of
                                // workers and employers with random
                                // work-site rules of behavior.

    For (G = 1,...,1000) {      // ENTER THE GENERATION CYCLE LOOP

                                // GENERATION CYCLE:

        InitiateGen();         // Configure workers and employers
                                // with user-supplied parameter values
                                // (initial expected utility assessments,
                                // minimum tolerance levels,...)

        For (TC = 1,...,150) {  // ENTER THE TRADE CYCLE LOOP

                                // TRADE CYCLE:
                                // Workers and employers determine
                                // their work-site partners, given
                                // their expected utility assessments,
                                // and record job search and
                                // inactivity costs.

            MatchTraders();

            Trade();           // Workers and employers engage
                                // in work-site interactions and
                                // record their work-site payoffs.

            UpdateExp();       // Workers and employers update their
                                // expected utility assessments, using
                                // newly recorded costs and work-site
                                // payoffs, and begin a new trade cycle.

        }

        AssessFitness();      // ENVIRONMENT STEP:
                                // Each worker and employer assesses
                                // his utility (fitness) level.

        EvolveGen();          // EVOLUTION STEP:
                                // Workers and employers separately
                                // evolve their work-site rules,
                                // and a new generation cycle begins.

    }
    Return 0;
}

```

Table 1.2: Flow of Activities in the ACE Labor Market

		Employer	
		c	d
Worker	c	(40,40)	(10,60)
	d	(60,10)	(20,20)

Table 1.3: Payoff Matrix for the Work-Site Prisoner's Dilemma Game

	Generation 12			Generation 50			Generation 1000		
ZeroT NEP=0	CoopAtt.	w 35.4 (6.6)	e 40.1 (5.5)	CoopAtt.	w 39.1 (1.5)	e 40.7 (1.4)	CoopAtt.	w 36.4 (2.7)	e 39.9 (3.6)
	IntDAtt.	32.2 (5.9)	33.1 (6.0)	IntDAtt.	31.8 (2.7)	32.8 (2.8)	IntDAtt.	33.2 (4.1)	34.2 (3.2)
	Overall	32.5 (5.9)	35.1 (5.8)	Overall	34.0 (2.3)	36.3 (2.5)	Overall	34.7 (3.6)	36.8 (3.6)
	Efficiency=85%			Efficiency=88%			Efficiency=90%		
LowT NEP=15	CoopAtt.	w 38.3 (2.3)	e 40.1 (2.3)	CoopAtt.	w 38.9 (1.4)	e 40.8 (1.4)	CoopAtt.	w 38.5 (0.8)	e 38.8 (1.5)
	IntDAtt.	33.6 (4.3)	36.7 (4.4)	IntDAtt.	28.0 (2.3)	40.5 (3.1)	IntDAtt.	34.2 (2.1)	33.9 (2.2)
	Overall	34.0 (3.6)	36.8 (3.6)	Overall	35.7 (2.0)	38.6 (2.0)	Overall	36.0 (1.5)	35.3 (2.0)
	Efficiency=88%			Efficiency=93%			Efficiency=88%		
HighT NEP=30	CoopAtt.	w 38.1 (2.0)	e 40.2 (1.6)	CoopAtt.	w 38.9 (0.7)	e 40.2 (0.5)	CoopAtt.	w 38.9 (0.5)	e 41.3 (0.5)
	CFailAtt.	30.2 (0.4)	30.3 (0.3)	CFailAtt.	30.0 (0.0)	30.1 (0.0)	CFailAtt.	30.0 (0.0)	30.1 (0.0)
	Overall	33.5 (0.9)	35.7 (1.1)	Overall	35.5 (0.5)	35.9 (0.4)	Overall	35.5 (0.6)	37.2 (0.5)
	Efficiency=50%			Efficiency=62%			Efficiency=67%		

Table 1.4: Utility and Efficiency Outcomes: Average utility levels (with standard deviations) attained by workers and by employers in the two attractors and overall, together with the general efficiency level, for each non-employment payment (NEP) treatment.

Table 1.5: Non-Participation Rates, Work-Site Behaviors, and Welfare Outcomes---ZeroT Treatment, Generation 12

NetD	Run	NON-PARTICIPATION RATES AND WORK SITE BEHAVIORS										WELFARE OUTCOMES						NETWORK PATTERNS
		UnE-w	Vac-e	NPD-w	NPD-e	IntD-w	IntD-e	AllD-w	AllD-e	AllC-w	AllC-e	Util-w	Util-w SD	Util-e	Util-e SD	MPow-w	MPow-e	
5	40	0	0	0	8	9	7	0	2	3	3	32.5	3.5	37.6	4.1	-0.19	-0.06	Mix of latched pairs and recurrent relations
7	45	0	0	0	12	10	6	0	2	2	4	25.6	6.5	40.9	6.2	-0.36	0.02	
Average		0.00	0.00	0.00	10.00	9.50	6.50	0.00	2.00	2.50	3.50	29.1	5.0	39.3	5.2	-0.27	-0.02	
%		0%	0%	0%	83%	79%	54%	0%	17%	21%	29%							
12	30	0	0	0	11	2	1	0	4	10	7	28.2	9.4	45.1	6.1	-0.29	0.13	Perfect latched pairs with high percentage of AllC Cooperators.
12	50	0	0	2	3	5	6	0	0	7	6	34.2	5.9	42.0	2.3	-0.15	0.05	
12	90	0	0	12	4	2	3	0	0	10	9	36.2	4.5	37.2	3.7	-0.10	-0.07	
12	65	0	0	2	0	0	0	2	0	10	12	42.8	6.6	36.2	9.9	0.07	-0.10	
Average		0.00	0.00	4.00	4.50	2.25	2.50	0.50	1.00	9.25	8.50	35.4	6.6	40.1	5.5	-0.12	-0.00	
%		0%	0%	33%	38%	19%	21%	4%	8%	77%	71%							
12	60	0	0	7	12	11	11	0	1	1	0	28.1	5.7	39.6	5.3	-0.28	-0.01	Perfect latched pairs with high percentage of Int-Defectors
12	35	0	0	1	12	9	9	1	1	2	2	30.9	5.6	36.3	4.7	-0.23	-0.09	
12	95	0	0	5	12	10	9	0	1	2	2	30.4	5.3	35.6	2.8	-0.24	-0.11	
12	75	0	0	12	12	11	9	0	2	1	1	28.0	7.6	35.1	8.3	-0.30	-0.12	
12	15	0	0	6	11	10	9	1	2	1	1	30.9	6.6	35.1	6.5	-0.23	-0.12	
12	55	0	0	7	6	12	12	0	0	0	0	31.6	3.3	34.8	3.7	-0.21	-0.13	
12	85	0	0	6	8	11	12	0	0	1	0	31.5	3.6	34.7	4.6	-0.21	-0.13	
12	5	0	0	5	12	11	11	1	1	0	0	29.7	5.2	33.9	5.1	-0.26	-0.15	
12	10	0	0	1	10	7	8	2	0	3	4	36.6	7.0	33.0	8.2	-0.09	-0.18	
12	20	0	0	11	8	9	10	2	0	1	2	35.9	5.4	30.7	6.4	-0.10	-0.23	
12	80	0	0	5	8	10	10	1	0	1	2	37.4	7.2	30.5	6.5	-0.07	-0.24	
12	25	0	0	8	5	9	11	2	0	1	1	36.0	6.1	29.3	8.6	-0.10	-0.27	
12	0	0	0	2	11	8	8	4	3	0	1	31.5	8.2	29.0	6.4	-0.21	-0.28	
12	70	0	0	11	12	9	10	3	2	0	0	32.2	5.7	25.1	6.3	-0.20	-0.37	
Average		0.00	0.00	6.21	9.93	9.79	9.93	1.21	0.93	1.00	1.14	32.2	5.9	33.1	6.0	-0.20	-0.17	
%		0%	0%	52%	83%	82%	83%	10%	8%	8%	10%							
Total Average		0.00	0.00	5.15	8.85	8.25	8.10	0.95	1.05	2.80	2.85	32.5	5.9	35.1	5.8	-0.19	-0.12	
Total %		0%	0%	43%	74%	69%	68%	8%	9%	23%	24%							

Table 1.6: Persistent Relationship Type Counts---ZeroT Treatment, Generation 12

NetD	Run	MUTUALITY			MIXED CASES (w - e)						MARKET POWER		NETWORK PATTERNS
		M-IntD	M-AllD	M-AllC	IntD-AllD	IntD-AllC	AllD-IntD	AllD-AllC	AllC-IntD	AllC-AllD	MPow-w	MPow-e	
5	40	10	1	6	0	0	0	0	1	1	-0.19	-0.06	Mix of latched pairs and recurrent relations
7	45	7	1	3	2	2	0	2	1	0	-0.36	0.02	
Average		8.50	1.00	4.50	1.00	1.00	0.00	1.00	1.00	0.50	-0.27	-0.02	
12	30	0	0	7	2	0	0	0	1	2	-0.29	0.13	Perfect latched pairs with high percentage of AllC Cooperators
12	50	5	0	6	0	0	0	0	1	0	-0.15	0.05	
12	90	2	0	9	0	0	0	0	1	0	-0.10	-0.07	
12	65	0	0	10	0	0	0	2	0	0	0.07	-0.10	
Average		1.75	0.00	8.00	0.50	0.00	0.00	0.50	0.75	0.50	-0.12	-0.00	
12	60	10	0	0	1	0	0	0	1	0	-0.28	-0.01	Perfect latched pairs with high percentage of Int-Defectors
12	35	9	1	2	0	0	0	0	0	0	-0.23	-0.09	
12	95	9	2	0	1	0	0	0	0	0	-0.24	-0.11	
12	75	9	0	0	1	1	0	0	0	1	-0.30	-0.12	
12	15	9	1	0	0	1	0	0	0	1	-0.23	-0.12	
12	55	12	0	0	0	0	0	0	0	0	-0.21	-0.13	
12	85	11	0	0	0	0	0	0	1	0	-0.21	-0.13	
12	5	11	1	0	0	0	0	0	0	0	-0.26	-0.15	
12	10	7	0	3	0	0	1	1	0	0	-0.09	-0.18	
12	20	8	0	1	0	1	2	0	0	0	-0.10	-0.23	
12	80	10	0	1	0	0	0	1	0	0	-0.07	-0.24	
12	25	9	0	1	0	0	2	0	0	0	-0.10	-0.27	
12	0	8	3	0	0	0	0	1	0	0	-0.21	-0.28	
12	70	9	2	0	0	0	1	0	0	0	-0.20	-0.37	
Average		9.36	0.71	0.57	0.21	0.21	0.43	0.21	0.14	0.14	-0.20	-0.17	
Total Average		7.75	0.60	2.45	0.35	0.25	0.30	0.35	0.35	0.25	-0.19	-0.12	

Table 1.7: Non-Participation Rates, Work-Site Behaviors, and Welfare Outcomes---HighT Treatment, Generation 12

NetD	Run	NON-PARTICIPATION RATES AND WORK-SITE BEHAVIORS										WELFARE OUTCOMES						NETWORK PATTERNS
		UnE-w	Vac-e	NPD-w	NPD-e	IntD-w	IntD-e	AllD-w	AllD-e	AllC-w	AllC-e	Util-w	Util-w SD	Util-e	Util-e SD	MPow-w	MPow-e	
0	45	0	0	0	1	1	1	0	0	11	11	38.6	0.6	40.7	1.0	-0.04	0.02	Mostly recurrent relations. High percentage of AllC Cooperators.
1	0	0	1	0	1	0	0	0	0	12	11	35.4	0.4	40.5	1.7	-0.12	0.01	
1	65	0	1	0	1	0	0	0	0	12	11	35.4	0.4	40.5	1.5	-0.12	0.01	
2	35	0	2	0	2	0	0	0	0	12	10	33.2	0.7	39.9	2.5	-0.17	-0.00	
5	75	0	3	0	5	1	2	0	0	11	7	32.6	0.7	40.5	1.7	-0.19	0.01	
6	40	4	2	2	12	8	10	0	0	0	0	34.0	1.8	37.2	1.9	-0.15	-0.07	
Average		0.67	1.50	0.33	3.67	1.67	2.17	0.00	0.00	9.67	8.33	34.9	0.8	39.9	1.7	-0.13	-0.00	
%		6%	13%	3%	31%	14%	18%	0%	0%	81%	69%							
12	5	0	0	0	0	0	0	0	0	12	12	39.8	0.0	41.0	0.0	-0.01	0.03	Mostly latched relations. High percentage of AllC Cooperators.
13	10	1	1	1	0	0	0	0	0	11	11	39.1	2.3	40.2	1.6	-0.02	0.01	
13	50	1	1	1	0	0	0	0	0	11	11	39.0	2.7	40.2	1.6	-0.03	0.01	
13	80	1	1	1	0	0	0	0	0	11	11	39.0	2.7	40.2	2.1	-0.03	0.01	
14	85	3	4	5	4	3	5	0	0	6	3	33.3	2.5	39.6	2.7	-0.17	-0.01	
Average		1.20	1.40	1.60	0.80	0.60	1.00	0.00	0.00	10.20	9.60	38.1	2.0	40.2	1.6	-0.04	0.01	
%		10%	12%	13%	7%	5%	8%	0%	0%	85%	80%							
22	20	10	10	12	10	2	2	0	0	0	0	30.9	1.9	31.3	0.2	-0.23	-0.22	
23	15	11	11	12	11	1	1	0	0	0	0	30.5	1.6	30.8	0.1	-0.24	-0.23	
24	25	12	12	12	12	0	0	0	0	0	0	30.0	0.0	30.1	0.0	-0.25	-0.25	
24	30	12	12	12	12	0	0	0	0	0	0	30.0	0.0	30.1	0.0	-0.25	-0.25	
24	55	12	12	12	12	0	0	0	0	0	0	30.0	0.0	30.1	0.0	-0.25	-0.25	
24	60	12	12	12	12	0	0	0	0	0	0	30.0	0.0	30.1	0.0	-0.25	-0.25	
24	70	12	12	12	12	0	0	0	0	0	0	30.0	0.0	30.1	0.0	-0.25	-0.25	
24	90	12	12	12	12	0	0	0	0	0	0	30.0	0.0	30.1	0.0	-0.25	-0.25	
24	95	12	12	12	12	0	0	0	0	0	0	30.0	0.0	30.1	0.0	-0.25	-0.25	
Average		11.67	11.67	12.00	11.67	0.33	0.33	0.00	0.00	0.00	0.00	30.2	0.4	30.3	0.0	-0.25	-0.24	
%		97%	97%	100%	97%	3%	3%	0%	0%	0%	0%							
Total Average		5.75	6.05	5.90	6.55	0.80	1.05	0.00	0.00	5.45	4.90	33.5	0.9	35.7	0.9	-0.16	-0.11	
Total %		48%	50%	49%	55%	7%	9%	0%	0%	45%	41%							

Table 1.8: Persistent Relationship Type Counts---HighT Treatment, Generation 12

NetD	Run	MUTUALITY			MIXED CASES (w-e)						MARKET POWER		PATTERNS
		M-IntD	M-AllD	M-AllC	IntD-AllD	IntD-AllC	AllD-IntD	AllD-AllC	AllC-IntD	AllC-AllD	MPow-w	MPow-e	
0	45	1	0	37	0	0	0	0	0	0	-0.04	0.02	Mostly recurrent Relations. High percentage of AllC Cooperators.
1	0	0	0	79	0	0	0	0	0	0	-0.12	0.01	
1	65	0	0	80	0	0	0	0	0	0	-0.12	0.01	
2	35	0	0	99	0	0	0	0	0	0	-0.17	-0.00	
5	75	1	0	60	0	0	0	0	1	0	-0.19	0.01	
6	40	21	0	0	0	0	0	0	0	0	-0.15	-0.07	
Average		3.83	0.00	59.17	0.00	0.00	0.00	0.00	0.17	0.00	-0.13	-0.00	
12	5	0	0	12	0	0	0	0	0	0	-0.01	0.03	Mostly latched relations. High percentage of AllC Cooperators.
13	10	0	0	11	0	0	0	0	0	0	-0.02	0.01	
13	50	0	0	11	0	0	0	0	0	0	-0.03	0.01	
13	80	0	0	11	0	0	0	0	0	0	-0.03	0.01	
14	85	3	0	4	0	0	0	0	2	0	-0.17	-0.01	
Average		0.60	0.00	9.80	0.00	0.00	0.00	0.00	0.40	0.00	-0.04	0.01	
22	20	2	0	0	0	0	0	0	0	0	-0.23	-0.22	Almost complete coordination failure.
23	15	1	0	0	0	0	0	0	0	0	-0.24	-0.23	
24	25	0	0	0	0	0	0	0	0	0	-0.25	-0.25	
24	30	0	0	0	0	0	0	0	0	0	-0.25	-0.25	
24	55	0	0	0	0	0	0	0	0	0	-0.25	-0.25	
24	60	0	0	0	0	0	0	0	0	0	-0.25	-0.25	
24	70	0	0	0	0	0	0	0	0	0	-0.25	-0.25	
24	90	0	0	0	0	0	0	0	0	0	-0.25	-0.25	
24	95	0	0	0	0	0	0	0	0	0	-0.25	-0.25	
Average		0.33	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.25	-0.24	
Total Average		1.45	0.00	20.20	0.00	0.00	0.00	0.00	0.15	0.00	-0.16	-0.11	