Non-Employment Benefits and the Evolution of Worker-Employer Cooperation: Experiments with Real and Computational Agents

Mark Pingle
University of Nevada at Reno

Leigh Tesfatsion
Iowa State University

ISU Economic Report Series No. 55
June 2001

ISU Economic Report No. 55

Non-Employment Benefits and the Evolution of Worker-Employer Cooperation: Experiments with Real and Computational Agents

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

and

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

12 June 2001

Abstract

Experiments with real and computational agents are used to examine the impact of changing the level of a *non-employment payoff* on the evolution of cooperation between workers and employers participating in a sequential employment game with incomplete contracts. Workers either direct work offers to preferred employers or choose unemployment and receive the non-employment payoff. Subject to capacity limitations, employers either accept work offers from preferred workers or remain vacant and receive the non-employment payoff. Matched workers and employers participate in an employment relationship modeled as a prisoner's dilemma game. In both types of experiments, increases in the non-employment payoff result in higher unemployment and vacancy rates while at the same time encouraging higher rates of cooperation among the workers and employers who do form matches. However, the behaviors exhibited by the computational agents are coordinated to a higher degree than the behaviors of the human subjects. This difference raises challenging questions for both human-subject and computational experimentalists.

Keywords: Labor market; unemployment rate; vacancy rate; evolution of cooperation; efficiency wage; search and matching; endogenous interaction networks; evolutionary game; human-subject experiments; computational experiments; agent-based computational economics.

1. Introduction

1.1. Overview

It is now commonly understood that the complexity of most employment relationships forces the typical employment contract to be incomplete. If the contract does not enforce the desired level of cooperation, it is reasonable to think other institutions might arise to do the job. Using experiments with both real and computational agents, we examine the possibility that the level of non-employment benefits affects the level of cooperation between workers and employers, thereby impacting the unemployment rate, the productivity of labor, and a variety of other economic outcomes.

Our framework is consistent with that of MacLeod and Malcomson (1998). Both workers and employers can shirk. Intuitively, we can think of a worker as offering hard work in exchange for a bonus, and an employer as offering a bonus in exchange for hard work. We model the employment relationship between a worker and employer as an iterated prisoner's dilemma game. Axelrod (1984) shows that cooperation can evolve in this environment if the perceived probability of future interactions with any given current partner is sufficiently high, and if agents do not discount this "shadow of the future" too heavily.

A distinctive feature of our employment game relative to previous studies, however, is that matches between workers and employers are not determined exogenously by random selection or some other extraneous device. Rather, workers pro-actively direct work offers to employers, which the employers can either accept or reject. Over successive *trade cycles* (work periods), workers form assessments regarding which employers are more preferable to make offers to, and employers form assessments regarding which workers are more preferable to accept offers from. At any time, workers and employers can also choose to stop participating in the labor market and instead collect the non-employment payoff.

Because we view this analysis as a baseline, we consider a labor market with a balanced structure. There are equal numbers of workers and employers. In each trade cycle, each worker has one work offer and each employer has one job opening. Any worker or employer who does not enter an employment relationship receives the non-employment payoff. Workers and employers make their shirking choices simultaneously so that neither has a strategic informational advantage. One assumption that does bias the market slightly in favor of employers, however, is that a worker must pay a small offer cost (search cost) each time that he makes a work offer.

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

A major difference between experiments with human subjects and experiments with computational agents is that learning processes must be fully articulated for computational agents. Our computational workers and employers are self-interested expected utility maximizers. In analogy to human subjects, each worker and employer is permitted to evolve his work-site strategy over time. This evolution involves both inductive learning (experimentation with new strategies) and social learning (mimicking the strategies of more successful agents of one's own type).

Because human subjects may not be expected utility maximizers and may select their strategies in manners quite different from what is programmed into our computational agents, we should not expect the results from our computational experiment to mirror in precise form the results of our human-subject experiment. Nevertheless, as stressed by Duffy (2001), we can determine in principle the source of any observed regularity arising in the computational experiment. Thus, both similarities and differences in results can be informative.

Two important findings are obtained from both the computational and human-subject experiments. First, an increase in the non-employment payoff increases both the unemployment rate and the vacancy rate while at the same time encouraging cooperation among the workers and employers who manage to match. In particular, a high non-employment payment channels agents either towards productive cooperative behavior or towards inactivity. Second, in the short and intermediate run, a low non-employment payoff appears preferable to either a zero non-employment payoff or a high non-employment payoff. A zero non-employment payoff encourages too much shirking, while a high non-employment payoff results in too high a risk of lost production due to coordination failure.

The results of the two experiments also differ in interesting ways. For example, the outcomes for the computational experiment display a much more stable and structured response to changes in the non-employment payoff than the outcomes for the human-subject experiment. In each treatment, outcomes in the computational experiment typically gravitate towards one of two "attractor states." In the treatments with zero and low non-employment payoffs, the two attractor states are similar: the first is characterized by latched pairs of mutually cooperative workers and employers, while the second is characterized by latched pairs of workers and employers who intermittently defect and cooperate. Alternatively, when the non-employment payoff is high, the first attractor state is characterized by latched pairs of mutually cooperative workers and employers while the second attractor state is a state of economic collapse in which each worker and employer ultimately becomes inactive. In contrast, in the human-subject experiment, most relationships that form between workers and employers are short-lived, ending when one or both of the participants shirk.

This latter finding 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 our relatively short human-subject trials, biasing behaviors towards unknown past points of reference? If so, the computational experiment may indeed provide an accurate prediction of what would happen in the human-subject experiment over a longer time span. Alternatively, the two experiments may differ structurally in some fundamental way so that differences would be observed regardless of time scale. In particular, we might ask, “Is the representation of agent learning in the computational experiment too inaccurate to permit valid comparisons with human subjects?” These and other challenging issues are addressed in the final sections of this study.

The design of the human-subject experiment is presented in Section 2, and results for this experiment are reported in Section 3. Section 4 outlines the design of the computational experiment, and Section 5 reports the results for this experiment. The results for the two experiments are compared and contrasted in Section 6. Section 7 contains concluding remarks.

1.2. Related Literature

Because no worker or employer knows for sure in our employment game how any potential partner will behave, the employment contract is incomplete. Mutually cooperative employment relationships yield rents relative to employment relationships of mutual defection and relative to the non-employment payoff. However, shirking is preferable to not shirking for each individual in any single trade cycle, which is particularly apparent in the final trade cycle. If it is common knowledge that all are self-interested, the incentive to cooperate recursively breaks down, beginning with the final trade cycle and ending with the first. The subgame perfect equilibrium is therefore a state where, in all trade cycles, all those in employment relationships shirk. When the non-employment payoff is higher than the mutual defection payoff, the unemployment and job-vacancy rates in this equilibrium are 100% and all agents earn the non-employment payoff. Conversely, when the non-employment payoff is below the mutual defection payoff, the unemployment and job-vacancy rates are both 0% and all agents earn the mutual defection payoff.

To the extent cooperation can evolve, agents can improve their welfare relative to the common knowledge Nash equilibrium where all are self interested. Kreps et al. (1982) have shown that cooperative play can be supported in a finitely repeated prisoner's dilemma if an agent believes there is at least a small probability that the other may be cooperative; e.g., a Tit-for-Tat player. The rent that can be earned is an external incentive that motivates the cooperation, rationally leading the agent to expect the other to be cooperative. The seminal work of Shapiro and Stiglitz (1984) uses this external incentive to explain involuntary unemployment and efficiency wages: Efficiency wages arise so as to create the unemployment that must develop so that employment yields the rent necessary to deter shirking. Extending this work, Albrecht and Vroman (1998) show that shirking can be an equilibrium phenomenon and multiple wage levels can result when workers are heterogeneous with respect to how much disutility they associate with work. Carmichael and MacLeod (1997) show that a gift-giving custom or an equivalent job search cost can promote cooperation because they increase the rent associated with a long-term employment relationship.

As an alternative to rent seeking, Rabin (1993) has shown that a preference for "fair" outcomes can motivate cooperation. Here, an internal incentive motivates the cooperation, associated with the belief that defecting (shirking) is unfair. Fehr, Kirchsteiger, and Riedl (1993) adopt this approach and present a "fair-wage effort" hypothesis as an alternative to the shirking version of the efficiency-wage hypothesis of Shapiro and Stiglitz. In their experiment, they found support for the fair-wage effort hypothesis in that it was common for a subject to behave cooperatively even when there was no rent to seek. To date, numerous experiments have demonstrated that both fairness concerns and self-interest motivate people. Fehr and Gachter (1998) review a number of studies. Goeree and Holt (2000), Bolton and Ockenfels (2000), and Fehr and Schmidt (1999) each demonstrate that adding inequity aversion to a self-interest model can help explain human behavior in a variety of experimental settings.

In an experiment comparable to ours, Falk, Gachter, and Kovacs (1999) find that human subjects exhibit fairness concerns and this increases worker effort levels. They also find that, compared to a one-shot interaction, repeated interaction with the same employer enhances worker effort. While our framework allows for repetition, it also introduces the realistic feature that any long-term employment relationship between a worker and an employer must arise endogenously. A priori, then, it is not clear whether the long-term relationships that can facilitate cooperation will arise. We focus on whether or not a change in the non-employment payoff can affect the evolution of cooperation by affecting the willingness of employers and workers to enter into cooperative employment relationships.

2. The Human-Subject Experiment

The human-subject experiment was implemented by having human subjects play an *employment game*.

2.1. Subjects

The experiment contained three treatments. Five sessions were run for each treatment, and each session included six new (i.e., inexperienced) subjects. Thus, a total of 90 human subjects participated. All subjects were University of Nevada student volunteers. The time of play for a session ranged from 45 minutes to 70 minutes. Subjects were paid based upon their performance. Subject earnings averaged $14.40, ranging from $5 to $26.

Subjects completed a demographics questionnaire so we could determine whether particular behaviors or outcomes were correlated with particular demographic characteristics. The 90 subjects ranged in age from 18 to 63, with the mean age being 24.6. They were 57 percent male and 43 percent female. The racial background was 74.4 percent Caucasian, 6.6 percent Hispanic, and 18.8 percent Asian. Marital status was 70.0 percent single, 16.6 percent married, 5.5 percent divorced, and 7.8 percent living with someone. Forty-three percent were receiving financial aid. The number of hours worked per week ranged from 0 to 60, with a mean of 18.3 hours. Work experience ranged from zero to 47 years, with a mean of 7.7 years. The number of college credits taken ranged from three to 160, with a mean of 67.7 credits. The GPA on college credits of the participants ranged from 2.0 to 4.0, with a mean GPA of 3.3.

2.2. Apparatus

The employment game was constructed using network programming languages. Worker and employer subjects communicated through a series of interconnected web pages. This allowed for the quick transfer of labor market information between employers and workers, ease in recording the labor market information, and control over the process through which workers and employers interacted. Worker subjects were assigned the identities Worker1, Worker2, and Worker3, while employer subjects were assigned the identities Employer1, Employer2, and Employer3. The six subjects participating in a session were randomly assigned to these six roles.

All of the subjects were in the same computer lab, stationed at six different computers. Subjects could see each other. However, the computer terminals were located far enough apart that the only information a participant could obtain was from his or her own computer terminal. Subjects were forbidden from revealing their individual employer or worker identities. Verbal communication with other participants during the session was also forbidden. Subjects were asked not to make any verbal comments or gestures during the game that would give away their identity. Subjects could only communicate with the administrator and could only do so privately.

2.3. Procedures

After being randomly assigned a role, participating subjects read a detailed set of game instructions. After any questions were privately answered, the employment game was played. Subjects then completed a demographics questionnaire and were paid for their performance.

The employment game consisted of twelve *trade cycles*. Each trade cycle provided an opportunity for workers and employers to obtain earnings. A trade cycle consisted of two parts: A matching process and an employment process. The matching process determined whether or not a particular worker obtained a job and whether or not a particular employer filled a job opening. The employment process determined the earnings obtained by employed workers and active employers.

The matching process used was a variant of the Gale-Shapley (1962) matching process. Each worker had one work offer that could become employed. Each employer had one job opening that could be filled. In the first *matching round*, each worker either submitted their offer to one of the employers or selected the *no offer* option. Whenever the no offer option was selected, the worker exited the matching process.

The offers submitted by workers accumulated on the *work offer lists* of the employers. Because there were three workers, an individual employer could receive 0, 1, 2, or 3 work offers. An employer not receiving any offers sat out during the given round. Employers receiving offers had to decide whether to accept or reject them. Because each employer had only one opening, any employer receiving more than one offer had to reject at least one offer.

Whenever an offer was rejected, a new matching round commenced. In the new round, each worker who had his or her offer rejected had to either redirect the offer to another employer (or back to the same employer) or exit the matching process by selecting the no offer option. An employer who received a redirected offer then had to either accept or reject it. To accept, the employer either had to have a vacancy or had to displace an existing offer. When an offer was displaced, the worker originally submitting the offer was informed and had to decide what to do with the offer in the next matching round. Additional matching rounds commenced until each work offer was either accepted or redirected to the no offer option. That is, the matching process ended when no new offers were made. Thus, the number of matching rounds occurring during a trade cycle was endogenously determined.

At the end of the matching process, workers and employers who were not matched received the *non-employment payoff*. The job openings of any unmatched employers where labeled *vacant* for the given trade cycle. The non-employment payoff received by an employer can therefore be interpreted as an employer subsidy paid to mitigate the hardship associated with not obtaining earnings from an employment relationship. An employer receiving the non-employment payoff was labeled *inactive* if the employer never received a work offer and *discouraged* if work offers were received but rejected. A worker receiving the non-employment payoff was labeled *inactive* if the no offer option was chosen in the first round. Because the inactive worker's offer is never actually made to any employer, the inactive worker does participate in the labor force. A worker receiving the non-employment payoff was labeled *unemployed* if the no offer option was chosen in a round subsequent to the first trade cycle. The unemployed worker actively seeks work but does not obtain a job in the trade cycle. For the inactive worker, the non-employment payment can be thought of as a welfare payment. For the unemployed worker, the non-employment payment can be thought of as unemployment compensation. In each case this payment to the worker mitigates the hardship associated with not obtaining earnings from an employment relationship.

Each matched worker-employer pair entered into an employment relationship. The amount of earnings generated during the trade cycle by the relationship and the distribution of the earnings between the employer and worker depended upon the type of relationship that formed. Both employer and worker had a simple choice to make, “To shirk or not to shirk.” If neither employer nor worker shirked, 80 points were generated and both employer and worker received 40. If both employer and worker shirked, 40 points of earnings were generated and both employer and worker received 20. Under this payment scheme, the *cooperation* associated with not shirking is more productive than the *defection* associated with shirking. For both employer and worker, the motivation to not shirk stemmed from the fact that cooperation led to higher productivity and consequently higher income. The motivation to shirk was introduced by assuming that shirking by only one of the two parties skews the distribution of the earnings toward the shirker. Specifically, when one shirked and one did not, the earnings generated was 70 points, of which the shirker received 60 and non-shirker received 10. In summary, the employment relationship entered by a matched employer-worker pair was a prisoner’s dilemma. Cooperation, in the form of non-shirking behavior, offered potential rewards, but was also risky relative to defection in the form of shirking.

2.4. Design

For self-interested subjects, it is clear that the values of 10 and 20 for the non-employment payoff are critical. If the non-employment payoff were less than 10, self-interested workers and employers would always prefer participating in an employment relationship to being inactive. As long as the non-employment payoff is less than 20, an employment relationship where worker and employer shirk yields a higher payoff to worker and employer than inactivity, and attempting to enter an employment relationship is not risky relative to inactivity unless the subject chooses not to shirk. Alternatively, if the non-employment payoff exceeds 20, attempting to enter an employment relationship is necessarily risky relative to choosing inactivity, for 20 or less is earned when the other shirks. We have then three qualitatively different situations. To examine these three situations, we set the non-employment payoff equal to 0 for our *ZeroT* treatment, 15 for our *LowT* treatment, and 30 for our *HighT* treatment. Table 1 presents a summary of our experimental design.

-- INSERT TABLE 1 ABOUT HERE --

Given this design, predicting subject behavior is not easy. Whether or not the non-employment payoff affects the willingness of subjects to enter into employment relationships or affects the willingness of subjects in employment relationships to cooperate depends upon the unobservable beliefs and strategies used by the subjects.

3. Results from the Human-Subject Experiment

We are interested in how the three non-employment payoff treatments affected the outcomes experienced by the human subjects. However, outcomes depend upon subject behavior, and subject behavior may depend upon the types of trading networks that form. Therefore, we begin by presenting results on network formation. We then present results on subject behavior and outcomes.

3.1. Network Formation

We characterized network formation in the human subject experiment by examining the last four of the twelve trade cycles in each session. The relationship between a given worker and given employer in a session was classified as latched, recurrent, or transient. If the worker and employer entered into an employment relationship with each other in the last two trade cycles, the relationship was classified as *latched*. If the worker and employer did not enter into an employment relationship with one another in any of the last 4 trade cycles, the relationship was classified as *transient*. If a relationship was neither latched nor transient, it was classified as *recurrent*.

Because there were three workers and three employers in each session, there were nine worker-employer relationships to classify per session. Thus, over the five sessions in a treatment, there were 45 relationships to classify. Table 2 shows the percentage of the 45 relationships that fall into each of the three categories.

-- INSERT TABLE 2 ABOUT HERE --

Given our definitions, at most 33 percent of all potential relationships can actually be latched, for if a worker becomes latched to one employer, he must be recurrent or transient with the other two employers. We found that relationships were more likely to be transient when the non-employment payoff was higher. Below, we will see this is primarily because a higher percentage of subjects moved toward inactivity as the non-employment payoff was increased. The relationship between the level of the non-employment payoff and the prevalence of latched relationships was not monotonic. Latched relationships were least prevalent in the HighT treatment and most prevalent in the LowT treatment. Recurrent relationships were most prevalent in the ZeroT treatment. In relative terms, then, we found that the HighT non-employment payoff level promoted a transient trading network, the LowT level promoted a latched network, and the ZeroT level promoted a recurrent network.

In the HighT treatment, network formation varied very little across the 5 sessions. Two-thirds of the relationships were transient in four of the five sessions, and the fifth session was nearly the same. In the ZeroT treatment, four of the five sessions were comparable in that there were no latched relationships and 78 to 89 percent of relationships were recurrent. One session in the ZeroT treatment was quite different in that 33 percent of the relationships were latched. None of these latched relationships was mutually cooperative, however. Network formation in the LowT treatment varied the most. In one session, 22 percent of the relationships were latched and 11 percent were recurrent, while in another no relationships were latched and 89 percent were recurrent. Except for the latter session, what was consistent in the LowT treatment was that at least 11 percent of the relationships in each session were latched and mutually cooperative.

Another way to examine relationship formation in the human-subject experiment is to examine the number of employment relationships that lasted at least two consecutive periods. There were 30 of these relationships in the ZeroT treatment, 26 in the LowT treatment, and 17 in the HighT treatment. Of these relationships, 17 percent were mutually cooperative in the ZeroT treatment, 23 percent in the LowT treatment, and 18 percent in the HighT treatment. Alternatively, 29 percent were relationships of mutual defection in the ZeroT treatment, 15 percent in the LowT treatment, and 30 percent in the HighT treatment. The mean relationship length was 2.56 in the HighT treatment, 3.00 in the LowT treatment, and 2.89 in the ZeroT treatment. If we define a long-term relationship as one lasting three trade cycles or more, 18 percent were long-term in the HighT treatment, 31 percent in the LowT treatment, and 20 percent in the ZeroT treatment.

Summarizing, we found that increasing the non-employment payoff both decreased the rate of relationship formation and increased the likelihood that a relationship would be transient. However, the impact of the non-employment payoff was not monotonic regarding the types of relationships that formed. The LowT treatment had the highest incidence of latched relationships, long term relationships, and mutually cooperative relationships.

3.2. Subject Behavior

In any given trade cycle, we classified the behavior of human subjects who entered employment relationships into four mutually exclusive categories:

1. Unprovoked Defection: Shirking at the first meeting or shirking when the other did not shirk at the last meeting.

2. Reciprocal Defection: Shirking when the other shirked at the last meeting.

3. Unprovoked Cooperation: Not shirking at the first meeting or not shirking when the other shirked at the last meeting.

4. Reciprocal Cooperation: Not shirking when the other did not shirk at the last meeting.

The subjects in this trade cycle who did not enter any employment relationships were classified into three additional mutually exclusive categories.

5. Unemployed Worker: A worker not entering into an employment relationship after having pursued one.

6. Discouraged Employer: An employer not entering into an employment relationship even though at least one work offer was received.

7. Inactivity: For workers, never making a worker offer. For employers, never receiving a work offer.

If cooperation is to evolve, reciprocal cooperation must not be extinguished. Rather, it must become the norm. For the three treatments, Figure 1 shows the mean rate of reciprocal cooperation as it evolved over the twelve trade cycles. For each treatment, the mean was calculated over the 30 subjects (15 employers and 15 workers) in all five sessions. We combined employers and workers into one group because we found no statistically significant difference in the shirking tendencies of workers and employers.

-- INSERT FIGURE 1 ABOUT HERE --

By definition, reciprocal cooperation cannot occur in the first trade cycle. By the third trade cycle, the rate of reciprocal cooperation was clearly higher in the ZeroT and LowT treatments than in the HighT treatment. The rate of reciprocal cooperation was comparable in the ZeroT and LowT treatments until the final trade cycle was approached. At that point, subjects in the ZeroT treatment moved away from reciprocal cooperation. However, the rate of reciprocal cooperation persisted in the LowT treatment.

Multinomial logistic regression is a useful tool for determining what factors influence the prevalence of any one of the seven types of behavior relative to any other. This regression method requires that one of the behavioral categories be selected as a reference group. Because we are interested in examining the evolution of cooperation, we selected reciprocal cooperation as our reference behavior. Table 3 presents the *relative risk ratios* associated with pertinent independent variables.

-- INSERT TABLE 3 ABOUT HERE --

Regression 1 informs us of the prevalence of the various types of behavior relative to reciprocal cooperation. A positive number indicates the particular behavior was more prevalent than reciprocal cooperation. Thus, we find that, of the four types of behavior that can be exhibited by employed subjects, reciprocal cooperation was generally the least prevalent, regardless of the treatment. We do find some differences across the treatments, however. Reciprocal defection was significantly more prevalent in the

ZeroT and HighT treatments than in the LowT treatment. Unprovoked cooperation was significantly more prevalent in the HighT treatment than in the LowT and ZeroT treatments.

There are three types of non-employment behavior. Regression 1 indicates the prevalence of unemployed workers, discouraged employers, and inactivity were significantly lower than reciprocal cooperation in the LowT and ZeroT treatments. In the HighT treatment, the prevalence of inactivity was significantly higher than reciprocal cooperation, while the prevalence of unemployed workers and discouraged employers was approximately equal to the prevalence of reciprocal cooperation. Thus, the prevalence of unemployed workers, discouraged employers, and inactivity were significantly lower in the LowT and ZeroT treatments than in the HighT treatment.

Regression 2 adds the variable CYCLE to regression 1, where CYCLE takes on the trade cycle values 1 through 12. Regression 2 indicates that subjects react to what they experience and significantly change their behavior over the twelve trade cycles. Only the prevalence of unemployment relative to reciprocal cooperation was unaffected. Relative to reciprocal cooperation, unprovoked defection and unprovoked cooperation each become significantly less prevalent, while reciprocal defection, discouragement among employers, and inactivity each become significantly more prevalent.

For the most part, we found demographic variables had no significant impact on the behavioral choices of the subjects. In particular, age, sex, race, marital status, financial aid status, weekly work hours, years of work experience, number of college credits, and whether or not the subject had been an employer in the real world had no impact. As shown in Regression 3, however, we did find that the student's GPA had an influence. Relative to reciprocal cooperation, higher GPA students were significantly more likely to exhibit unprovoked defection than lower GPA students and were also significantly more likely to become unemployed.

In addition to differing with regard to demographic characteristics such as sex, age, and GPA, human subjects can differ in terms of what motivates their behavior. Fehr and Gachter (1998) contrast *homo oeconomicus* and *homo reciprocans*, the former being self-interested and the latter being motivated by fairness concerns. Fehr and Schmidt (2000) found it useful to categorize decision-makers as *fair* or *selfish* when attempting to explain choices in situations with incomplete contracts.

Following this approach, we assume that unprovoked defection is *unkind*, while unprovoked cooperation is *kind*. We assume the fair type will not tend to exhibit unprovoked defection, but will exhibit unprovoked cooperation. Defecting is the Nash equilibrium choice for a single play of our prisoner's dilemma employment game. Thus, we assume the selfish type will tend to exhibit unprovoked defection. Of course, because our experiment involves repeated play, a self-interested subject may exhibit cooperative behavior for some time in an effort to encourage others to be cooperative, only to defect at the end. However, because we cannot discern the dynamic intentions of subjects, we simplistically assume the selfish type will not tend to exhibit unprovoked cooperation. Likewise, we assume the selfish type will not tend to exhibit reciprocal cooperation. We assume both selfish and fair types will exhibit reciprocal defection.

These assumptions allowed us to classify the human subjects according to their behavioral choices over the twelve trade cycles. Specifically, we classified a subject as *selfish* if there were more acts of aggressive defection than combined acts of aggressive cooperation and reciprocal cooperation. A subject was classified as *fair* if the opposite were true. A subject was classified as *half-selfish and half-fair* if the number of acts of unprovoked defection equaled the number of acts of unprovoked cooperation and reciprocal cooperation. The variable FAIR is the average percentage of session subjects who were classified as fair. Over all five sessions, an average of 47 percent of the subjects in a session were classified as fair, a percentage consistent with summary findings from a number of previous studies reported by Fehr and Gachter (1998).

Regression 3 indicates that, when a session contained a higher percentage of fair subjects, the likelihood of reciprocal cooperation was higher relative to all other behaviors. This is, of course, not too surprising given our definitions of selfish and fair types. What is more interesting is that the relative risk ratios on the FAIR variable are especially large (in absolute value) for the non-employment behaviors. This indicates that a stronger presence of fair types tends to be associated with a higher employment rate.

Using Regression 3, it is now desirable to re-examine the treatment effects after controlling for the significant behavioral impacts of the variables CYCLE, GPA, and FAIR. With only a single exception, we now find that the four employment behaviors are equally prevalent. That is, once we control for a subject's GPA and for the extent to which FAIR types are present in his session, the subject has a statistically equal probability of exhibiting unprovoked defection, reciprocal defection, unprovoked cooperation, or reciprocal cooperation. The single exception is that unprovoked cooperation is significantly more prevalent in the HighT treatment.

3.3. Outcomes

We next compare a variety of outcomes for the three treatments. For workers, we examine labor force participation rates, unemployment rates, the accumulation of offer costs, and market power. For employers, we examine job vacancy rates, discouragement rates, and market power. For all subjects combined, we examine the shirking rate, the rates of mutual cooperation, mutual defection, and single defection, mean production and total production, and mean utility.

There were three workers per session and five sessions per treatment. Each worker had to decide how to allocate a single work offer each trade cycle. Thus, there were 15 offers per trade cycle, or 180 offers over the twelve trade cycles in a given treatment. We define the labor force participation rate for a trade cycle to be (15-number of inactive workers)/15. To be unemployed the worker must seek work. Thus, we define the unemployment rate for a trade cycle by the ratio that includes the number of unemployed in the numerator and 15 minus the number of inactive workers in the denominator. Figures 2 and 3 present the labor force participation rates and unemployment rates for the three treatments.

-- INSERT FIGURE 2 ABOUT HERE --

-- INSERT FIGURE 3 ABOUT HERE --

Regression analysis indicates that there was no significant change in the labor force participation and unemployment rates over the twelve trade cycles in the LowT and ZeroT treatments. However, in the HighT treatment the rate of participation significantly decreased while the rate of unemployment significantly increased. By the twelfth trade cycle, the rate of participation in the HighT treatment was substantially lower than in the other two treatments, while the rate of unemployment was substantially higher. Even though the difference in the rate of participation between the LowT treatment and ZeroT treatment appears slight, regression analysis indicates that this difference was significant at a 5% level. The difference in the unemployment rates between the LowT and ZeroT treatments was also significant.

One might expect a higher payoff for non-employment to discourage labor force participation and encourage unemployment. More interesting is the fact that increasing the non-employment payment from zero to the low level had only a small impact, while increasing it from the low level to the high level had a much more significant impact. It is also significant that the difference between the HighT treatment and the other two treatments increased over the twelve trade cycles. This indicates learning may play a role in determining the extent to which a higher non-employment payoff discourages labor force participation.

Because the number of potential job openings equals the number of potential work offers, the job vacancy rate must equal the sum of the unemployment rate and the worker inactivity rate. Consequently, we find that the vacancy rate was positively correlated with the non-employment payoff. Moreover, whereas the vacancy rates in the LowT and ZeroT treatments remained fairly stable, the vacancy rate in the HighT treatment increased over successive trade cycles. A job vacancy in the ZeroT treatment was quite rare.

A job may be vacant either because an employer never receives a work offer or because an employer refuses all work offers he receives. We refer to the latter situation as *employer discouragement*. The discouragement rate is the number of vacancies caused by discouragement divided by the number of potential job openings. We find that discouragement was consistently rare in both the LowT and ZeroT treatments, while discouragement became increasingly common in the HighT treatment. In the last three trade cycles of the HighT treatment, employer discouragement was associated with roughly 40 percent of the job vacancies. The remaining 60 percent of vacancies were associated with worker discouragement in the sense that some workers chose to be inactive or stop searching for a job without soliciting each employer.

The following estimated equation characterizes the mean offer cost incurred by workers for each treatment.

$$OC = \underset{(0.09)}{1.21}HIGHT + \underset{(0.09)}{1.38}LOWT + \underset{(0.09)}{1.59}ZEROT - \underset{(0.01)}{0.007}CYCLE, \quad R^2 = .39.$$

HIGHT, LOWT, and ZEROT are dummy variables for the three treatments, and CYCLE indexes the trade cycle number running from 1 through 12. Standard errors are presented in parentheses. The coefficients on the treatment dummy variables indicate that increasing the non-employment payment reduces the mean offer cost incurred by workers. These differences are statistically significant at the 5% level. The coefficient on CYCLE indicates that the mean offer cost decreases with experience. However, this is result is not statistically significant at the 5% level.

Figure 4 presents the mean shirking rates of employers and workers over successive trade cycles, by treatment, where each mean is calculated over the subjects who enter into employment relationships.

-- INSERT FIGURE 4 ABOUT HERE --

In the HighT and ZeroT treatments, the shirking rate increased over successive trade cycles while the shirking rate in the LowT treatment decreased. To examine the significance of these trends, we regressed the shirking rate S on the treatment dummy variables HIGHT, LOWT, ZEROT, and on the interactions of these variables with the trade cycle variable CYCLE. The results are presented as Regression 1 in Table 4. Only the increasing trend exhibited in the ZeroT treatment is significant at the 5% level. However, the trends exhibited by the HighT and LowT treatments are significant at the 13.4% and 10.0% levels, respectively. By the twelfth and final trade cycle, the mean shirking rate was highest in the ZeroT treatment and by far the lowest in the LowT treatment.

-- INSERT TABLE 4 ABOUT HERE --

Shirking could occur in a context where both workers and employers shirk, or it could occur in a context where one type of agent shirks and the other does not. To examine the origin of shirking behavior and the evolution of cooperation in more detail, we define three possible relationship types: *Mutual cooperation* (*MC*), *mutual defection* (*MD*), and *single defection* (*SD*).

Under mutual cooperation, neither type of agent shirks. Regression 2 indicates that the ZeroT treatment was most effective in promoting mutual cooperation in the initial trade cycles. However, under this treatment the rate of mutual cooperation decreased significantly over successive trade cycles. In the final trade cycle 12, there was not a single incidence of mutual cooperation even though the labor force participation rate was 100% and the unemployment rate was 0%. In the HighT treatment, in contrast, the incidence of mutual cooperation was relatively low (about 14% of employment relationships) and did not exhibit a significant trend. Of interest, however, is the fact that the actual rate of mutual cooperation dropped from its peak level of 33% in trade cycle 11 to a low of 0% in trade cycle 12, an indication that subjects looked forward to their choice on the last trade cycle. In contrast to the other two treatments, the rate of mutual cooperation significantly increased in the LowT treatment. Even in the final trade cycle, 40% of the employment relationships in this treatment were mutually cooperative.

Under mutual defection, both types of agent shirk. Regression 3 indicates that mutual defection was rare in the ZeroT treatment in the initial trade cycles. However, the rate of mutual defection significantly increased so that, by trade cycle 12, mutual defection was more prevalent in the ZeroT treatment than in the other two treatments. The rate of mutual defection also significantly increased over successive trade cycles in the HighT treatment. In contrast, the rate of mutual defection did not significantly change from one trade cycle to the next in the LowT treatment. In the final trade cycle 12, the actual rate of mutual defection was 87% in the ZeroT treatment, 40% in the HighT treatment, and 17% in the LowT treatment.

When only one agent shirks in a work-site interaction between a worker and an employer, the shirker receives the high payoff 60 and the non-shirker receives the “sucker” payoff 10. The shirker could be either the worker or the employer. However, we found no statistically significant difference in the rates with which workers and employers engaged in such unilateral defections. Thus, we combined all cases into one *single defection* category. Regression 4 indicates that the likelihood of a sucker payoff differs little across the three treatments in the initial trade cycles. However, the likelihood of a sucker payoff decreases in successive trade cycles in both the HighT treatment and the ZeroT treatment, while it does not significantly change in the LowT treatment. Consequently, the regression model indicates that the rate of sucker payoffs is higher in the LowT treatment than the other two treatments during the final few trade cycles. An interesting anomaly, however, is that the sucker payoff rate was highest in the HighT treatment during the final trade cycle 12 (60 percent), as a number of subjects who had been in mutually cooperative relationships defected.

Workers and employers could earn points in two ways: namely, from an employment relationship and from non-employment payments. We can think of the points generated by employment relationships as output produced, while we can think of non-employment payment points as value that a government must obtain elsewhere. How many points are produced from employment relationships depends upon how many employment relationships are formed and upon the behavioral choices made by the worker and the employer in these relationships. For a given employment relationship, the most points are produced under mutual cooperation (40+40=80). The least points are produced under mutual defection

(20+20=40). A single defection produces an intermediate number of points (60+10=70). Of course, no points are produced when a worker chooses to make no offer and an employer's job is left vacant.

The total production level over the 106 employment relationships that developed in the HighT treatment was 7,480, an average of 70.6 per relationship. In the LowT treatment, 157 employment relationships produced 10,320 points, an average of 65.8 per relationship. Finally, in the ZeroT treatment, 176 employment relationships produced 9,953 points, an average of 56.5 per relationship. Thus, we find that increasing the non-employment payoff discourages the formation of employment relationships, but increased the productivity of the average relationship. Total production was highest in the LowT treatment.

Figures 5 and 6 show how mean production and total production varied over the twelve trade cycles. Mean production in the HighT and LowT treatments remained roughly constant, while mean production in the ZeroT treatment decreased significantly. The mean for the HighT treatment was higher than for the LowT treatment. This difference is significant at a 10% level but not at a 5% level.

-- INSERT FIGURE 5 ABOUT HERE --

-- INSERT FIGURE 6 ABOUT HERE --

Total production significantly decreased in the HighT and LowT treatments, but for opposite reasons. The productivity of the average employment relationship was maintained in the HighT treatment. However, total product decreased because the number of employment relationships declined. In the ZeroT treatment, the number of employment relationships was maintained, but the productivity of the average relationship decreased. By the final few trade cycles, total production was highest in the LowT treatment. In these final few trade cycles, there were fewer employment relationships in the LowT treatment than in the ZeroT treatment, but the productivity of the average relationship was higher. The productivity of the average employment relationship was lower in the LowT treatment than in the HighT treatment, but there were more employment relationships.

We can use the average number of points earned per trade cycle as a measure of the *utility* of a subject. There were six subjects in five sessions in three treatments, each earning points in twelve trade cycles. This provides us with 1080 observations. Half of these observations are on employers and half on workers. The following estimated equation shows how a subject's utility level U was influenced by the treatment, by experience over the trade cycles, and by the subject's type (worker or employer).

$$\underset{}{U} = \underset{(1.85)}{28.0HIGHT} + \underset{(1.85)}{29.3LOWT} + \underset{(1.85)}{32.3ZEROT} - \underset{(0.24)}{.07HC} - \underset{(0.24)}{.32LC} - \underset{(0.24)}{.93ZC} + \underset{(0.96)}{1.95SUB}, \quad R^2 = .0199$$

The variables HIGHT, LOWT and ZEROT are treatment dummy variables, while HC, LC, and ZC are interaction variables obtained by multiplying the treatment dummy variables and the trade cycle variable

CYCLE. SUB is a dummy variable that is 0 for a worker and 1 for an employer. Standard errors are shown in parentheses.

The very low R^2 value indicates that the variation in subject utility levels is not well explained by the factors we are examining here. We do learn, however, that employers enjoyed a significantly higher utility level than the workers. (This is presumably because workers had to pay an offer cost, whereas employers did not have to pay any comparable cost.) Controlling for this worker-employer difference, the average subject in each treatment received about the same utility outcome in the initial trade cycle. Over the twelve trade cycles, the utility level of the average subject decreases, regardless of the treatment. However, only the decrease in the ZeroT treatment is significant at a 5% level. Thus, we learn that, by trade cycle 12, the average subject in the ZeroT treatment is significantly worse off than the average subject in either the LowT or the HighT treatment. At a 5% level of significance, we cannot reject the null hypothesis that the subjects in the HighT treatment attain the same average utility level as the subjects in the LowT treatment.

We can examine the *market power* of a subject by examining how the subject's utility level compares to what would be received in a *competitive equilibrium*. We assume all agents would trade costlessly in a cooperative manner in competitive equilibrium, receiving the 40-point payoff associated with mutual cooperation in each trade cycle. This allows us to define a market power measure MPow = (U-40)/40, which is the percentage by which a subject's utility level U exceeds the competitive payoff outcome. The following estimated equation shows how a subject's market power was influenced by the treatment, by experience over the trade cycles, and by the subject's type (worker or employer).

$$\underset{}{MPow} = \underset{(0.5)}{-.30HighT} \underset{(0.05)}{-.26LowT} \underset{(0.05)}{-.19ZeroT} \underset{(0.006)}{- .002HC} \underset{(0.006)}{- .008LC} \underset{(0.006)}{- .023ZC} \underset{(0.024)}{+ .049SUB}, \; R^2 = .0199$$

Because the market power measure MPow is a linear transformation of the utility level U, the statistical properties of this regression are the same as the last. The regression indicates that the average subject does worse in each of our treatments than he would do in a competitive equilibrium. The market power of workers is significantly lower than the market power of employers. The highest average market power level is attained by employers in the initial trade cycles of the ZeroT treatment, though at a 5% level of significance this market power level is not statistically different than that of employers in the ZeroT and LowT treatments. The lowest average market power level is attained by workers in the final trade cycle 12 for the ZeroT treatment, and this result is statistically significant at a 5% level. We therefore cannot reject the null hypothesis that there is no difference in the average market power levels observed in the LowT and HighT treatments.

4. The Computational Experiment

4.1. The Computational Labor Market

The computational labor market closely parallels the general structure of the human-subject labor market set out in Section 2. Thus, the market comprises equal numbers of workers and 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. Also, the workers and employers repeatedly seek preferred work-site partners using a modified Gale-Shapley matching mechanism, engage in efficiency-wage work-site interactions modeled as prisoner's dilemma games, and evolve their work-site behaviors over time. Advantage is taken, however, of the ability to consider greater numbers of workers and employers interacting over longer periods of time. In addition, unlike the case for human subjects, the learning processes of the computational workers and employers must be explicitly specified.

The computational experiment was 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 animation, table, and chart displays.[1] The specific TNG parameter settings for the experiment at hand are described in the next section. All other TNG parameter settings are the same as in Tesfatsion (2001).

4.2. Implementation Details

The computational labor market comprises twelve workers and twelve employers. Market activities are divided into a sequence of 1000 *generations*. Each generation in turn is divided into two parts: (a) a *trade cycle loop* consisting of successive *trade cycles*; and (b) an *evolution step*.

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"[2] with 16 internal states. Thus, the set of feasible work-site rules for each worker and employer, while extremely large, is nevertheless finite.

---

[1] For research, tutorials, user instructions, source code, and executables related to the TNG Lab, visit the TNG home page at http://www.econ.iastate.edu/tesfatsi/tnghome.htm.

[2] A *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

Each worker and employer in the initial generation also has an initial expected utility assessment for each of his potential work-site partners. For simplicity, in the experiments at hand, these assessments are all set equal to the mutual cooperation payoff. [3]

The workers and employers in the initial generation participate in 150 successive trade cycles. During 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 expected utility assessments; and (2) an employment process, during which each matched worker-employer pair engages in one work-site interaction. In addition, during each process, each agent updates his current expected utility assessment for a particular potential partner every time that he obtains a payoff from an interaction with this potential partner.

The matching process is the same as described in Section 2 for the human-subject experiment, with one exception. Each computational worker and employer has a *minimum tolerance level,* assigned as part of the agent configuration process for the initial generation.[4] If the expected utility assessment assigned to an employer by a worker ever falls below this 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 this level, the employer will stop accepting work offers from this worker. In the human-subject experiment, the minimum tolerance levels are not under the control of the experimenter.

Given these minimum tolerance levels, the manner in which workers direct work offers to employers in the matching process for each trade cycle follows the modified Gale-Shapley procedure outlined in Section 2. Each worker and employer has a preference ranking over potential partners, determined by his current expected utility assessments. Each worker starts by directing a work offer to his most preferred tolerable employer, if any such employer exists. 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.

---

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.

[3] This 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 6.

[4] In the current experiment, these minimum tolerance levels are consistently set equal to the non-employment payoff. This specification is meant to capture the opportunity cost idea that entering into a risky work-site interaction is viewed by agents 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.

The employment process for the computational experiment is the same as for the human-subject experiment. Once a worker and employer are matched, their subsequent work-site interaction is modeled as a prisoner's dilemma game. 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, as in the human-subject experiment, a worker incurs an *offer cost* OC=1.0 each time that 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 payoff* NEP for the trade cycle. Each worker and employer keeps complete track of each payoff he receives during the course of each trade cycle, including work-site payoffs, negative payoffs due to offer costs, and non-employment payoffs.

Each worker and employer uses a simple reinforcement learning algorithm to update his expected utility assessments for potential partners in response to new payoffs. Every worker and employer initially associates $U^0 = C$ with each potential work-site partner with whom he has not yet interacted. Subsequently, each time an agent v interacts with an agent z, agent v forms an updated expected utility assessment for z by summing U^0 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 both work-site payoffs and any negative payoffs due to offer costs. Consequently, an updated expected utility assessment for any agent z is simply an average of all payments received to date in interactions with z, augmented by the "prior" payoff expectation U^0. This method for updating expected utility assessments is consistent. That is, if an agent interacts repeatedly with another agent z, then eventually his expected utility assessment for z will converge to his true average payoff level from interactions with z.[5]

At the end of the initial generation, the *utility level* of each worker and employer is calculated to be the average total net payoffs per trade cycle that the agent obtained during the course of the preceding trade cycle loop, i.e., his total net payoffs divided by 150. The workers and employers then enter into an evolution step during which they separately evolve (structurally update) their work-site rules based in part on these actual attained utility levels.

In the evolution step, each worker and employer engages in inductive learning by experimentation with the use of new work-site rules. In addition, each worker and employer also engages in social mimicry learning by mimicking aspects of the work-site rules used by more successful (higher utility)

[5] See McFadzean and Tesfatsion (1999) for more details. Briefly, consistency follows from the finite state automaton representation for work-site rules, which ensures that the action patterns between any two agents must eventually enter into a cycle as the number of their interactions becomes sufficiently large.

agents of his 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).[6]

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 potential 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.[7]

4.3. Experimental Design

The computational experiment parallels the human-subject experiment in focusing on only one treatment variable, the non-employment payoff NEP. The three tested values for NEP are 0, 15, and 30, which are referred to as treatment ZeroT, treatment LowT, and treatment HighT, respectively.

For each treatment, 20 runs were generated using 20 different seeds for the TNG Lab pseudo-random number generator: namely, {0,5,10,…,95}. In the data tables reported in Section 5, each run is identified

---

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

[7] A 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 to zero in the TNG Lab, the software used to implement the computational experiment. Thus, to retain the non-employment payoff 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 payoff 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 NEP*=0 is used in place of NEP. In all data tables presented in Section 5, 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.

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 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. Table 5 provides a summary of these basic experimental design features.

-- INSERT TABLE 5 ABOUT HERE --

5. Results for the Computational Experiment

5.1. Overview

In contrast to the results for the human-subject experiment reported in Section 3, 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 non-employment payoff treatments, the twenty runs essentially 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 as a function of the non-employment payoff treatment.

Before reporting these results 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. These descriptive statistics capture many of the same aspects as reported for the human-subject experiment. With computational agents, however, we can be much more precise about distinguishing persistent from transient relationships and continuous relationships from relationships that are intermittent or random in nature.

5.2. Descriptive Statistics

Definition of a Persistent Relationship

As noted in Section 4, work-site rules in the computational experiment 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 can be summarized in the form of a *work-site history* H:P, where the *handshake* H is a (possibly null) string of work-site actions that form a non-repeated pattern and 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*.

Classification of Networks by Competitive Distance

We will next construct a measure that permits the classification of experimentally observed interaction networks into alternative types. This measure calculates the distance between an observed interaction network among workers and employers and an idealized interaction network capable of supporting a competitive (full employment) market outcome.

A possible interaction pattern among the workers and employers in generation G of a run R is referred to as an *(interaction) network*, denoted generically by N. Each network N 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.

Let N(G,R) denote the network depicting the actual interaction pattern among the workers and employers in generation G of a run R. 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).

Consider the following *competitive interaction pattern* between workers and employers: *Each worker is recurrently directing work offers to employers, and every worker and employer has at least one* *persistent relationship.* This interaction pattern is the pattern that would result under "competitive market conditions" in which every worker and employer is behaving cooperatively, each worker has the same expected employment rate as any other worker, and each employer has the same expected vacancy rate as any other employer.

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, this measure indicates the extent to which PN(G,R) deviates from the "null hypothesis" of a competitive market network. This network distance measure provides a useful way to classify different types of persistent networks observed to arise in the computational experiment. Note, in particular, that a completely recurrent persistent network has a network distance of 0, a completely latched persistent network has a network distance of 12, and a completely disconnected persistent network (no persistent relationships) has a network distance of 24.

Measurement of Market Non-Participation Rates

A worker or employer who fails to form any persistent relationship during generation G in run R is classified as *persistently non-employed* for that generation and run. The percentage of workers who are persistently non-employed constitutes the *persistent unemployment rate* for that generation and run. Similarly, the percentage of employers who are persistently non-employed constitutes the *persistent vacancy rate* for that generation and run.

A persistently non-employed worker is essentially an unemployed worker in the sense of Section 3, and a persistently non-employed employer is essentially a discouraged employer in the sense of Section 3. No agent in the computational experiment is inactive in the sense of Section 3 when inactivity is measured in terms of an entire trade cycle loop. In particular, then, the labor force participation rate is 100% over each trade cycle loop.[8]

---

[8] In the computational experiment, the workers' initial expected utility assessments for potential employers are all set equal to the mutual cooperation payoff level, C. In each treatment, C exceeds the minimum tolerance level, which is set equal to the non-employment payoff. It follows that each worker at the beginning of the trade cycle loop for generation G will always direct at least one work offer to an employer, even if that worker eventually becomes persistently non-employed. Thus, every persistently non-employed worker is essentially an unemployed worker in the sense of Section 3, and no worker is inactive in the sense of Section 3, when measured in terms of an entire trade cycle loop. In addition, since all employers in the computational experiment are initially viewed indifferently by workers at the beginning of a trade cycle loop, and potential work offers equal potential job offers, every employer will receive at least one work offer in the initial trade cycle of a trade cycle loop. Thus, no employer is inactive in the sense of Section 3 when measured in terms of an entire trade cycle loop.

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. An NPD is an unprovoked defector in the sense defined in Section 3 for the human-subject experiment, but an unprovoked defector need not be an NPD.

An agent (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 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

As in Section 3, 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 computational labor market to be a situation in which each worker is recurrently directing work offers to employers, and each agent is a persistent cooperator (AllC).

Ignoring offer costs, the utility level that each worker and employer would attain in these competitive market conditions is simply the mutual cooperation payoff level, C. Therefore, as in Section 3, we define the *market power* 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 mutual defectors (AllDs), the relationship is classified as *mutual defection* (*M-AllD*). If both participants are persistent mutual 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.

These relationship classifications parallel the reciprocity classifications introduced in Section 3 for human-subject relationships, with two major differences. First, the relationship classifications for the computational experiment focus only on persistent relationships; transient relationships are ignored. Second, these relationships are classified in accordance with the types of behaviors expressed by each participant in these relationships, where these behaviors are holistically described in terms of the categories IntD, AllD, and AllC. That is, the focus is on relationships as a whole, not on action-pair counts within relationships.

This is done because agents in the computational experiment evolve their work-site rules holistically, on the basis of their entire past interaction histories to date. It would therefore be highly misleading to assume that an action by one agent represents a response only to the immediate past action of his partner. Indeed, as will be reported in Section 5.3, there is an enormous amount of coordination in the relationships observed for the computational experiment, in the sense that mutuality relationships totally dominate. This indicates that agents are engaging in a high degree of relationship coordination in a holistic manner.

In contrast, in the human-subject experiment it is much more difficult to get a handle on which relationships are truly persistent. The interaction horizons are simply too short. Moreover, it is difficult to determine how the workers and employers are actually determining their work-site behaviors. Consequently, it makes sense to break down observed relationships into action-pair counts in an attempt to get a handle on the characterization of these behaviors.

With these cautions in mind, the concept of "reciprocal cooperation" in Section 3 is essentially captured in holistic terms by the concept of an M-AllC relationship. Similarly, the concept of "reciprocal defection" in Section 3 is essentially captured in holistic terms by M-AllD relationships and by *balanced* M-IntD relationships in which the attained market power levels of workers and employers are essentially equal, taking offer costs into account.

Finally, the concept of a "selfish agent" in Section 3 is captured in the computational experiment by a participant in a persistent relationship who has a higher shirking rate than his partner. This can occur in a variety of ways. For example, a worker participating in a relationship of the form AllD-AllC, AllD-IntD or IntD-AllC is clearly shirking at a higher rate than the participating employer. Moreover, a worker participating in an M-IntD relationship can be inferred to be shirking at a higher rate than the participating employer if the worker is attaining a relatively higher market power level, taking offer costs into account.

5.3. Experimental Findings

Network Formation in the Three Treatments

For each of the twenty runs corresponding to each non-employment 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 5.2, 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 7. Network distance is measured along the horizontal axes and the number of runs corresponding to this network distance is indicated on the vertical axes. Recall that a network distance of 0 corresponds to a perfectly recurrent persistent network, a network distance of 12 corresponds to a perfectly latched persistent network, and a network distance of 24 corresponds to a completely disconnected network (no persistent relationships).

-- INSERT FIGURE 7 ABOUT HERE --

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).

Market Non-Participation Rates, Work-Site Behaviors, and Market Power in the Three Treatments

Table 6 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).

-- INSERT TABLE 6 ABOUT HERE --

The twenty runs 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, all exhibiting perfectly latched persistent networks. The remaining two runs 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, very high percentages of the workers and the employers are AllCs. Despite the prevalence of AllC agent types, the employers attain an average mean market power level (-0.02) that is markedly higher than the corresponding level obtained by the workers (-0.14). 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, IntD, and AllD agent types among employers.

In the second, more dominant attractor state comprising fourteen runs, very high percentages of the workers and the employers are both NPDs and IntDs. Interestingly, the workers and employers obtain similar average mean market power levels in this second attractor state (-0.22 for workers and -0.20 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.

Table 7 reports persistent relationship type counts for treatment ZeroT, sampled at generation 12. As in Table 6, data are reported for individual runs grouped into the two attractor states.

-- INSERT TABLE 7 ABOUT HERE --

The most striking aspect of Table 7 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 7 reveals that the first attractor state is dominated by mutual cooperation (M-AllC) whereas the second attractor state is dominated by mutual intermittent defection (M-IntD). Mutual defection (M-AllD) is almost entirely absent.

The mean market power levels reported in Table 7 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 attractor state are very close, this hides an underlying volatility in relative shirking rates across runs.

The characteristics reported above 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, agents on average 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 runs in the first attractor state are characterized by perfect latching and a high percentage of AllC agent types in M-AllC relationships. The runs in the second attractor state are characterized by almost perfect latching and a high percentage of IntD agent types in M-IntD relationships. The 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 7. 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 -0.11, whereas the workers attain a markedly lower level of -0.19. 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 6 and Table 7 for treatment ZeroT sampled at generation 12. The first attractor state comprises nine runs, the second attractor state comprises seven runs, and only four runs appear as transition states. In the first attractor state the workers and employers attain average mean market power levels of -0.06

and -0.05, respectively. In the second attractor state the workers and employers attain uniformly lower but balanced average mean market power levels of -0.17 and -0.18, respectively. As in Table 7, this balance hides an underlying volatility in shirking rates across runs. Finally, the first attractor state is strongly dominated by M-AllC relationships, and the second attractor state is strongly dominated by M-IntD relationships. Mixed types of relationships are almost entirely absent in the two attractor states.

Table 8 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 9 reports persistent relationship type counts for these same runs, again sampled at generation 12.

-- INSERT TABLE 8 ABOUT HERE --

-- INSERT TABLE 9 ABOUT HERE --

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 payoff, agents are opting for non-employment rather than choosing to remain in M-IntD relationships.

It is interesting to observe how these outcomes for treatment HighT vary from generation 12 to generation 1000. 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. This suggests a movement over time toward increased coordination on the first attractor state in which M-AllC relationships are dominant.

Summarizing the relative market power outcomes of workers and employers 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 workers. 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 -0.05 to -0.14. 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 8 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 -0.01 to -0.06.

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.

Finally, pooling all runs together for each treatment at each sampled generation, the following results are obtained for MP, the average market power level attained per agent. For all sampled generations, MP is uniformly lower in treatment ZeroT than in treatments LowT or HighT. For generations 12 and 50, MP is highest in treatment LowT. By generation 1000, however, MP is highest in treatment HighT. The latter finding reflects the previously noted observation that agents in treatment HighT become increasingly more successful at coordinating on persistent mutual cooperation (the first attractor state) rather than persistent non-employment (the second attractor state) in each successive sampled generation.

Never-Provoked Defection in the Three Treatments

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 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 6 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 payoff 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 about these two treatments 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 payoff 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 payoff 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 payoff 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 8, 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 each agent type engages 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 picker 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.

6. Comparison of Human-Subject and Computational Experiment Results

6.1. Network formation in the Two Experiments

In both the human-subject and computational experiments, the number of employment relationships decreases as the non-employment payoff increases. An increasing number of workers and employers choose inactivity in preference to engaging in risky work-site interactions.

Nevertheless, the interaction networks among those who choose to remain in the labor market exhibit substantial differences in the two experiments. In the computational experiment, as seen in Figure 7, almost every employed worker ends up in a continuous long-term relationship with a single employer. In contrast, as seen in Table 2, recurrent relationships are dominant in the human-subject experiment; latched relationships are only rarely observed.

A key explanation for this difference lies in the specification of the offer cost in relation to the scale of each experiment. In an attempt to run parallel experiments, the same low value of 1.0 was used for the offer cost in each experiment. Nevertheless, in the computational experiment, we chose to include twelve workers and twelve employers in order to sustain "genetic diversity" in the pool of work-site rules for genetic-algorithm learning purposes. In contrast, for the human-subject experiment, we chose to include only three workers and three employers in order to keep the experiment from becoming unwieldy.

The offer cost is incurred per work offer. Consequently, the difference in scale between the two experiments implies that a worker in a recurrent network can end up making a substantially larger number of work offers per trade cycle in the computational experiment than in the human-subject experiment. This arises because, in the computational experiment, there are four times as many employers per worker. Consequently, if workers are randomly or intermittently directing work offers across employers without latching, and each employer has only one job opening, there is a much greater risk of accumulating a large number of refusals before employment is secured. It follows that a recurrent network is much more costly for workers to maintain in the computational experiment. Indeed, an example of this is seen in Table 9 for treatment HighT sampled at generation 12. Workers attain a relatively low mean market power level in the transition-state runs in which recurrent relationships dominate, despite the fact that a very high percentage of the workers and employers in these recurrent relationships are engaging in mutually cooperative behavior.

This potential for accumulating large numbers of refusals, hence offer costs, in a recurrent network also explains why latched networks are so much more prevalent in the computational experiment. Refusals lower the expected utility assessments of the employers doing the refusing and encourage workers to latch on to any employer who does hire them.

A conjecture that follows from these observations is that latching behavior should become more prevalent in the human-subject experiment as the offer cost is increased or, alternatively, as the number of workers and employers is proportionally increased while retaining zero excess capacity.

6.2. Market Non-Participation Rates in the Two Experiments

Unemployment and vacancy rates in the computational experiment are easily characterized. In treatments ZeroT and LowT, both rates are close to 0% for all runs in all sampled generations lying in either of the two attractor states. See, for example, Table 6.

In treatment HighT, the unemployment and vacancy rates show some evolution over time. When runs are sampled at generation 12, the average unemployment and vacancy rates are 10% and 12%, respectively, in the first attractor state dominated by mutual cooperation, and both average rates are approximately 97% in the second attractor state dominated by complete or near-complete coordination failure. (See Table 8.) When the runs for treatment HighT are instead sampled at generation 50 and generation 1000, the average unemployment and vacancy rates are close to 0% in all runs in the first attractor state dominated by mutual cooperation and equal to 100% in the second attractor state characterized by complete coordination failure.

Qualitatively similar results are obtained in the human-subject experiment. For example, as depicted in Figure 3, the unemployment rate (averaged over all workers in all five sessions) is substantially higher in treatment HighT than in treatments ZeroT and LowT.

On the other hand, Figure 3 also shows that the higher unemployment rate for treatment HighT in the human-subject experiment is particularly marked in the last few trade cycles; that is, the average unemployment rate in treatment HighT seems to increase as the human subjects become more experienced. This phenomenon is not observed for the computational experiment. Indeed, to the contrary, the computational workers and employers in treatment HighT become better at coordinating on mutual cooperation and avoiding coordination failure in each successive sampled generation. This is indicated by a monotonic growth in the size of the cooperative first attractor state.

This latter difference may perhaps be arising due to the relatively optimistic specification of initial expected utility assessments in the computational experiment. For the first generation of agents, initial expected utility assessments are set equal to the mutual cooperation payoff, and these optimistic initial assessments are then reconstituted for each successive generation of agents. This reconstituted optimism ensures that each successive generation of agents in the computational experiment keeps attempting to participate in the labor market even if they have experienced bad payoff outcomes in the past.

In contrast, the participants in the human-subject experiment appear to become discouraged rather quickly when they cooperate against a defecting partner. In particular, once subjects in treatment HighT opted for inactivity during a particular trade cycle, they tended to remain inactive for the remainder of their sessions.

6.3. Unprovoked Behavior and Reciprocity in the Two Experiments

In the human subject experiment, unprovoked defection (UD) and unprovoked cooperation (UC) represent actions taken unilaterally in relationships with less familiar partners, while reciprocal defection (RD) and reciprocal cooperation (RC) are mutuality actions that can arise in relationships with more familiar partners. In the computational experiment, never-provoked defection (NPD) characterizes actions taken unilaterally with strangers or in (typically early) interactions with fully cooperative partners, while mutual cooperation (M-AllC), mutual defection (M-AllD), and mutual intermittent defection (M-IntD) represent mutuality patterns that can emerge in persistent relationships with more familiar partners.

Examining behavior when the partner is less familiar, the incidence of UD in the human-subject experiment was higher in treatment HighT than in treatments LowT and ZeroT. The results of the computational experiment were similar in that the move to the high non-employment payoff in treatment HighT led to an extremely high incidence of NPD in the second attractor state associated with persistent non-employment. Thus, a consistent message obtained from our experiments is that a high non-employment payoff leads to a greater incidence of aggressive defection behavior, which then leads to non-employment.

In addition, however, in every treatment for the computational experiment, mutual cooperation was the dominant behavior in the first attractor state. This, then, highlights a second possibility. Namely, if a worker and employer are "nice" or "fair" and do not aggressively defect in their initial few interactions with each other, then the outcome will tend to be a productive cooperative employment relationship rather than persistent intermittent defection or non-employment. Some of this mutual cooperation was indeed observed between pairs of agents in the human-subject experiment in both the LowT and HighT treatments.

Examining human-subject behavior when the partner is more familiar, the rates of reciprocal cooperation in the final trade cycle 12 were comparable for the HighT and ZeroT treatments, while the rate of reciprocal cooperation was significantly higher in the LowT treatment. These results are generally consistent with those obtained for the computational experiment when we consider average behavior across both attractor states for each treatment.

More precisely, in the computational experiment, reciprocal cooperation (in the form of M-AllC relationships) dominated the first attractor state in all three treatments. However, mutual intermittent defection was the dominant relationship in the second attractor state in treatments ZeroT and LowT whereas persistent non-employment was the dominant outcome in the second attractor state in treatment HighT. Sampled at generation 12, for example, the average incidence of M-AllC relationships across both attractor states was 0.4 per run for treatment ZeroT (see Table 7) and 3.5 per run for treatment HighT (see Table 9). In contrast, sampled at generation 12, the average incidence of M-AllC relationships across both attractor states was 6.0 per run for treatment LowT.

Thus, another consistent message obtained from the two experiments is that a zero non-employment payoff may deter reciprocal cooperation by promoting too much intermittent defection whereas a high non-employment payoff may deter reciprocal cooperation by luring agents into inactivity.

One interesting issue raised by these comparisons is the treatment of purely transient relationships. In the computational experiment, the focus is on whole relationships rather than actions within relationships, and only persistent relationships are considered in the counts of mutual (reciprocal) cooperation and defection. In contrast, in the human-subject experiment, reciprocal cooperation and reciprocal defection are measured on the basis of incidences of paired-action behaviors within relationships, even if these relationships are transient in nature.

The implications of these differences in definition are vividly illustrated by the different conclusion reached in the two experiments regarding reciprocal defection in treatment HighT. For the computational experiment, as reported in Table 9, the conclusion reached is that reciprocal defection (in the form of M-AllD or balanced M-IntD relationships) is virtually absent. For the human-subject experiment, as reported in Table 3 (regression 1), the conclusion reached is that reciprocal defection is prevalent. The apparent discrepancy is resolved, however, by the following observation. Although almost every action taken by the computational workers and employers in the second attractor in treatment HighT is a defection, none of these defections contributes toward the measure of reciprocal defection because no agent manages to form any persistent relationships. In contrast, any defection in response to a defection would count as an incidence of reciprocal defection in the human-subject experiment, whether or not the relationship as a whole was transient or persistent.

The important point raised by these observations for human-subject and computational experimentalists is that time scale matters greatly. From the vantagepoint of the 150^{th} trade cycle in generation 1000 in the computational experiment, characterizing the behaviors expressed in transient relationships terminating in earlier trade cycles, let alone earlier generations, may seem to be a waste of time. However, as stressed by Roth and Erev (1995), such transient relationships may be where life in the real world predominantly takes place.

6.4. Market Power Outcomes in the Two Experiments

Our market power definition measures the utility level of an agent relative to the utility level that the agent would attain in a theoretical competitive equilibrium with no transaction costs. In both types of experiments, the mean market power levels attained by workers and employers are generally negative. This finding stems from the presence of transaction costs in the form of work offer costs and from the inability of agents to achieve the level of cooperation assumed in the competitive equilibrium.

A second consistency between the two experiments is that the mean market power of employers tends to exceed that of workers, even when the employers and workers are mutually cooperating. This occurs because workers bear all of the transaction costs associated with the formation and maintenance of employment relationships.

A third fundamental consistency between the two experiments is that agents in treatment ZeroT tend to attain lower average market power levels than agents in treatments LowT or HighT. Two factors explain this finding. First, agents in treatment ZeroT who choose to be non-employed do not receive a positive non-employment benefit. Second, the absence of a positive non-employment benefit encourages agents to engage in more shirking and less cooperation in their work-site interactions.

Of particular interest is the finding in the computational experiment that the mean market power of an average agent is highest in treatment LowT for generations 12 and 50 but highest in treatment HighT for generation 1000. Two factors explain this pattern. First, the increase in the non-employment payoff tends to increase the incidence of mutual cooperation in any employment relationships that form. Second, as time goes by, more agents learn that a higher payoff can be attained through mutual cooperation in an employment relationship than in non-employment, so that the absolute number of employment relationships increases.

This finding casts an interesting light on the human-subject experiment. In the latter, the mean market power levels were comparable for treatments LowT and HighT, but subjects in treatment HighT did exhibit a significantly higher rate of unprovoked cooperation. That is, some subjects in treatment HighT were exceptionally willing to risk receiving the lowest "sucker" payoff in order to pursue the rent associated with cooperative employment. The long-run mean success of treatment HighT agents in the computational experiment suggests that the aggressive pursuit of cooperative employment exhibited by HighT subjects in the human-subject experiment may have long-run implications. In particular, given a long enough time span, the mean welfare of the human subjects in treatment HighT might exceed the welfare of subjects in treatment LowT as well as in treatment ZeroT.

7. Conclusion

The results from the human-subject and computational experiments illustrate how the level of a non-employment payoff can influence economic outcomes by influencing the evolution of cooperation. Of particular interest is the finding that increasing the non-employment payoff effectively filters out employers and workers who are more likely to shirk, so that the productivity of the average employment relationship increases. In addition, as the non-employment payoff increases, the computational experiment suggests that a greater number of agents learn over time to coordinate on mutual cooperation and avoid coordination failure, so that overall efficiency increases as well.

While we did find evidence in the human-subject experiment that some subjects seem to be "naturally" more cooperative than others, findings from both the human-subject and computational experiments suggest that the level of cooperation tends to evolve as agents learn from their environment. The computational experiment illustrates that environment and chance can both have an influence on this evolutionary process. When the non-employment payoff is high, the computational labor market either evolves towards a highly productive economy where everyone is cooperative or towards economic collapse with 100 percent unemployment. The difference between approaching one attractor state versus another depends upon which agents happen to pair up in the early stages of the evolutionary process, and how they happen to behave in these pairings. This path dependence suggests that, while a change in the level of the non-employment payoff may have substantial effects on key economic outcomes such as unemployment and average productivity, the standard deviation of these outcomes could be quite large.

References

Albrecht, J. and S. Vroman, 1998. Nash equilibrium efficiency wage distributions, International Economic Review 39(1), 183-203.

Axelrod, R., 1984. The Evolution of Cooperation (Basic Books, New York).

Bolton, G.E., and A. Ockenfels, 2000. A theory of equity, reciprocity and competition, American Economic Review 90(1), March, 166-193.

Carmichael, H. L., and W.B. MacLeod, 1997. Gift giving and the evolution of cooperation. International Economic Review 38(1), 485-509.

Duffy, J., 2001. Learning to speculate: Experiments with artificial and real agents, Journal of Economic Dynamics and Control, 25(3-4), 295-319

Falk, A., Gachter, S. and J. Kovacs, 1999. Intrinsic motivation and extrinsic incentives in a repeated game with incomplete contracts, Journal of Economic Psychology 20, 251-284.

Fehr, E. and S. Gachter, 1998. Reciprocity and economics: The economic implications of Homo Reciprocans, European Economic Review 42, 845-859.

Fehr, E., G. Kirchsteiger and A. Riedl, 1993. Does fairness prevent market clearing? An experimental investigation, Quarterly Journal of Economics 108 (2), 437-459.

Fehr, E. and K. Schmidt, 1999. A theory of fairness, competition, and cooperation. Quarterly Journal of Economics 114, 817-868.

Fehr, E., and K. Schmidt, 2000. Fairness, incentives, and contractual choices, European Economic Review, vol. 44 (4-6), 1057-1068.

Gale, D. and L. Shapley, 1962. College admissions and the stability of marriage. American Mathematical Monthly 69, 9-15.

Goeree, J.K. and C.A. Holt, 2000. Asymmetric inequality aversion and noisy behavior in alternating offer bargaining games, European Economic Review 44, 1079-1089.

Kreps, D., Milgrom, P., Roberts, J., and R. Wilson, 1982. Rational cooperation in the finitely repeated prisoners dilemma, Journal of Economic Theory 27, 245-252.

McFadzean, D., Stewart, D. and L. Tesfatsion, 2001. A computational laboratory for evolutionary trade networks, IEEE Transactions on Evolutionary Computation, forthcoming.

McFadzean, D. and L. Tesfatsion, 1999. A C++ platform for evolutionary trade networks, Computational Economics 14, 108-134.

MacLeod, W.B. and J. Malcomson, 1998. Motivation and markets, American Economic Review 88 (3), June, 388-411.

Orbell, J. M. and R. M. Dawes, 1993. Social welfare, cooperators' advantage, and the option of not playing the game. American Sociological Review 58, 787-800.

Rabin, M., 1993. Incorporating fairness into game theory and economics, American Economic Review 83(5), December, 1281-1302.

Rabin, M. and J. Shrag, 1999. First impressions matter: A model of confirmatory bias. Quarterly Journal of Economics 114, 37-82.

Roth, A. and I. Erev, 1995. Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior 8, 164-212.

Shapiro, C. and J. Stiglitz, 1984. Equilibrium unemployment as a worker discipline device, American Economic Review 74 (3), 433-444.

Tesfatsion, L., 1997. A trade network game with endogenous partner selection, In: H. Amman, B. Rustem, and A.B. Whinston (Eds.), Computational Approaches to Economic Problems, Kluwer Academic Publishers, Dordrecht, The Netherlands, 249-69.

Tesfatsion, L. 2001. Structure, behavior and market power in an evolutionary labor market with adaptive search. Journal of Economic Dynamics and Control 25, 419-457.

Figure 1

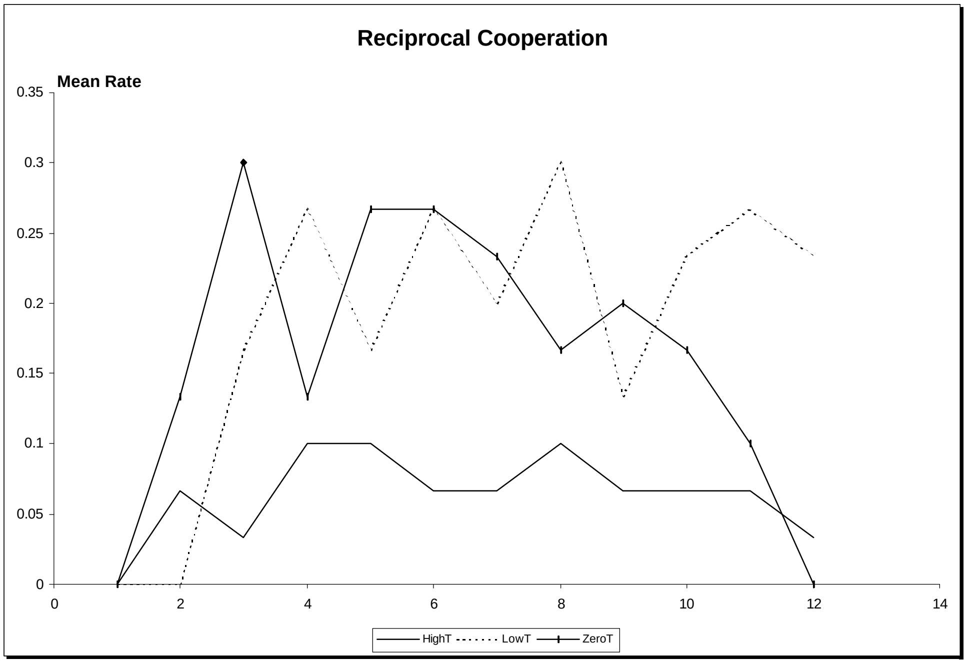

Figure 2

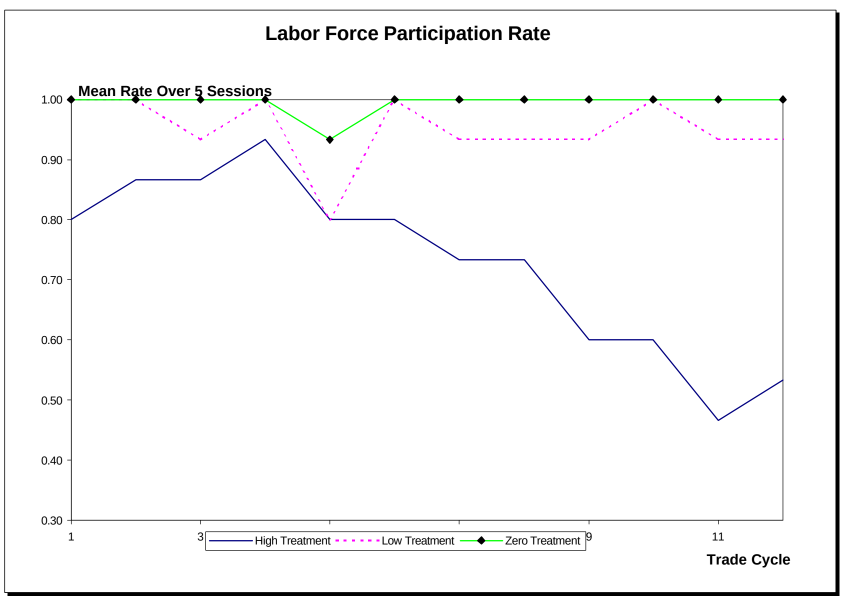

Figure 3

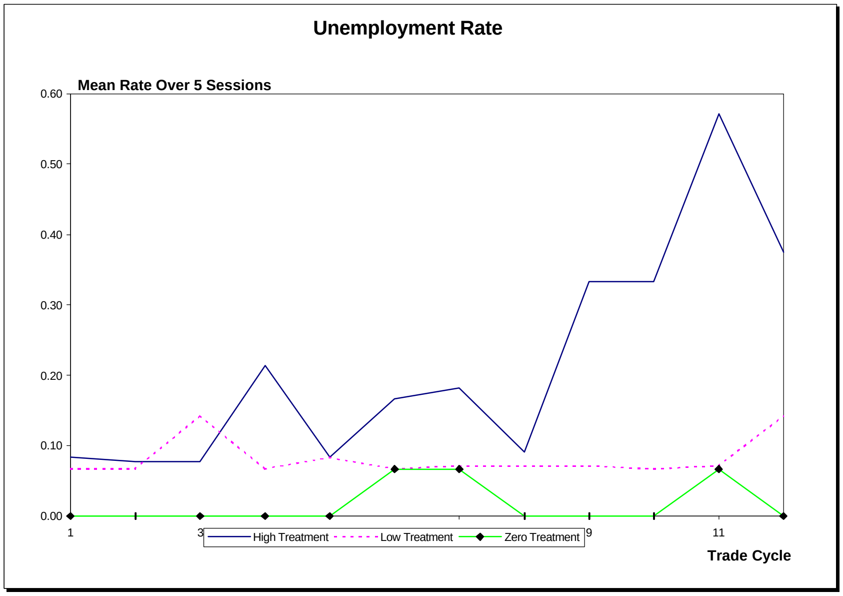

Figure 4

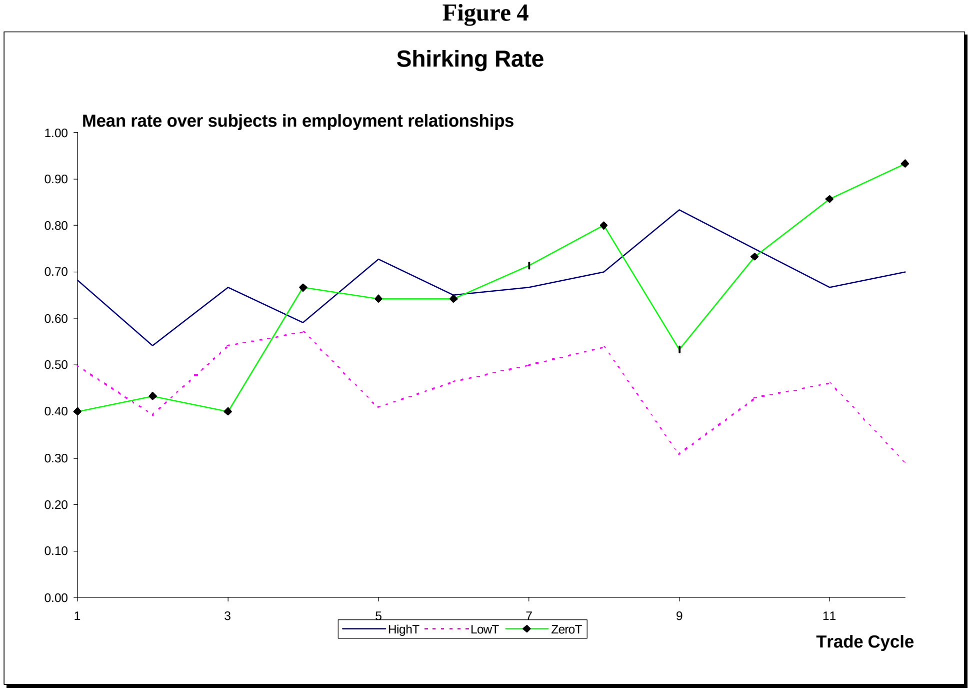

Figure 5

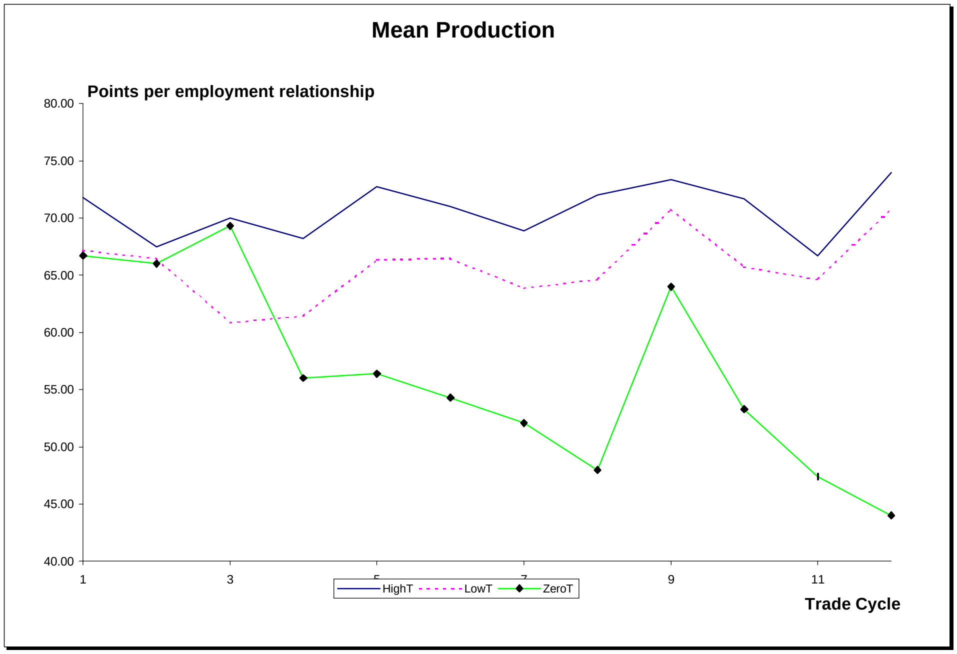

Figure 6

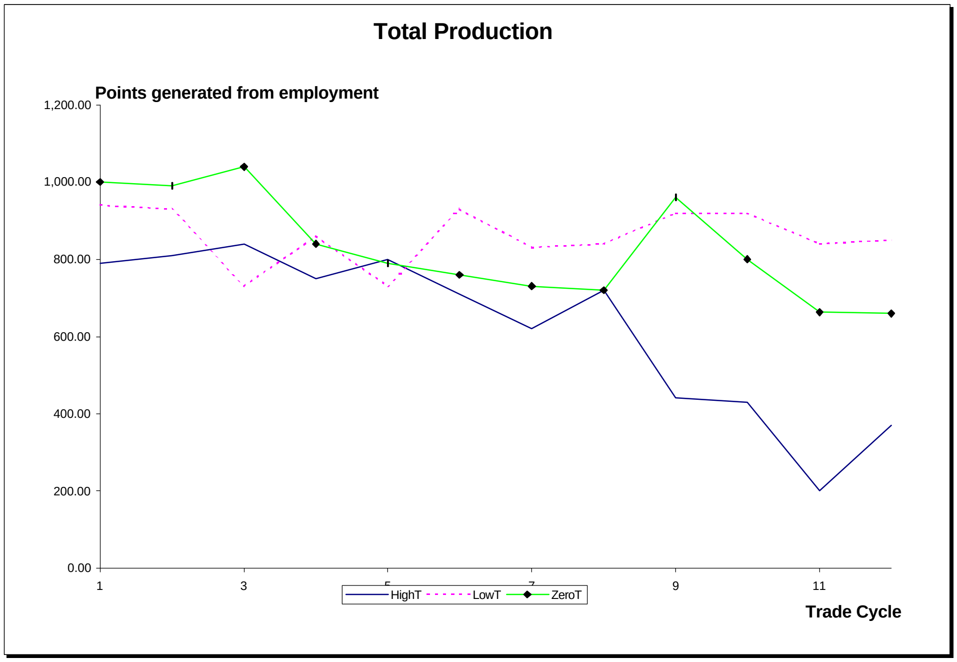

Figure 7: Network Distribution by Treatment and by Generation for the Computational Experiment

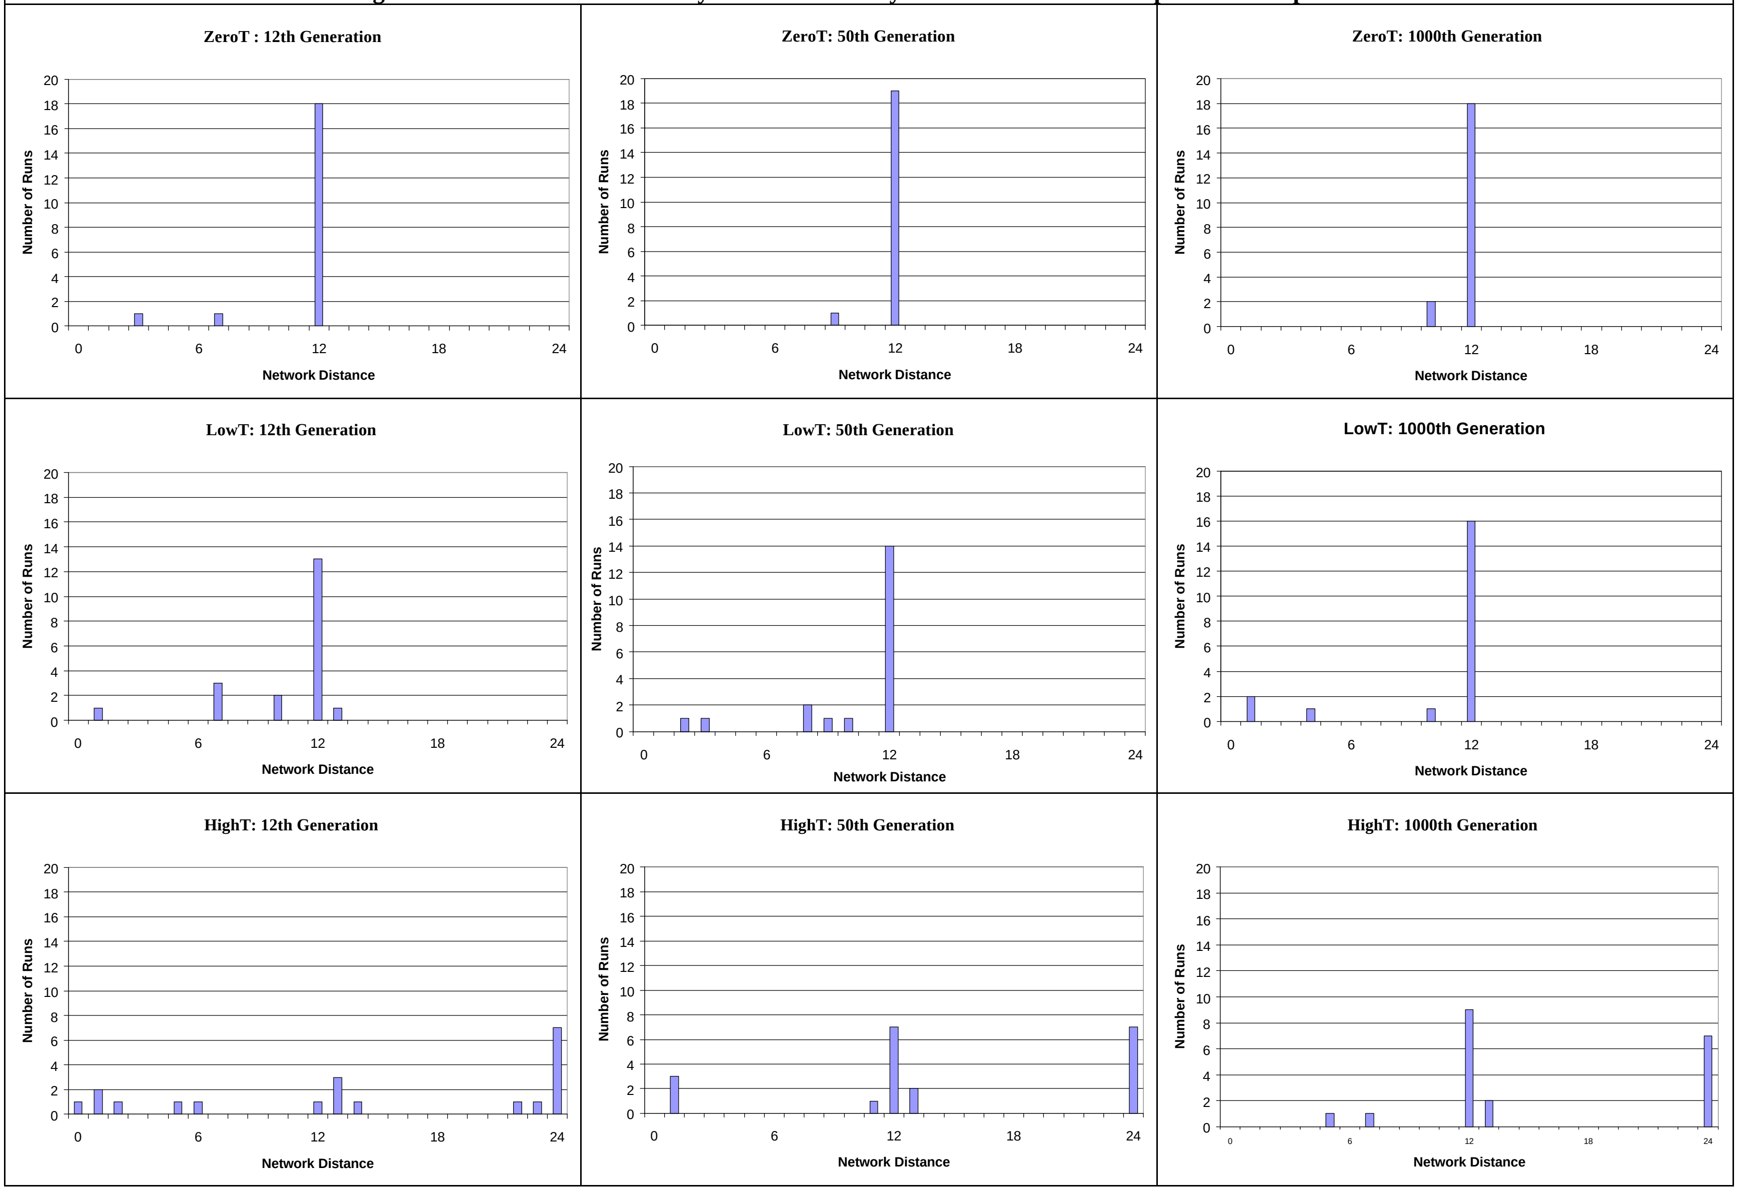

Table 1: Human-Subject Experimental Design			
	Treatment		
	ZeroT	**LowT**	**HighT**
Non-employment Payoff	0	15	30
Sessions per treatment	5	5	5
Workers per session	3	3	3
Employers per session	3	3	3
Trade cycles per session	12	12	12
Matching rounds per trade cycle	Endogenous	Endogenous	Endogenous

Table 2: Human-Subject Network Formation			
	Latched	**Recurrent**	**Transient**
HighT	2.2	28.9	68.9
LowT	13.3	48.9	37.8
ZeroT	6.7	71.1	22.2

Table 3: Determinants of Human-Subject Work-Site Behavior			
	Regression 1	**Regression 2**	**Regression 3**
Unprovoked Defection			
ZEROT	.51**	1.39**	.92
LOWT	.36**	1.26**	.92
HIGHT	1.38**	2.21**	1.46
CYCLE		-.14**	-.14**
GPA			.69**
FAIR			-3.70**
Reciprocal Defection			
ZEROT	.78**	-.14	.72
LOWT	.13	-.81**	.23
HIGHT	.78**	-.11	.37
CYCLE		.12**	.12**
GPA			.26
FAIR			-3.41**
Unprovoked Cooperation			
ZEROT	.11	1.46**	1.63
LOWT	.11	1.50**	1.83
HIGHT	.73**	2.02**	1.93*
CYCLE		-.24**	-.24**
GPA			.39
FAIR			-2.74**
Unemployed Worker			
ZEROT	-2.98**	-3.28**	-6.55**
LOWT	-1.64**	-1.95**	-5.52**
HIGHT	.04	-.25	-3.81**
CYCLE		.04	.04
GPA			2.00**
FAIR			-7.76**
Discouraged Employer			
ZEROT	-4.08**	-4.92	-1.21
LOWT	-2.60**	-3.46	.19
HIGHT	.23	-.58	2.16
CYCLE		.11**	.12**
GPA			-.04
FAIR			-7.68**
Inactivity			
ZEROT	-2.39**	-3.42**	-1.09
LOWT	-.87**	-1.62**	.94
HIGHT	1.41**	.72*	2.51**
CYCLE		.10**	.10**
GPA			-.12
FAIR			-3.88**
R^2	.155	.202	.225

**Significant at 5% level of significance
*Significant at 10% level of significance

Table 4: Human-Subject Shirking Rate and Relationship Types				
	Regression 1	**Regression 2**	**Regression 3**	**Regression 4**
Independent Var.	S	MC	MD	SD
HIGHT	.61**	.14**	.36**	.25**
LOWT	.53**	.24**	.30**	.23**
ZEROT	.38**	.33**	.08	.29**
HIGHT*CYCLE	.011	.002	.023**	-.013*
LOWT*CYCLE	-.012*	.018**	-.005	-.006
ZEROT*CYCLE	.042**	-.022**	.061**	-.019**
R^2	.76	.59	.67	.16

**Significant at 5% level of significance
*Significant at 10% level of significance

Table 5: Computational Experimental Design			
	Treatment		
	ZeroT	**LowT**	**HighT**
Non-Employment Payoff Level	0	15	30
Runs per Treatment	20	20	20
Workers per Run	12	12	12
Employers per Run	12	12	12
Generations per Run	1000	1000	1000
Trade Cycles per Generation	150	150	150
Matching Rounds per Trade Cycle	Endogenous	Endogenous	Endogenous
Generations Sampled	12, 50, 1000	12, 50, 1000	12, 50, 1000

Table 6: 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	31.5	3.5	36.6	4.1	-0.21	-0.09	Mix of latched pairs and recurrent relations
7	45	0	0	0	12	10	6	0	2	2	4	24.6	6.5	39.9	6.2	-0.39	0.00	
Average		**0.00**	**0.00**	**0.00**	**10.00**	**9.50**	**6.50**	**0.00**	**2.00**	**2.50**	**3.50**	**28.1**	**5.0**	**38.3**	**5.2**	**-0.30**	**-0.04**	
%		**0%**	**0%**	**0%**	**83%**	**79%**	**54%**	**0%**	**17%**	**21%**	**29%**							
12	30	0	0	0	11	2	1	0	4	10	7	27.2	9.4	44.1	6.1	-0.32	0.10	Perfect latched pairs with high percentage of AllC Cooperators.
12	50	0	0	2	3	5	6	0	0	7	6	33.2	5.9	41.0	2.3	-0.17	0.03	
12	90	0	0	12	4	2	3	0	0	10	9	35.2	4.5	36.2	3.7	-0.12	-0.09	
12	65	0	0	2	0	0	0	2	0	10	12	41.8	6.6	35.2	9.9	0.04	-0.12	
Average		**0.00**	**0.00**	**4.00**	**4.50**	**2.25**	**2.50**	**0.50**	**1.00**	**9.25**	**8.50**	**34.4**	**6.6**	**39.1**	**5.5**	**-0.14**	**-0.02**	
%		**0%**	**0%**	**33%**	**38%**	**19%**	**21%**	**4%**	**8%**	**77%**	**71%**							
12	60	0	0	7	12	11	11	0	1	1	0	27.1	5.7	38.6	5.3	-0.32	-0.04	Perfect latched pairs with high percentage of Int-Defectors
12	35	0	0	1	12	9	9	1	1	2	2	29.9	5.6	35.3	4.7	-0.25	-0.12	
12	95	0	0	5	12	10	9	0	1	2	2	29.4	5.3	34.6	2.8	-0.27	-0.14	
12	75	0	0	12	12	11	9	0	2	1	1	27.0	7.6	34.1	8.3	-0.33	-0.15	
12	15	0	0	6	11	10	9	1	2	1	1	29.9	6.6	34.1	6.5	-0.25	-0.15	
12	55	0	0	7	6	12	12	0	0	0	0	30.6	3.3	33.8	3.7	-0.24	-0.16	
12	85	0	0	6	8	11	12	0	0	1	0	30.5	3.6	33.7	4.6	-0.24	-0.16	
12	5	0	0	5	12	11	11	1	1	0	0	28.7	5.2	32.9	5.1	-0.28	-0.18	
12	10	0	0	1	10	7	8	2	0	3	4	35.6	7.0	32.0	8.2	-0.11	-0.20	
12	20	0	0	11	8	9	10	2	0	1	2	34.9	5.4	29.7	6.4	-0.13	-0.26	
12	80	0	0	5	8	10	10	1	0	1	2	36.4	7.2	29.5	6.5	-0.09	-0.26	
12	25	0	0	8	5	9	11	2	0	1	1	35.0	6.1	28.3	8.6	-0.13	-0.29	
12	0	0	0	2	11	8	8	4	3	0	1	30.5	8.2	28.0	6.4	-0.24	-0.30	
12	70	0	0	11	12	9	10	3	2	0	0	31.2	5.7	24.1	6.3	-0.22	-0.40	
Average		**0.00**	**0.00**	**6.21**	**9.93**	**9.79**	**9.93**	**1.21**	**0.93**	**1.00**	**1.14**	**31.2**	**5.9**	**32.1**	**6.0**	**-0.22**	**-0.20**	
%		**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**	**31.5**	**5.9**	**34.1**	**5.8**	**-0.21**	**-0.15**	
Total %		**0%**	**0%**	**43%**	**74%**	**69%**	**68%**	**8%**	**9%**	**23%**	**24%**							

Table 7: 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	-0.21	-0.09	Mix of latched pairs and recurrent relations
7	45	7	1	3	2	2	0	2	1	0	-0.39	0.00	
Average		**8.50**	**1.00**	**4.50**	**1.00**	**1.00**	**0.00**	**1.00**	**1.00**	**0.50**	**-0.30**	**-0.04**	
12	30	0	0	7	2	0	0	0	1	2	-0.32	0.10	Perfect latched pairs with high percentage of AllC Cooperators
12	50	5	0	6	0	0	0	0	1	0	-0.17	0.02	
12	90	2	0	9	0	0	0	0	1	0	-0.12	-0.10	
12	65	0	0	10	0	0	0	2	0	0	0.04	-0.11	
Average		**1.75**	**0.00**	**8.00**	**0.50**	**0.00**	**0.00**	**0.50**	**0.75**	**0.50**	**-0.14**	**-0.02**	
12	60	10	0	0	1	0	0	0	1	0	-0.32	-0.04	Perfect latched Pairs with high percentage of Int-Defectors.
12	35	9	1	2	0	0	0	0	0	0	-0.25	-0.12	
12	95	9	2	0	1	0	0	0	0	0	-0.27	-0.14	
12	75	9	0	0	1	1	0	0	0	1	-0.33	-0.15	
12	15	9	1	0	0	1	0	0	0	1	-0.25	-0.15	
12	55	12	0	0	0	0	0	0	0	0	-0.24	-0.16	
12	85	11	0	0	0	0	0	0	1	0	-0.24	-0.16	
12	5	11	1	0	0	0	0	0	0	0	-0.28	-0.18	
12	10	7	0	3	0	0	1	1	0	0	-0.11	-0.20	
12	20	8	0	1	0	1	2	0	0	0	-0.13	-0.26	
12	80	10	0	1	0	0	0	1	0	0	-0.09	-0.26	
12	25	9	0	1	0	0	2	0	0	0	-0.13	-0.29	
12	0	8	3	0	0	0	0	1	0	0	-0.24	-0.30	
12	70	9	2	0	0	0	1	0	0	0	-0.22	-0.40	
Average		**9.36**	**0.71**	**0.57**	**0.21**	**0.21**	**0.43**	**0.21**	**0.14**	**0.14**	**-0.22**	**-0.20**	
Total Average		**7.75**	**0.60**	**2.45**	**0.35**	**0.25**	**0.30**	**0.35**	**0.35**	**0.25**	**-0.21**	**-0.15**	

Table 8: 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	37.6	0.6	39.7	1.0	-0.06	-0.01	Mostly recurrent relations. High percentage of AllC Cooperators.
1	0	0	1	0	1	0	0	0	0	12	11	34.4	0.4	39.5	1.7	-0.14	-0.01	
1	65	0	1	0	1	0	0	0	0	12	11	34.4	0.4	39.5	1.5	-0.14	-0.01	
2	35	0	2	0	2	0	0	0	0	12	10	32.2	0.7	38.9	2.5	-0.20	-0.03	
5	75	0	3	0	5	1	2	0	0	11	7	31.6	0.7	39.5	1.7	-0.21	-0.01	
6	40	4	2	2	12	8	10	0	0	0	0	33.0	1.8	36.2	1.9	-0.18	-0.09	
Average		**0.67**	**1.50**	**0.33**	**3.67**	**1.67**	**2.17**	**0.00**	**0.00**	**9.67**	**8.33**	**33.9**	**0.8**	**38.9**	**1.7**	**-0.15**	**-0.03**	
%		**6%**	**13%**	**3%**	**31%**	**14%**	**18%**	**0%**	**0%**	**81%**	**69%**							
12	5	0	0	0	0	0	0	0	0	12	12	38.9	0.0	40.0	0.0	-0.03	0.00	Mostly latched relations. High percentage of AllC Cooperators.
13	10	1	1	1	0	0	0	0	0	11	11	38.1	2.3	39.2	1.6	-0.05	-0.02	
13	50	1	1	1	0	0	0	0	0	11	11	38.0	2.7	39.2	1.6	-0.05	-0.02	
13	80	1	1	1	0	0	0	0	0	11	11	38.0	2.7	39.2	2.1	-0.05	-0.02	
14	85	3	4	5	4	3	5	0	0	6	3	32.3	2.5	38.6	2.7	-0.19	-0.04	
Average		**1.20**	**1.40**	**1.60**	**0.80**	**0.60**	**1.00**	**0.00**	**0.00**	**10.20**	**9.60**	**37.1**	**2.0**	**39.2**	**1.6**	**-0.07**	**-0.02**	
%		**10%**	**12%**	**13%**	**7%**	**5%**	**8%**	**0%**	**0%**	**85%**	**80%**							
22	20	10	10	12	10	2	2	0	0	0	0	29.9	1.9	30.3	0.2	-0.25	-0.24	Almost complete coordination failure.
23	15	11	11	12	11	1	1	0	0	0	0	29.5	1.6	29.8	0.1	-0.26	-0.26	
24	25	12	12	12	12	0	0	0	0	0	0	29.0	0.0	29.1	0.0	-0.28	-0.27	
24	30	12	12	12	12	0	0	0	0	0	0	29.0	0.0	29.1	0.0	-0.28	-0.27	
24	55	12	12	12	12	0	0	0	0	0	0	29.0	0.0	29.1	0.0	-0.28	-0.27	
24	60	12	12	12	12	0	0	0	0	0	0	29.0	0.0	29.1	0.0	-0.28	-0.27	
24	70	12	12	12	12	0	0	0	0	0	0	29.0	0.0	29.1	0.0	-0.28	-0.27	
24	90	12	12	12	12	0	0	0	0	0	0	29.0	0.0	29.1	0.0	-0.28	-0.27	
24	95	12	12	12	12	0	0	0	0	0	0	29.0	0.0	29.1	0.0	-0.28	-0.27	
Average		**11.67**	**11.67**	**12.00**	**11.67**	**0.33**	**0.33**	**0.00**	**0.00**	**0.00**	**0.00**	**29.2**	**0.4**	**29.3**	**0.0**	**-0.27**	**-0.27**	
%		**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**	**32.5**	**0.9**	**34.7**	**0.9**	**-0.19**	**-0.13**	
Total %		**48%**	**50%**	**49%**	**55%**	**7%**	**9%**	**0%**	**0%**	**45%**	**41%**							

Table 9: 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	-0.06	-0.01	Mostly recurrent Relations. High percentage of AllC Cooperators.
1	0	0	0	79	0	0	0	0	0	0	-0.14	-0.01	
1	65	0	0	80	0	0	0	0	0	0	-0.14	-0.01	
2	35	0	0	99	0	0	0	0	0	0	-0.20	-0.03	
5	75	1	0	60	0	0	0	0	1	0	-0.21	-0.01	
6	40	21	0	0	0	0	0	0	0	0	-0.18	-0.09	
Average		**3.83**	**0.00**	**59.17**	**0.00**	**0.00**	**0.00**	**0.00**	**0.17**	**0.00**	**-0.15**	**-0.03**	
12	5	0	0	12	0	0	0	0	0	0	-0.03	0.00	Mostly latched relations. High percentage of AllC Cooperators.
13	10	0	0	11	0	0	0	0	0	0	-0.05	-0.02	
13	50	0	0	11	0	0	0	0	0	0	-0.05	-0.02	
13	80	0	0	11	0	0	0	0	0	0	-0.05	-0.02	
14	85	3	0	4	0	0	0	0	2	0	-0.19	-0.04	
Average		**0.60**	**0.00**	**9.80**	**0.00**	**0.00**	**0.00**	**0.00**	**0.40**	**0.00**	**-0.07**	**-0.02**	
22	20	2	0	0	0	0	0	0	0	0	-0.25	-0.24	Almost complete coordination failure.
23	15	1	0	0	0	0	0	0	0	0	-0.26	-0.26	
24	25	0	0	0	0	0	0	0	0	0	-0.28	-0.27	
24	30	0	0	0	0	0	0	0	0	0	-0.28	-0.27	
24	55	0	0	0	0	0	0	0	0	0	-0.28	-0.27	
24	60	0	0	0	0	0	0	0	0	0	-0.28	-0.27	
24	70	0	0	0	0	0	0	0	0	0	-0.28	-0.27	
24	90	0	0	0	0	0	0	0	0	0	-0.28	-0.27	
24	95	0	0	0	0	0	0	0	0	0	-0.28	-0.27	
Average		**0.33**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**-0.27**	**-0.27**	
Total Average		**1.45**	**0.00**	**20.20**	**0.00**	**0.00**	**0.00**	**0.00**	**0.15**	**0.00**	**-0.19**	**-0.13**	