

## Hysteresis in an Evolutionary Labor Market with Adaptive Search\*

Leigh Tesfatsion

Department of Economics

Iowa State University, Ames, IA 50011-1070

<http://www.econ.iastate.edu/tesfatsi/>

[tesfatsi@iastate.edu](mailto:tesfatsi@iastate.edu)

**Abstract:** This study undertakes a systematic experimental investigation of hysteresis (path dependency) in an agent-based computational labor market framework. It is shown that capacity asymmetries between work suppliers and employers can result in two distinct hysteresis effects, network and behavioral, when work suppliers and employers interact strategically and evolve their worksite behaviors over time. These hysteresis effects result in persistent heterogeneity in earnings and employment histories across agents who have no observable structural differences. At a more global level, these hysteresis effects are shown to result in a one-to-many mapping between treatment factors and experimental outcomes. These hysteresis effects may help to explain why excess earnings heterogeneity is commonly observed in real-world labor markets.

**Keywords:** Labor markets; network hysteresis; behavioral hysteresis; capacity asymmetries; agent-based computational economics.

\* To appear in S.-H. Chen (ed.), *Evolutionary Computation in Economics and Finance*, Springer-Verlag.

## 1. INTRODUCTION

In the empirical labor economics literature, a labor market is said to exhibit *hysteresis* if temporary shocks appear to have persistent effects on earnings and employment histories.<sup>1</sup> A key concern of empirical labor economists has been the identification of possible propagation mechanisms through which hysteresis might occur.

To date, attention has largely been focused on apparent hysteresis in aggregate unemployment: namely, the protracted effects that unemployment shocks appear to have on the “natural” rate of unemployment. Blanchard and Summers (1990) discuss three distinct types of propagation mechanisms that have been advanced as possible explanations: lag effects arising from the difficulty of adjusting physical capital stocks; long-term labor supply effects arising from the human capital erosion resulting from unemployment; and insider-outsider effects arising from the preferential treatment of actual employees relative to potential employees in the wage bargaining process. Although Blanchard and Summers identify insider-outsider effects as the most promising explanation for hysteresis in European labor markets, they also caution (pp. 270–271) about small sample problems that make this hypothesis difficult to test.

In contrast, this study focuses on a form of hysteresis routinely observed for individual work suppliers and employers in micro panel data: namely, observationally equivalent work suppliers and employers have markedly different earnings and employment histories [see, e.g., Abowd et al. (1999)]. The basic question addressed in this study is whether temporary shocks in the form of idiosyncratic worksite interactions can propagate up into sustained differences in earnings and employment histories for observationally equivalent workers and employers.

---

<sup>1</sup>As pointed out by Piscatelli et al. (1999), the term hysteresis has been used in economic and econometric theory to refer to two distinct phenomena: persistence in deviations from equilibria, possibly followed by an eventual return to a previous equilibrium state; and the presence of unit/zero roots in systems of linear difference or differential equations, implying that a single temporary shock permanently changes the equilibrium path of the system. In empirical economics, however, hysteresis is used more loosely to mean that temporary shocks are observed to result in a persistent change from a previously persistent system state, even though this previously persistent system state cannot be verified to be an equilibrium and the persistent change cannot be verified to be permanent. The latter usage is followed in the current computational study.

Two interdependent aspects of worksite interactions are considered. Who works for whom, and with what regularity? And how do work suppliers and employers behave in these worksite interactions?

In real world labor markets, the behavioral characteristics expressed by work suppliers and employers in their worksite interactions, such as trustworthiness and diligence, depend on who is working for whom. In turn, who is working for whom depends on the behavioral characteristics that have been expressed by work suppliers and employers in their past worksite interactions. Moreover, as stressed in the efficiency-wage literature [Akerloff and Yellon (1986), Yellon (1991)], the behavioral characteristics of work suppliers and employers can also be important determinants of worksite productivity. These behavioral characteristics thus have potentially strong effects on earnings and employment histories. Unfortunately, individual data on the behavioral characteristics of workers and employers are difficult to obtain. The potential effects of these behavioral characteristics are thus usually ignored in micro panel data studies of labor market earnings and employment; typically only observable structural attributes such as training, education, and gender are included as possible explanatory variables.

Using recently developed agent-based programming tools, however, computational labor market frameworks can be constructed in which work suppliers and employers adaptively choose and refuse their potential worksite partners and evolve their worksite behaviors over time on the basis of past worksite interactions. Consequently, the following hypothesis can now be subjected to systematic experimental investigation:

**Worksite Interaction Hysteresis (WIH) Hypothesis:** Temporary shocks in the form of idiosyncratic worksite interactions can result in persistently heterogeneous earnings and employment histories for work suppliers and employers with identical observable structural attributes.

This study investigates the WIH hypothesis in the context of a dynamic computational labor

market framework with strategically interacting work suppliers and employers.<sup>2</sup> As will be clarified below, the labor market framework is a flexible computational laboratory permitting experiments with a wide variety of alternative specifications for the exogenous aspects of market structure and agent attributes. The primary purpose of this study, however, is to take a first cut at the computational study of the WIH hypothesis by specifying these exogenous aspects in relatively simple terms. Thus, as implemented for this study, the labor market framework comprises a fixed equal number of work suppliers and employers. These work suppliers and employers repeatedly participate in costly searches for worksite partners on the basis of continually updated expected utility, engage in efficiency-wage worksite interactions modelled as prisoner's dilemma games, and evolve their worksite strategies over time on the basis of the earnings secured by these strategies in past worksite interactions.

Work suppliers have identical observable structural attributes, and similarly for employers. In particular, each work supplier is assumed to have the same capacity  $wq$ , where  $wq$  is the maximum number of potential work offers that each work supplier can make. Similarly, each employer is assumed to have the same capacity  $eq$ , where  $eq$  is the maximum number of job openings that each employer can provide. Work suppliers and employers are heterogeneous with regard to their worksite strategies. However, a work supplier and employer engaged in a worksite interaction are not able to directly observe each other's strategies; they only observe the behavior and earnings outcomes flowing from the use of these strategies.

The experimental design of the study consists of the systematic variation, from high to low, of *job capacity* as given by the ratio  $eq/wq$ . Jobs are in excess supply when job capacity exceeds one, in balanced supply when job capacity is equal to one, and in tight supply when job capacity is less than

---

<sup>2</sup>This labor market framework was first presented in preliminary fashion in Tesfatsion (1998) as a special case of the Trade Network Game (TNG) model developed in Tesfatsion (1997a,b) for studying the evolution of buyer-seller trade networks. It is used in Tesfatsion (1999) to study market power effects in labor markets. The framework is an example of agent-based computational economics (ACE) modelling. ACE is the computational study of economies modelled as evolving decentralized systems of autonomous interacting agents. For various ACE-related resources, including surveys, readings, software, and pointers to research groups, see the ACE web site at <http://www.econ.iastate.edu/tesfatsi/ace.htm>.

one. For each tested job capacity ratio, twenty different runs are generated using twenty different pseudo-random number seed values.<sup>3</sup> In examining the resulting run histories, particular attention is focused on the experimental determination of correlations between job capacity and the formation of persistent networks among work suppliers and employers, and between network formations and the types of persistent worksite behaviors and earnings outcomes that these networks support.

A key finding of this study is that the WIH hypothesis is strongly supported. In the presence of job capacity asymmetries, idiosyncratic worksite interactions tend to result in persistent network patterns and/or persistent behavioral patterns that support persistently heterogeneous earnings levels across employed work suppliers and across nonvacant employers. These persistent network and behavioral patterns are intermediate hysteresis effects of interest in their own right. It is therefore useful to introduce the following formal definitions:

**Network Hysteresis:** Temporary shocks in the form of idiosyncratic worksite interactions result in persistently heterogeneous network relationships for agents who have identical observable worksite behaviors and structural attributes.

**Behavioral Hysteresis:** Temporary shocks in the form of idiosyncratic worksite interactions result in persistently heterogeneous worksite behaviors for agents who have identical observable structural attributes.

As will be clarified more carefully in subsequent sections, one reason that network hysteresis arises in the labor market framework is that job search is costly. Work suppliers bear the costs of wasted time spent in submitting unsuccessful work offers to employers during the course of job search. In the presence of capacity asymmetries, these sequentially incurred job search costs can induce path-dependent networks among work suppliers and employers that support persistently

---

<sup>3</sup>All labor market experiments reported in this study are implemented using version 105b of the Trade Network Game (TNG) source code developed by McFadzean and Tesfatsion (1999), which in turn is supported by SimBioSys, a general C++ class framework for evolutionary simulations developed by McFadzean (1995). Source code for both the TNG and SimBioSys can be downloaded as freeware at the current author's web site, along with extensive user instructions.

heterogeneous earnings levels across employed work suppliers and across nonvacant employers even when each matched work supplier and employer pair expresses the same type of worksite behavior (e.g., mutual cooperation). Since this earnings heterogeneity arises from a structural asymmetry (e.g., tight job capacity) and not from any deficiency in the worksite strategies of the agents *per se*, it cannot be remedied by evolutionary selection pressures acting upon these strategies.

Behavioral hysteresis arises in the labor market framework for two reasons: differences in own worksite strategies; and differences in the strategies of worksite partners. The first reason is easy to understand. If two work suppliers have different worksite strategies, then in general they will exhibit different worksite behaviors even if they are interacting with the same employer. The second reason is more interesting and stems from the following observation: The behavior an agent expresses in a worksite interaction is a function of the behavior that is expressed by his worksite partner. For example, a single work supplier interacting with two different employers can end up in a mutually cooperative relation with one employer and a completely hostile relation with the other, all triggered by some difference in the employers' expressed behaviors (e.g., one employer initially cooperates and the other initially defects). Thus, even if two work suppliers have identical worksite strategies and are in an identical network pattern with employers (e.g., each is working continuously for one employer), there is no guarantee they will express identical worksite behaviors unless the employers they are interacting with have identical worksite strategies.

As will be clarified below, due to the relatively greater mobility of work suppliers and to evolutionary selection pressures, work suppliers and employers tend to exhibit behavioral hysteresis in their worksite interactions only in conditions of excess job capacity. In conditions of balanced and tight job capacity, the behaviors of the agents within each agent type tend to coordinate rapidly into similar or even identical patterns. On the other hand, neither mobility nor evolutionary selection pressures can eliminate the substantial network hysteresis that tends to arise when there is tight or excess job capacity.

At a more global level, network and behavioral hysteresis result in a one-to-many mapping

between treatment factors and experimental outcomes. That is, for each particular treatment, as the initial random seed value is varied across experimental runs, a small but multiple number of distinct network formations are observed to arise and persist among work suppliers and employers across runs, each supporting a distinct pattern of worksite behaviors and earnings outcomes. This finding is consistent with the many analytical two-sided labor market studies, such as Diamond (1982), that establish the existence of multiple steady-state search equilibria. In the current process study, however, a histogram is obtained for each treatment showing the proportion of runs that evolve each type of network formation, which provides suggestive information regarding the size and importance of their basins of attraction.

The labor market framework is described in Section 2. In Section 3, descriptive statistics are constructed for the ex post classification of network formations, worksite behaviors, and earnings outcomes. The experimental design of the study is outlined in Section 4, and a detailed discussion of experimental findings is presented in Section 5. Concluding remarks are given in Section 6.

## **2. Labor Market Framework**

The labor market framework differs in several essential respects from standard labor market models. First, it is a dynamic process model defined algorithmically in terms of the internal states and behavioral rules of work suppliers and employers rather than by the usual system of demand, supply, and equilibrium equations. The only equations that arise in the model are those used by the agents themselves to summarize observed aspects of their world and to implement their behavioral rules. Second, agents attempt to learn about the behavioral rules of other agents even as these rules are coevolving over time. Third, starting from given initial conditions, all events are contingent on agent-initiated interactions and occur in a path-dependent time line. The analogy to a culture growing in a petri dish, observed by an interested researcher but not disturbed, is apt.

The labor market framework comprises an equal number  $M$  of work suppliers who make work offers and employers who receive work offers, where  $M$  can be any positive integer. Each work

supplier can have work offers outstanding to no more than  $wq$  employers at any given time, and each employer can accept work offers from no more than  $eq$  work suppliers at any given time, where the work offer quota  $wq$  and the employer acceptance quota  $eq$  can be any positive integers.<sup>4</sup>

As seen in Table 1, work suppliers and employers are modelled as autonomous endogenously-interacting agents with internalized social norms, internally stored state information, and internal behavioral rules. Each agent, whether a work supplier or an employer, has this same general internal structure. However, work suppliers differ from employers in terms of their specific market protocols, fixed attributes, and initial endowments; and all agents can acquire different state information and evolve different worksite behavioral rules<sup>5</sup> over time on the basis of their past experiences. Note, in particular, that all agents have stored addresses for other agents together with internalized market protocols for communication. These features permit agents to communicate state-dependent messages to other agents at event-triggered times, a feature not present in standard economic models. As will be clarified below, the work suppliers and employers depend on this communication ability to seek out and secure worksite partners on an ongoing adaptive basis.

**Insert Table 1 about here**

As outlined in Table 2, activities in the labor market framework are divided into a sequence of *generations*. Each work supplier and employer in the initial generation is assigned a randomly generated rule governing his worksite behavior together with initial expected utility assessments regarding potential worksite partners. The work suppliers and employers then enter into a *trade cycle loop* during which they repeatedly search for worksite partners on the basis of their current expected utility assessments, engage in efficiency-wage worksite interactions modelled as prisoner's dilemma games, and update their expected utility assessments to take into account newly incurred job search costs and worksite payoffs. At the end of the trade cycle loop, the work suppliers and

---

<sup>4</sup>When  $wq$  exceeds 1, each work supplier can be interpreted as some type of information service provider (e.g., broker or consultant) that is able to supply services to at most  $wq$  employers at a time or as some type of union organization that is able to oversee work contracts with at most  $wq$  employers at a time.

<sup>5</sup>In principle, agents could evolve any or all of their behavioral rules, but for current study purposes only the evolution of worksite behavioral rules is considered.

employers each separately evolve (structurally modify) their worksite behavioral rules based on the past utility outcomes secured with these rules, and a new generation commences.

**Insert Table 2 about here**

The particular module specifications used in all experiments reported below will now be described in roughly the order depicted in Table 2.<sup>6</sup>

Matches between work suppliers and employers are determined using a one-sided offer auction, a modified version of the “deferred acceptance mechanism” originally studied by Gale and Shapley (1962).<sup>7</sup> Under the terms of this auction, hereafter referred to as the *deferred choice and refusal* (DCR) mechanism, each work supplier submits work offers to a maximum of  $wq$  employers he ranks as most preferable on the basis of expected utility and who he judges to be tolerable in the sense that their expected utility is not negative. Similarly, each employer selects up to  $eq$  of his received work offers that he finds tolerable and most preferable on the basis of expected utility and he places them on a waiting list; all other work offers are refused. Work suppliers redirect refused work offers to tolerable preferred employers who have not yet refused them, if any such employers exist. Once employers stop receiving new work offers, they accept all work offers currently on their waiting lists.

A work supplier incurs a job search cost in the form of a negative *refusal payoff*  $R$  each and every time that an employer refuses one of his work offers during a trade cycle; the employer who does the refusing is not penalized.<sup>8</sup> A work supplier or employer who neither submits nor accepts

---

<sup>6</sup>All experiments reported in this paper are implemented using version 105b of the Trade Network Game (TNG) source code developed by McFadzean and Tesfatsion (1999). The latter study provides a detailed discussion of all module implementations. In addition, the TNG source code (with extensive comment statements and user instructions) can be downloaded as freeware from the current author’s web site, permitting all module implementations to be specifically viewed in source code form.

<sup>7</sup>See Roth and Sotomayor (1990) for a careful detailed discussion of Gale-Shapley deferred acceptance matching mechanisms, including a discussion of the way in which the Association of American Medical Colleges since WWII has slowly evolved such an algorithm (the National Intern Matching Program) as a way of matching interns to hospitals in the United States.

<sup>8</sup>This modelling for job search costs is equivalent to assuming: (i) each work supplier must pay a job search cost in amount  $-R$  for each work offer he makes to an employer; and (ii) each possible worksite payoff for work suppliers is increased by the amount  $-R$ , so that a work supplier is able to recoup the job search costs he incurs in making a work offer if and only if this work offer is accepted.

work offers during a trade cycle receives an *inactivity payoff* 0 for the entire trade cycle. The refusal and inactivity payoffs are each assumed to be measured in utility terms.

If an employer accepts a work offer from a work supplier in any given trade cycle, the work supplier and employer are said to be *matched* for that trade cycle. Each match constitutes a mutually agreed upon contract stating that the work supplier shall supply labor services at the worksite of the employer until the beginning of the next trade cycle. These contracts are risky in that outcomes are not assured.

Specifically, each matched work supplier and employer engage in a worksite interaction modelled as a two-person prisoner's dilemma game reflecting the basic efficiency wage hypothesis that work effort levels are affected by overall working conditions (e.g., wage levels, respectful treatment, safety considerations). The work supplier can either cooperate (exert high work effort) or defect (engage in shirking). Similarly, the employer can either cooperate (provide good working conditions) or defect (provide substandard working conditions).

The range of possible worksite payoffs is assumed to be the same for each worksite interaction in each trade cycle: namely, as seen in Table 3, a cooperator whose worksite partner defects receives the lowest possible payoff  $L$  (sucker payoff); a defector whose worksite partner also defects receives the next lowest payoff  $D$  (mutual defection payoff); a cooperator whose worksite partner also cooperates receives a higher payoff  $C$  (mutual cooperation payoff); and a defector whose worksite partner cooperates receives the highest possible payoff  $H$  (temptation payoff).

**Insert Table 3 about here**

The worksite payoffs in Table 3 are assumed to be measured in utility terms, and to be normalized about the inactivity payoff 0 so that  $L < D < 0 < C < H$ . Thus, a work supplier or employer that ends up either as a sucker with payoff  $L$  or in a mutual defection relation with payoff  $D$  receives negative utility, a worse outcome than inactivity (unemployment or vacancy). These worksite payoffs are also assumed to satisfy the usual prisoner's dilemma regularity condition  $(L + H)/2 < C$

guaranteeing that mutual cooperation dominates alternating cooperation and defection on average.

Each agent, whether a work supplier or an employer, uses a simple learning algorithm to update his expected utility assessments on the basis of new payoff information. Specifically, an agent  $v$  assigns an exogenously given initial expected utility  $U^o$  to each potential worksite partner  $z$  with whom he has not yet interacted. Each time an interaction with  $z$  takes place,  $v$  forms an updated expected utility assessment for  $z$  by summing  $U^o$  together with all payoffs received to date from interactions with  $z$  (including both worksite payoffs and refusal payoffs) and then dividing this sum by one plus the number of interactions with  $z$ .

The rule governing the worksite behavior of each agent, whether work supplier or employer, is represented as a finite-memory pure strategy for playing a prisoner's dilemma game with an arbitrary partner an indefinite number of times, hereafter referred to as a *worksite strategy*. At the commencement of each trade cycle loop, agents have no information about the worksite strategies of other agents; they can only learn about these strategies by engaging other agents in repeated worksite interactions and observing the behavioral and utility outcomes that ensue. In consequence, each agent's choice of an action in a current worksite interaction with another agent is determined entirely on the basis of his own past interactions with this other agent plus his initial expected utility assessment of the agent. Each agent thus keeps separate track of his interaction history with each potential worksite partner.

At the end of each trade cycle loop, the *utility (fitness)* of each work supplier and employer is measured by the average payoff he attained over this trade cycle loop. Average payoff is calculated as total net payoffs (negative refusal payoffs plus worksite payoffs) divided by the total number of payoffs received. The worksite strategies of workers and employers are then separately evolved by means of standardly specified genetic algorithms involving recombination, mutation, and elitism operations that are biased in favor of more fit agents.<sup>9</sup> This evolution is meant to reflect the

---

<sup>9</sup>More precisely, for each agent type (work supplier or employer), the genetic algorithm evolves a new collection of agent worksite strategies from the existing collection of agent worksite strategies by applying the following four steps: (1) *Evaluation*, in which a fitness score is assigned to each strategy in the existing strategy collection; (2)

formation and transmission of new ideas by mimicry and experimentation, not reproduction in any biological sense. That is, if a worksite strategy successfully results in high fitness for an agent of a particular type, then other agents of the same type are led to modify their own strategies to more closely resemble the successful strategy.

An important caution is in order here, however. The information that work suppliers and employers are currently permitted to have access to in the evolution step is substantial: namely, complete knowledge of the collection of strategies used by agents of their own type in the previous trade cycle loop, ranked by fitness. The evolution step is thus more appropriately interpreted as an iterative stochastic search algorithm for determining possible strategy configuration attractors rather than as a social learning mechanism per se. The resulting outcomes will be used in subsequent work as a yardstick against which to assess the performance of more realistically modelled social learning mechanisms.

### 3. DESCRIPTIVE STATISTICS

Each of the labor market experiments reported in this study results in a one-to-many mapping between structural characteristics and outcomes. That is, when each particular experimental treatment is repeated for a range of pseudo-random number seed values, a distribution of behavioral, network, and utility outcomes is generated. Consequently, the mapping between treatment factors and outcomes must be characterized statistically.

This section explains the descriptive statistics that have been constructed to aid in the experimental determination of correlations between treatment factors and network formations, and between network formations and the types of worksite behaviors and utility outcomes that these

---

*Recombination*, in which offspring (new ideas) are constructed by combining the genetic material (structural characteristics) of pairs of parent strategies chosen from among the most fit strategies in the existing strategy collection; (3) *Mutation*, in which additional variations (new ideas) are constructed by mutating the structural characteristics of each offspring strategy with some small probability; and (4) *Replacement*, in which the most fit (elite) strategies in the existing strategy collection are retained for the new collection of strategies and the least fit strategies in the existing strategy collection are replaced with offspring strategies. See McFadzean and Tesfatsion (1999) for a more detailed discussion of this use of genetic algorithms in the Trade Network Game (TNG), and see Sargent (1993) for a more general discussion of genetic algorithm design and use.

networks support. Networks depict who is working for whom, and with what regularity. Worksite behavior refers to the specific actions undertaken by workers and employers in their worksite interactions. Finally, utility refers to the average payoff levels attained by work suppliers and employers as a result of job search and worksite interactions.

### 3.A *Classification of Contractual Networks by Distance*

First introduced is a distance measure on persistent networks that permits the classification of these networks into alternative types. This distance measure calculates the extent to which an observed pattern of persistent agent relationships deviates from an idealized pattern that specifies relationships among agent types without consideration for the identity of individual agents within agent types. As will be seen in Section 5, this distance measure permits networks to be distinguished on the basis of the differential worksite behaviors and utility outcomes that they support. In addition, as a by-product, it provides a useful indicator of the extent to which heterogeneity in attained utility levels arises from network hysteresis.

All labor market experiments reported in this study were implemented using version 105b of the TNG source code developed by McFadzean and Tesfatsion (1999). Let  $s$  denote a seed value for the pseudo-random number generator incorporated in this source code, and let  $E$  denote a *potential economy*, i.e., an economy characterized structurally by the source code together with specific values for all source code parameters<sup>10</sup> apart from  $s$ . The *sample economy* generated from  $E$ , given the seed value  $s$ , is denoted by  $(s, E)$ .

Worksite strategies are represented as finite state machines,<sup>11</sup> hence the actions undertaken by any agent  $v$  in repeated worksite interactions with another agent  $z$  must eventually cycle. Consequently, these actions can be summarized in the form of a *worksite history*  $H:P$ , where the

---

<sup>10</sup>A complete annotated listing of these source code parameters is given in Section 4 (Table 4), below.

<sup>11</sup>A *finite state machine* is a system comprising a finite collection of internal states together with a state transition function that gives the next internal state the system will enter as a function of the current state and other current inputs to the system. For the application at hand, the latter inputs are the actions selected by a worker and an employer engaged in a worksite interaction. See McFadzean and Tesfatsion (1999) for a more detailed discussion and illustration of how finite state machines are used to represent worksite strategies in the TNG source code.

*handshake*  $H$  is a (possibly null) string of worksite actions that form a non-repeated pattern and the *persistent portion*  $P$  is a (possibly null) string of worksite actions that are cyclically repeated. For example, letting  $c$  denote cooperation and  $d$  denote defection, the worksite history  $ddd:dc$  indicates that agent  $v$  defected against agent  $z$  in his first three worksite interactions with  $z$  and thereafter alternated between defection and cooperation.

A work supplier  $w$  and employer  $e$  are said to exhibit a *persistent relationship* during a given trade cycle loop  $T$  of a sample economy  $(s, E)$  if the following two conditions hold: (a) their worksite histories with each other during the course of  $T$  take the form  $H_w:P_w$  and  $H_e:P_e$  with nonnull  $P_w$  and  $P_e$ ; and (b) accepted work offers between  $w$  and  $e$  do not permanently cease during  $T$  either by choice (a permanent switch to strictly preferred partners) or by refusal (one agent becoming intolerable to the other because of too many defections). A persistent relationship between  $w$  and  $e$  is said to be *latched* if  $w$  works for  $e$  continuously (in each successive trade cycle), and it is said to be *recurrent* if  $w$  works for  $e$  randomly or periodically.

A possible pattern of relationships among the work suppliers and employers in the final generation of a potential economy  $E$  is referred to as a *network*, denoted generically by  $K(E)$ . Each network  $K(E)$  is represented in the form of a directed graph in which the vertices  $V(E)$  of the graph represent the work suppliers and employers, the edges of the graph (directed arrows) represent work offers directed from work suppliers to employers, and the edge weight on any edge denotes the number of accepted work offers between the work supplier and employer connected by the edge.

Let  $K(s, E)$  denote the network depicting the actual pattern of relationships among the work suppliers and employers in the final generation of the sample economy  $(s, E)$ . The reduced form network  $K^p(s, E)$  derived from  $K(s, E)$  by eliminating all edges of  $K(s, E)$  that correspond to nonpersistent relationships is referred to as the *persistent network* for  $(s, E)$ .

Let  $V^o(E)$  denote a *base network pattern* that partially or fully specifies a potential pattern of relationships among the work suppliers and employers in the potential economy  $E$  by placing general

constraints on the relationships among agent types without regard for the individual identity of agents within each type. For example,  $V^o(E)$  could designate that each work supplier directs work offers to at least two employers. The collection of all networks whose edges conform to the base network pattern  $V^o(E)$  is referred to as the *base network class*, denoted by  $K^o(E)$ .

The *distance*  $D^o(s, E)$  between the persistent network  $K^p(s, E)$  and the base network class  $K^o(E)$  for a sample economy  $(s, E)$  is then defined to be the number of vertices (work suppliers and employers) for  $K^p(s, E)$  whose edges (persistent relationships) fail to conform to the base network pattern  $V^o(E)$ . As will be demonstrated in Section 5, this distance measure provides a useful way to classify the different types of persistent networks observed to arise for a given value of  $E$  as the seed value  $s$  is varied.

### 3.B Classification of Worksite Behaviors and Utility Outcomes

Let a sample economy  $(s, E)$  be given. A work supplier or employer in the final generation of  $(s, E)$  is referred to as an *aggressive agent* if he engages in at least one defection against another agent that has not previously defected against him. The  $1 \times 2$  vector giving the percentages of work suppliers and employers in the final generation of  $(s, E)$  that are aggressive is referred to as the *aggressive profile* for  $(s, E)$ . The aggressive profile measures the extent to which work suppliers and employers behave opportunistically in worksite interactions with partners who are either strangers<sup>12</sup> or who so far have been consistently cooperative.

A work supplier or employer in the final generation of  $(s, E)$  is referred to as *persistently inactive* if he constitutes an isolated vertex of the persistent network  $K^p(s, E)$ . The  $1 \times 2$  vector giving the percentages of work suppliers and employers in the final generation of  $(s, E)$  who are persistently inactive is referred to as the *p-inactive profile* for  $(s, E)$ . The p-inactive profile measures the extent to which work suppliers and employers in this final generation fail to establish any persistent rela-

---

<sup>12</sup>The importance of stance toward strangers and first impressions for determining subsequent outcomes in path dependent contexts such as the labor market framework has been stressed by Orbell and Dawes (1993) and by Rabin and Schrag (1999).

tionships. The p-inactive percentage for work suppliers constitutes their persistent unemployment rate, whereas the p-inactive percentage for employers constitutes their persistent vacancy rate.

A work supplier or employer in the final generation of  $(s, E)$  is referred to as a *repeat defector* if he establishes at least one persistent relationship for which the persistent portion  $P$  of his worksite history  $H:P$  includes a defection  $d$ . Defections for work suppliers correspond to shirking episodes, and defections for employers correspond to the provision of poor working conditions. The  $1 \times 2$  vector giving the percentages of work suppliers and employers in the final generation of  $(s, E)$  who are repeat defectors is referred to as the *r-defector profile* for  $(s, E)$ . The r-defector profile measures the extent to which work suppliers and employers in the final generation of  $(s, E)$  engage in recurrent or continuous defections.

If, instead, a work supplier or employer in the final generation of  $(s, E)$  establishes at least one persistent relationship and his worksite history for each of his persistent relationships has the general form  $H:c$ , he is referred to as *persistently nice*. The  $1 \times 2$  vector giving the percentages of work suppliers and employers in the final generation of  $(s, E)$  who are persistently nice is referred to as the *p-nice profile* for  $(s, E)$ . The p-nice profile measures the extent to which work suppliers and employers in this final generation establish persistent relationships characterized by fully cooperative behavior.

By construction, each work supplier and employer in the final generation of  $(s, E)$  must either be a persistently inactive agent, a repeat defector, or a persistently nice agent.

Finally, the  $1 \times 2$  vector giving the average utility (fitness) levels attained by work suppliers and employers in the final generation of  $(s, E)$  is referred to as the *utility profile* for  $(s, E)$ . The utility profile measures the distribution of welfare across agent types.

#### 4. EXPERIMENTAL DESIGN

The labor market experiments reported in Section 5 are for two-sided markets comprising 12 work suppliers and 12 employers. Each work supplier has the same offer quota,  $wq$ , and each employer

has the same acceptance quota,  $eq$ . Attention is focused on the effects of varying job capacity from high to low, where job capacity is measured by the ratio  $eq/wq$ . Four different settings for job capacity are tested: high excess job capacity ( $eq \gg wq$ ); balanced job capacity ( $eq = wq = 1$ ); tight job capacity ( $eq = 1$  and  $wq = 2$ ); and extremely tight job capacity ( $eq \ll wq$ ).

The values for all remaining parameters are maintained at fixed values throughout all experiments. Table 4 lists these fixed parameter values along with the specific  $wq$  and  $eq$  quota values for an experiment with high excess job capacity. The parameter values in Table 4, together with the TNG source code, constitute a potential economy  $E$  in the sense defined in the previous section.

**Insert Table 4 About Here**

For each tested  $E$ , twenty sample economies  $(s, E)$  were generated using twenty arbitrarily selected seed values  $s$  for the pseudo-random number generator included in the TNG source code.<sup>13</sup> For each run  $s$ , the persistent network  $K^p(s, E)$  was determined and graphically depicted, and the components for the four behavioral profiles (aggressive, p-inactive, r-defector, and p-nice) and the utility profile were calculated and recorded.

A base network pattern  $V^o(E)$  was then specified for each tested economy  $E$  that constrained the general relationships among agent types without relying on the individual identity of agents within each agent type. This base network pattern provides the 0 point for the distance measure  $D^o(\cdot, E)$  and hence is an intrinsically arbitrary normalization. However, its degree of specificity governs the dispersion of the resulting distance values  $D^o(s, E)$  across sample runs  $s$  and the extent to which these distance values display useful correlations with the components of the behavioral and utility profiles. In practice, then, the choice of the base network pattern for each tested economy was fine-tuned so that the resulting distance values provided a classification of networks into distinct types supporting distinct patterns of worksite behaviors and utility outcomes.

---

<sup>13</sup>These twenty seed values are as follows: 5, 10, 15, 20, 25, 30, 45, 65, 63, 31, 11, 64, 41, 66, 13, 54, 641, 413, 425, and 212. The final fourteen values were determined by random throws of two and three die. The TNG source code used to implement the labor market framework uses pseudo-random number values in the initialization of worksite strategies, in the matching process to break ties among equally preferred worksite partners, and in genetic algorithm recombination and mutation operations applied to worksite strategies in the evolution step.

Given  $V^o(E)$ , the distance  $D^o(s, E)$  of  $K^p(s, E)$  from  $K^o(E)$  was recorded for each run  $s$ , and a histogram for the distance values  $D^o(s, e)$  was constructed giving the percentage of runs  $s$  corresponding to each possible distance value. Finally, as a rough stability check, the number of generations was also increased to 100 for each tested economy  $E$  and the minimum, maximum, and average values for the utility levels attained by work suppliers and employers in each of the 100 generations were graphically generated for each sample economy  $(s, E)$ .

## 5. EXPERIMENTAL FINDINGS

Consider a two-sided potential economy  $E$  comprising 12 work suppliers and 12 employers with a work offer quota  $wq = 1$  and an employer acceptance quota  $eq = 12$ . These quota values imply there is high excess job capacity. Employers are forced to remain vacant unless work suppliers happen to direct work offers their way, hence employers face a substantial structural risk of vacancy. On the other hand, work suppliers face zero structural risk of having their work offers refused by employers because of limited job capacity.

As depicted in Figure 1(a), the base network pattern  $V^o(E)$  for this high excess job economy  $E$  is as follows: Each work supplier is latched to an employer, and each employer has at least one latched work supplier.

— **Insert Figure 1 About Here** —

Descriptive statistics for the twenty sample economies  $(s, E)$  corresponding to this high excess job economy  $E$  are presented in Table 5(a).<sup>14</sup> Note that 75% of the sample economies  $(s, E)$  lie in the distance cluster 3–9. The low mean utility level 0.35 attained by employers in this distance cluster is due to three factors: a high mean p-inactivity (vacancy) rate among employers due to high excess job capacity; a high mean aggression (initial defection) rate by work suppliers; and a

---

<sup>14</sup>In Table 5, for each distance cluster, the mean and standard deviation are calculated for each component of the three behavioral profiles (aggressive, p-inactive, and p-nice) and the utility profile across the sample runs lying in this distance cluster. Since the r-defector profiles can be derived from the p-inactive and p-nice profiles (see Section 3.B), they are omitted.

low mean p-nice (cooperation) rate by work suppliers that induces retaliatory defections by some employers.

— **Insert Table 5 About Here** —

The persistent networks that arise for the sample economies  $(s, E)$  in distance cluster 3–9 reveal strong network hysteresis, i.e., strong persistent differences in relationship patterns for both work suppliers and employers. The typical scenario is as follows: r-defector work suppliers latch on to a proper subset of p-nice employers, with anywhere from one to four work suppliers latched to the same employer, and drive down the utility levels of these employers to small positive values. Remaining employers are left vacant with utility levels at zero (the inactivity payoff). This scenario ensures that the worksite strategies of the p-nice employers are advantaged in the evolution step relative to the worksite strategies of the employers who are left vacant. Since work suppliers and employers evolve separately, the worksite strategies of the p-nice employers tend to reproduce into the next generation, which ensures the perpetuation of a cooperative set of employers whom the work suppliers can continue to opportunistically defect against.

At least some degree of persistent heterogeneity in utility levels across employed work suppliers and across nonvacant employers is observed for each sample economy  $(s, E)$  in the distance cluster 3–9. This persistent heterogeneity in utility levels across the active agents of each agent type is primarily due to behavioral hysteresis – specifically, persistent differences in worksite behaviors between different latched work supplier and employer pairs – rather than to network hysteresis. The different distance values observed for the persistent networks arising in these sample economies are essentially a count of the number of persistently vacant employers who have degenerated into p-inactivity primarily by bad luck but also occasionally by ostracism.

Table 5(a) also shows that the remaining 25% of the sample economies for this  $E$  lie in a second distance cluster  $D^o = 24$ . The mean utility level 1.02 attained by employers in this second distance cluster is much higher than that attained in distance cluster 3–9 due to the high mean percentage

of p-nice behavior exhibited by both work suppliers and employers. This mean utility level is nevertheless substantially below the mutual cooperation payoff level 1.40 due to the 5% mean p-inactivity (vacancy) rate among employers, a structural consequence of high excess job capacity that is independent of how cooperatively the employers behave in their worksite interactions. The typical pattern exhibited in this distance cluster is p-nice work suppliers randomly directing work offers among employers without latching. Note that the mean utility level 1.39 attained by work suppliers is very close to the mutual cooperation payoff level 1.40. No latching takes place in any of the sample economies in this distance cluster, and utility levels are largely homogeneous across employed work suppliers and across nonvacant employers.

Next consider the case in which the work offer quota remains at  $wq = 1$  but the employer acceptance quota is reduced to  $eq = 1$  so that job capacity is balanced. This change dramatically affects network formation.

Specifically, as depicted in Figure 1(b), the base network pattern now consists of disjoint latched pairings of one work supplier and one employer. As detailed in Table 5(b), 75% of the sample economies for this balanced job capacity economy  $E$  lie in distance cluster 0–2, implying that the base network pattern is by far the most predominant network formation observed. Heterogeneity in utility outcomes across employed work suppliers is essentially due to differences in job search costs incurred in the process of attaining the coordinated base network pattern. Once in this coordinated state, very little further heterogeneity in utility outcomes is observed either across employed work suppliers or across nonvacant employers.

More precisely, balanced job capacity favors employers over work suppliers, because work suppliers must bear the costs associated with job search. Nevertheless, the endogenous mobility of work suppliers protects them from overly opportunistic worksite behavior by employers. An employer who attempts to sustain too high a defection frequency against a work supplier will cause this work supplier to quit (redirect his future work offers elsewhere) if other employers are perceived as better earnings opportunities, or even to exit the labor force altogether. Although defections

by employers occur rather frequently in the handshake portions of worksite histories, in all but one of the sample economies the employers end up expressing largely p-nice behavior rather than attempting to exploit the work suppliers' vulnerability to job search costs by engaging in repeat defections.

On the other hand, work suppliers who fail to latch when job capacity is balanced tend to accumulate large job search costs (negative refusal payoffs). Being fired by an employer for aggressive or r-defection behavior can thus be very costly for work suppliers, and most display p-nice behavior in their worksite interactions. Nevertheless, even work suppliers who succeed in established a mutually p-nice latched relationship with a single employer typically accumulate two or three negative refusal payoffs from a wide range of employers on their way to attaining this coordinated state. These job search costs, together with the aggressive (initial defection) behavior of many employers, tend to lower the mean utility level of work suppliers relative to employers.

The stability checks conducted for this balanced job capacity case reveal that many of the sample economies exhibit unsettled mean utility outcomes over generations 1 through 100 in the form of persistent drifting, bubbling, or regime shifts. The reason for this appears to be that, with balanced job capacity, networks form in response to job search costs, yet they support largely p-nice or even *c:c* worksite behavior. In consequence, these networks are not robust to the entrance of new, initially cooperative worksite strategies introduced in the evolution step.

As job capacity keeps tightening, work suppliers have an increasingly difficult time forming persistent relationships with employers, a finding indicated in Figure 1 by the decreasing size of work supplier boxes relative to employer boxes as one moves from part (a) to part (d). This increased coordination failure is detailed in Table 5. Note, in particular, the growing mean percentage of work suppliers who become p-inactive (unemployed) as job capacity successively tightens.

More precisely, with tight job capacity, the typically observed experimental outcome is that each employer forms persistent relationships with a particular subset of work suppliers. These persistent relationships are recurrent in the following sense: In each trading period, the work offers received

by the employer from his persistent work suppliers exceed his job capacity limits, and he accepts only a portion of these work offers by random selection. The reason for the random selection is that these persistent work suppliers tend to be largely cooperative in their interactions with their employers, so that the employers are generally indifferent regarding whose work offers to accept. Since job openings are relatively scarce, this implies that work suppliers face a risk that their work offers will be refused by employers due to capacity limitations even if they have never defected against any employer in their past worksite interactions. On the other hand, for reasons elaborated above for the case of balanced job capacity, the endogenous mobility of work suppliers still induces largely p-nice behavior among employers.

A work supplier whose work offer is refused by an employer incurs a job search cost. This causes the work supplier to lower the utility he expects to attain from any next work offer to this employer, which in turn encourages the work supplier to direct his next work offer elsewhere. A work supplier who receives too many refusals from employers eventually ceases making work offers altogether because the expected utility he assigns to each prospective employer falls below zero, the inactivity payoff level. This complete discouragement tends to occur for work suppliers in the early stages of a trade cycle loop when work suppliers are spreading their work offers among many different employers and refusal rates tend to be high.

As discouraged work suppliers leave the labor force, however, refusal rates decline and the condition of the remaining work suppliers improves. Consequently, even when job capacity is extremely tight, over half the sample economies manage to evolve to a sustainable state in which work suppliers who remain in the labor force are able to find employment at a high enough frequency to sustain their utility levels at positive levels.

Heterogeneity in utility outcomes across employed work suppliers in conditions of tight and extremely tight job capacity thus principally arises from two sources. First, work suppliers experience differential numbers of accepted work offers, resulting in differential worksite payoffs. Second, work suppliers experience differential numbers of refused work offers, resulting in differential job search

costs.

These differences in worksite payoffs and job search costs result from network hysteresis. Job search costs incurred by chance determine a network of persistent relationships among work suppliers and employers in a highly path-dependent way. Some work suppliers manage to establish recurrent persistent relationships with so many employers that they are essentially guaranteed to achieve full employment in each trade cycle, whereas other “underemployed” work suppliers only manage to place a few of their potential work offers in each trade cycle. This network hysteresis supports persistent heterogeneity in utility outcomes across active work suppliers. Since this heterogeneity is fundamentally caused by a structural asymmetry (too few job openings) and not from differences in worksite strategies per se, it cannot be remedied by evolutionary selection pressures acting upon worksite strategies.

## 6. CONCLUDING REMARKS

Despite the structural simplicity of the computational labor market framework used in this study, the reported experimental findings highlight interesting network and behavioral hysteresis effects arising from idiosyncratic worksite interactions that may be important for understanding aspects of real-world labor markets. For example, it is seen that these hysteresis effects tend to support earnings outcomes that are more heterogeneous than would be predicted on the basis of the observable structural attributes of work suppliers and employers. This phenomenon is routinely observed in real-world labor markets, and is referred to as the “excess heterogeneity problem” [Abowd et al. (1999)]. The current study demonstrates the feasibility and potential usefulness of investigating network and behavioral hysteresis effects in an agent-based computational framework.

## REFERENCES

- Abowd, John M., Francis Kramarz, and David N. Margolis (1999) High wage workers and high wage firms. *Econometrica* 67, 251–333.

- Akerloff, George, and Janet Yellen, Eds. (1986) *Efficiency Wage Models of the Labor Market*, Cambridge University Press, Cambridge, UK.
- Blanchard, Olivier J., and Lawrence H. Summers (1990) Hysteresis and the European Unemployment Problem, In Lawrence H. Summers (ed.), *Understanding Unemployment*, The MIT Press, Cambridge, MA, Chapter 8 (pp. 227–285).
- Diamond, Peter (1982) Aggregate demand management in search equilibrium. *Journal of Political Economy* 90, pp. 881–895.
- Gale, David, and Lloyd Shapley (1962) College admissions and the stability of marriage. *American Mathematical Monthly* 69, pp. 9–15.
- McFadzean, David (1995) *SimBioSys: A Class Framework for Evolutionary Simulations*. Master’s Thesis, Department of Computer Science, University of Calgary, Alberta, Canada.
- McFadzean, David, and Leigh Tesfatsion (1999) A C++ platform for the evolution of trade networks. *Computational Economics* 14, 109–134.
- Orbell, John M., and Robin M. Dawes, 1993. Social welfare, cooperators’ advantage, and the option of not playing the game. *American Sociological Review* 58, 787–800.
- Piscitelli, Laura, Rod Cross, Michael Grinfeld, and Harbir Lamba (1999), A test for strong hysteresis, *Computational Economics*, to appear.
- Rabin, Matthew, and J. Schrag, 1999. First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics* 114, 37–82.
- Roth, Alvin, and Marilda A. O. Sotomayor (1990) *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge, UK: Cambridge University Press.
- Sargent, Thomas J. (1993) *Bounded Rationality in Macroeconomics*. Clarendon Press, Oxford.
- Tesfatsion, Leigh (1995) A trade network game with endogenous partner selection.” ISU Economic Report No. 36, April.

- Tesfatsion, Leigh (1997a) A trade network game with endogenous partner selection. In Hans Amman, Berc Rustem, and Andrew Whinston, eds., *Computational Approaches to Economic Problems*, pp. 249–269. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Tesfatsion, Leigh (1997b) How economists can get alive. In W. Brian Arthur, Steven Durlauf, and David Lane (eds.), *The Economy as an Evolving Complex System, II*, pp. 533–564. Proceedings Volume XXVII, Santa Fe Institute Studies in the Sciences of Complexity, Reading, MA: Addison-Wesley.
- Tesfatsion, Leigh (1998) Preferential partner selection in evolutionary labor markets: A study in agent-based computational economics. In V. William Porto, N. Saravanan, Don Waagen, A. E. Eiben, eds., *Evolutionary Programming VII*, pp. 15–24. Proceedings of the Seventh Annual Conference on Evolutionary Programming, Berlin: Springer-Verlag.
- Tesfatsion, Leigh (1999) Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search. ISU Economic Report 51. to appear in the *Journal of Economic Dynamics and Control*.
- Yellon, Janet (1991) Efficiency-wage models of unemployment. In N. Gregory Mankiw and David Romer (eds.), *New Keynesian Economics*, Volume 2, The MIT Press, Cambridge, MA, pp. 113–122.

Copyright © 2000 Leigh Tesfatsion. All Rights Reserved.

```
class Agent
{
  Internalized Social Norms:
    Market protocols for communicating with other agents;
    Market protocols for job search and matching;
    Market protocols for worksite interactions;
  Internally Stored State Information:
    My attributes;
    My endowments;
    My beliefs and preferences;
    Addresses I have for myself and for other agents;
    Additional data I have about other agents.
  Internal Behavioral Rules:
    My rules for gathering and processing new information;
    My rules for determining my worksite behavior;
    My rules for updating my beliefs and preferences;
    My rules for measuring my utility (fitness) level;
    My rules for modifying my rules.
};
```

Table 1: General Form of the Internal Structure of an Agent.

```

int main () {

    InitiateEconomy();           // Construct initial subpopulations of
                                // work suppliers and employers with
                                // random worksite strategies.

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

                                // GENERATION CYCLE:

        InitiateGen();          // Configure work suppliers and employers
                                // with user-supplied parameter values
                                // (initial expected utility levels, work offer
                                // quotas, employer acceptance quotas,...)

        For (I = 1,...,IMax) {  // Enter the Trade Cycle Loop

                                // Trade Cycle:
                                // Work suppliers and employers determine
                                // their worksite partners, given
                                // their expected utility assessments,
                                // and record job search and
                                // inactivity costs.
                                // Work suppliers and employers engage
                                // in worksite interactions and
                                // record their worksite payoffs.
                                // Work suppliers and employers update their
                                // expected utility assessments, using
                                // newly recorded costs and worksite
                                // payoffs, and begin a new trade cycle.

            MatchTraders();
            Trade();
            UpdateExp();
        }

                                // Environment Step:
                                // Work suppliers and employers
                                // assess their utility levels.

        AssessFitness();

                                // Evolution Step:
                                // Worksite strategies of work suppliers and
                                // employers are separately evolved, and
                                // a new generation cycle begins.

        EvolveGen();
    }
    Return 0;
}

```

Table 2: Logical Flow of the Labor Market Framework

		Employer	
		c	d
Work Supplier	c	(C,C)	(L,H)
	d	(H,L)	(D,D)

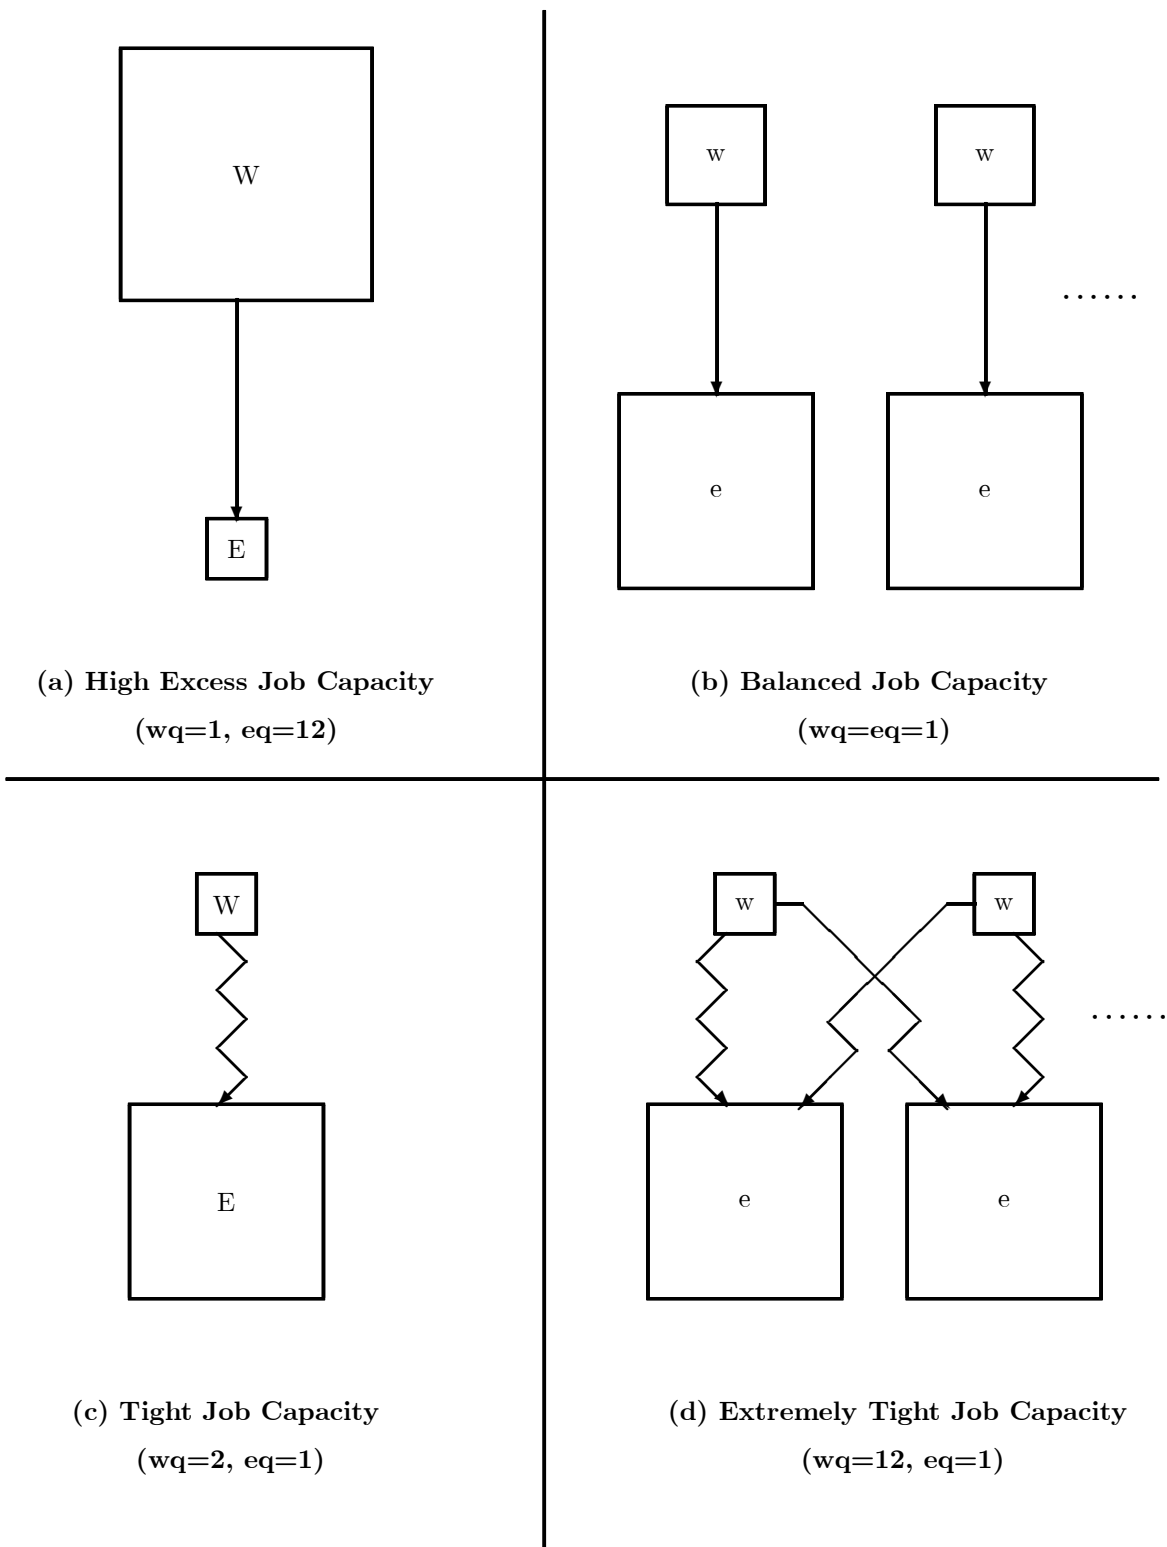
Table 3: Payoff Matrix for the Worksite Prisoner's Dilemma Game

```

// PARAMETER VALUES HELD FIXED ACROSS EXPERIMENTS
GMax = 50 // Total number of generations.
IMax = 150 // Number of trade cycles per trade cycle loop.
AgentCount = 24 // Total number of agents.
RefusalPayoff = -0.5 // Payoff R received by a refused agent.
InactivityPayoff = +0.0 // Payoff received by an inactive agent.
Sucker = -1.6 // Lowest possible worksite payoff, L.
BothDefect = -0.6 // Mutual defection worksite payoff, D.
BothCoop = +1.4 // Mutual cooperation worksite payoff, C.
Temptation = +3.4 // Highest possible worksite payoff, H.
InitExpPayoff = +1.4 // Initial expected utility level,  $U^0$ .
Elite = 67 // GA elite percentage for each agent type.
MutationRate = .005 // GA mutation rate (bit toggle probability).
FsmStates = 16 // Number of internal FSM states.
FsmMemory = 1 // FSM memory (in bits) for past move recall.
WorkSuppliers = 12 // Number of work suppliers.
Employers = 12 // Number of employers.
// PARAMETER VALUES VARIED ACROSS EXPERIMENTS
WorkerQuota = 1 // Work offer quota wq.
EmployerQuota = 12 // Employer acceptance quota eq.

```

Table 4: Parameter Values for a Labor Market with High Excess Job Capacity



**Figure 1: Base Network Patterns for Labor Markets with Differential Job Capacities.** A relatively larger box for either work suppliers (W) or employers (E) under a particular job capacity treatment indicates that this agent type attains a relatively higher average utility level in the sample economies whose networks approximate the depicted base network pattern. Straight edges indicate latched (continuous) persistent relationships and zig-zag edges indicate recurrent (random or periodic) persistent relationships.

$D^\circ$ Cluster	% Runs	AGGRESSIVE		P-INACTIVE		P-NICE		UTILITY	
		w	e	w	e	w	e	w	e
3-9	75%	97% (5%)	16% (34%)	2% (3%)	40% (12%)	3% (5%)	39% (28%)	1.74 (.27)	0.35 (.14)
24	25%	2% (3%)	5% (7%)	2% (3%)	5% (7%)	98% (3%)	95% (7%)	1.39 (.02)	1.02 (.03)

Table 5(a): High Excess Job Capacity ( $wq=1, eq=12$ )

$D^\circ$ Cluster	% Runs	AGGRESSIVE		P-INACTIVE		P-NICE		UTILITY	
		w	e	w	e	w	e	w	e
0-2	75%	16% (33%)	23% (39%)	1% (3%)	1% (3%)	94% (6%)	86% (26%)	1.10 (.14)	1.33 (.22)
4	10%	50% (50%)	54% (46%)	8% (8%)	8% (8%)	50% (50%)	46% (46%)	0.57 (.05)	0.86 (.57)
24	15%	0% (0%)	22% (20%)	0% (0%)	8% (0%)	89% (16%)	78% (20%)	0.24 (.08)	1.42 (.05)

Table 5(b): Balanced Job Capacity ( $wq=eq=1$ )

$D^\circ$ Cluster	% Runs	AGGRESSIVE		P-INACTIVE		P-NICE		UTILITY	
		w	e	w	e	w	e	w	e
0-7	55%	2% (3%)	5% (9%)	19% (10%)	4% (7%)	81% (10%)	96% (6%)	0.30 (.05)	1.35 (.09)
13-17	15%	100% (0%)	69% (43%)	47% (14%)	19% (18%)	8% (12%)	14% (20%)	0.32 (.04)	0.76 (.13)
24	30%	100% (0%)	100% (0%)	100% (0%)	100% (0%)	0% (0%)	0% (0%)	-0.10 (0)	-0.02 (0)

Table 5(c): Tight Job Capacity ( $wq=2, eq=1$ )

$D^\circ$ Cluster	% Runs	AGGRESSIVE		P-INACTIVE		P-NICE		UTILITY	
		w	e	w	e	w	e	w	e
0-6	35%	1% (3%)	1% (3%)	12% (4%)	1% (3%)	86% (7%)	96% (6%)	0.31 (.03)	1.37 (.06)
15-17	20%	10% (14%)	92% (14%)	35% (7%)	2% (4%)	17% (20%)	25% (34%)	0.35 (.17)	1.22 (.20)
24	45%	100% (0%)	100% (0%)	100% (0%)	100% (0%)	0% (0%)	0% (0%)	-0.10 (.00)	-0.01 (.00)

Table 5(d): Extremely Tight Job Capacity ( $wq=12, eq=1$ )

**Table 5. Experimental Findings for Differential Job Capacities**