Economic Report 48
Original Version: 25 October 1998
Revised: 15 November 1998

# Ex Ante Capacity Effects in Evolutionary Labor Markets with Adaptive Search[1]

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

**Abstract**: This study reports on computational experiments for an agent-based labor market model with adaptive choice and refusal of worksite partners and with endogenously evolving worksite behaviors. Two treatment factors are experimentally varied: market structure; and ex ante capacity constraints on potential work offers and job openings. Particular attention is focused on experimentally determined correlations between treatment factors and the formation of contractual networks among workers and employers, and between contractual network formation and the types of worksite interactions and welfare outcomes that these contractual networks support.

**Keywords**: Labor markets; search and matching; contractual networks; adaptation; evolutionary game; agent-based computational economics; C++ source code.

[1] The author is grateful to P. Orazem and R. Reed for helpful comments.

## I. Introduction

An interesting theoretical literature stressing job search and matching in labor markets has flourished following the influential work in this area by Jovanovic (1979), Diamond (1982), and Pissarides (1985) in the late nineteen seventies and early nineteen eighties. See, for example, Pissarides (1990), Hosios (1990), Mortenson and Pissarides (1994), Aghion and Howitt (1994), and MacLeod and Malcomson (1998). To achieve analytical tractability, however, researchers in this literature commonly postulate an aggregate matching function that proxies the complicated process of employer recruitment, worker search, and mutual evaluation. That is, the intense individualistic rivalry among workers and employers that characterizes many real-world labor markets is not modelled. Moreover, again for tractability, the competition of ideas within agents is generally not considered; attention is largely focused on steady-state behavior.[2]

In a preliminary study [Tesfatsion (1998b)], it was conjectured that some of the tractability problems encountered in analytical labor market studies might be alleviated by taking an agent-based computational approach. To explore this possibility, a labor market framework was designed that built on the computational trade network game (TNG) developed in Tesfatsion (1995, 1997a,b, 1998a) for studying the formation and evolution of buyer-seller trade networks under alternatively specified market structures.

The TNG labor market framework comprises a collection of workers and employers that evolve over time. Workers and employers repeatedly choose and refuse worksite partners on the basis of continually updated expected utility, engage in worksite interactions modelled as two-person games, and evolve their worksite strategies on the basis of past interaction experiences. This

[2]Important exceptions exist, of course. For example, Pissarides (1985) studies the dynamics of adjustment in labor markets in response to shocks. Montgomery (1991) explores the relationship between social networks and labor market outcomes when employers are able to use a stochastically determined social network among workers to relay job offers to the acquaintances of high-ability workers. In addition, motivated in part by an early seminal paper by Myerson (1977), an interesting analytical literature has developed in recent years focusing on the endogenous formation of agent linkages in social communication networks, which has potentially important implications for labor market modelling. See, for example, Kirman et al. (1986), van den Nouweland (1993), Jackson and Wolinsky (1996), and Watts and Strogatz (1998).

framework permits many aspects of labor markets to be determined endogenously on the basis of local agent interactions that analytical labor market studies commonly specify in more restricted ways, either as exogenously given attributes or as assumed aggregate steady-state relationships. Examples of these endogenized aspects include: worker preference orders over potential employers; employer preference orders over potential workers; intensity of job search activity; search opportunity costs; contractual networks among workers and employers; worker and employer worksite behaviors; worker compensations; employer earnings; quit rates; firing rates; vacancy rates; and unemployment rates.

This study reports on computational experiments for the TNG labor market framework.[3] As stressed by Diamond (1982) and Hosios (1990, pp. 285–288), entry/exit externalities arise in labor search and match models because, *ceteris paribus*, each additional worker makes it easier for employers to fill job vacancies but harder for workers to find job openings, and each additional employer makes it easier for workers to find job openings but harder for employers to fill job vacancies. The experimental design used in the current study examines the long-run effects of entry/exit externalities on worksite behavior, contractual network formation, and social welfare under alternative market structures as the nature of the entry/exit externalities is varied from being strongly biased in favor of workers to being strongly biased in favor of employers.

More precisely, the experimental design of this study involves two treatment factors: labor market structure; and ex ante capacity constraints on potential work offers and job openings. Three types of labor market structures are investigated: two-sided markets comprising pure workers and pure employers; partially fluid markets comprising pure workers, pure employers, and agents capable of functioning as both workers and employers; and endogenous-type markets in which each agent is capable of functioning as both a worker and an employer. For each labor market structure, various

[3] All labor market experiments reported in this study were implemented using version 105b of the TNG source code developed by McFadzean and Tesfatsion (1997), which in turn is supported by SimBioSys, a general C++ class framework for evolutionary simulations developed by McFadzean (1995). Version 105b of the TNG source code and SimBioSys are both available for downloading as freeware at the current author's Web site, along with extensive user instructions.

ex ante capacity constraints are considered ranging from high excess job capacity (many potential job openings relative to potential work offers) to extremely tight job capacity (few potential job openings relative to potential work offers).

For each specification of the two treatment factors, twenty runs ("economies") are generated for the TNG labor market using twenty different pseudo-random number seed values. In examining these sample economies, particular attention is focused on experimentally determined correlations between the treatment factors and the formation of contractual networks among workers and employers, and between contractual network formation and the types of strategic worksite interactions and welfare outcomes that these contractual networks support.

One interesting finding is that unemployment and vacancy rates are not as tightly correlated with ex ante capacity constraints as might be expected. Observed unemployment rates among workers do tend to increase as job capacity is incrementally tightened and observed vacancy rates among employers do tend to decline. On the other hand, workers and employers favored ex ante with greater market power tend to evolve more predacious worksite behaviors, which can lead to increased firing rates (employers refusing further work offers from predacious workers) and quit rates (exploited workers directing future work offers elsewhere).

Another interesting finding observed for each treatment is that multiple types of contractual networks can arise and persist, each supporting a distinct pattern of worksite behaviors and welfare outcomes. This finding is consistent with the findings of analytical two-sided labor market studies [e.g., Diamond (1982)] concerning 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 economies that evolve each type of network formation, which provides suggestive information regarding the size and importance of their basins of attraction.

Moreover, the number of contractual networks that arise and persist for each treatment tends to increase as one moves from endogenous-type markets to two-sided markets to partially-fluid markets. In particular, the presence in partially-fluid markets of a subset of flexible worker-employer

agents able to "crowd out" pure workers or pure employers tends to lead to increased volatility in network formation, worksite behaviors, and welfare outcomes.

Although multiple network formations are observed for each treatment, the distribution of these network formations and the nature of the worksite behaviors they support tend to be strongly correlated with the differential market power of workers and employers as implied by the specification of the treatment factors. For example, in conditions of tight job capacity in which pure workers have low match rates, the predominant types of networks that form tend to support aggressive and predacious employer behavior against pure workers. Conversely, in conditions of excess job capacity in which pure employers have low match rates, the predominant types of networks that form tend to support aggressive and predacious worker behavior against pure employers. These experimental findings provide support for the analytical studies of Pissarides (1985), Rubinstein and Wolinsky (1985), and Wolinsky (1987), among others, who explore search and match models in which the relative returns to the participants of a match are assumed to be affected by their relative match probabilities.

The exploitation opportunities that arise due to differential market power induce structural changes in agent worksite behaviors as the agents evolve their worksite strategies over time in response to past worksite experiences. In consequence, workers and employers who succeed in establishing persistent relationships (matches) with each other tend to express relationship-specific worksite behaviors; namely, an initial transient "handshake" involving some degree of non-cooperative behavior on the part of both match participants as attempted exploitation results in retaliation, followed by a persistent repeated action pattern in which the agent with the greater market power limits his exploitation to a level that the other participant is willing to tolerate. Interestingly, structurally distinct workers (both pure workers and worker-employers) who interact repeatedly with the same employer generally come to express the same handshake and persistent action patterns with that employer, so that the employer acts as a coordinating device across workers; and similarly for groups of employers (both pure employers and worker-employers) who interact with the same

worker.

Despite this variation at the level of individual agents, however, the average welfare levels achieved by each agent type (pure workers, pure employers, and worker-employers) across the sample of economies generated for each treatment tend to strongly reflect the differential market power of workers and employers as implied by the specification of the treatment factors. These average welfare levels typically fixate at approximately constant levels within the first twenty-five agent generations and persist at these levels through 100 agent generations, the maximum number of generations tested.

Another interesting related finding is that the efficiency and stability criteria commonly used to evaluate the performance of matching mechanisms in steady-state (or static) labor search and match models turn out to be highly incomplete indicators of welfare outcomes in the current TNG labor market context. The steady-state viewpoint hides the strong role played by market structure and ex ante capacity constraints in determining the types of persistent matching networks that evolve, the types of persistent interaction behaviors that these networks support, and the transactions costs and opportunity costs to agents that the achievement of these persistent networks and behaviors entails.

More precisely, as will be clarified in later sections, the matching outcomes generated for workers and employers in the TNG labor market are generated via a Gale-Shapley matching mechanism that guarantees on a period by period basis both Pareto efficiency for workers (the agents who make offers) and pairwise stability. These local optimality properties take agent preference rankings over potential partners as given, however, whereas in actuality there is a continuous feedback between preference rankings and agent interactions over time. In consequence, these local optimality properties have very little global predictive power; they tend to be overwhelmed by the strong channeling effects of market structure and ex ante capacity contraints. This suggests the need, in dynamic contexts, for more global optimality criteria for matching mechanisms that take into account the joint determination of matching networks and interaction behaviors over time.

The TNG labor market framework is described in Section II. Section III constructs ex post measures for the classification of network formations, worksite behaviors, and welfare outcomes. The experimental design of the study is outlined in Section IV, and a detailed discussion of experimental findings is presented in Section V. Concluding remarks are given in Section VI.

## II. The TNG Labor Market Framework

The TNG labor market consists of three disjoint (and possibly null) subpopulations of agents that separately evolve over time: pure workers who make work offers; pure employers who receive work offers; and worker-employers capable of both making and receiving work offers. The pure workers and worker-employers are collectively referred to as *workers*, and the worker-employers and pure employers are collectively referred to as *employers*. Each worker can have no more than $wq$ work offers outstanding to employers at any given time, and each employer can accept no more than $eq$ work offers from workers at any given time, where the work offer quota $wq$ and the acceptance quota $eq$ can be any positive integers.[4] Although highly simplified, these parametric specifications will be seen in Section IV, below, to permit the study of a variety of labor market structures operating under different ex ante capacity constraints.

As outlined in Table 1, each agent in the initial generation is constructed and assigned a random strategy governing worksite interactions. The agents then enter into a nested pair of generation cycle and trade cycle loops during which they repeatedly determine contractual partnerships, engage in worksite interactions, update their expected utility assessments for worksite partners based on newly recorded payoffs, and evolve their worksite strategies over time.

The TNG labor market framework facilitates the study of labor markets from an agent-based perspective in two key ways. First, as depicted in Table 2, each agent is instantiated as an autonomous, endogenously-interacting software agent with internally stored state information and

[4] When $wq$ exceeds 1, the workers can be interpreted as some type of information service provider (broker, consultant, ...) able to provide services to more than one employer at a time.

```
int main () {
  Init();                    // Construct initial subpopulations of pure workers,
                             //   worker-employers, and pure employers with
                             //   random worksite strategies
  For (G = 1,...,GMax) {     // Enter the generation cycle loop.
                             // Generation Cycle:
    InitGen();               //   Configure all agents with user-supplied
                             //     parameter values (initial expected utility
                             //     levels, work offer/acceptance quotas,...).
    For (I = 1,...,IMax) {   //   Enter the trade cycle loop.
                             //   Trade Cycle:
      MatchTraders();        //     Determine worksite partners,
                             //       given expected utilities,
                             //       and record refusal and
                             //       wallflower payoffs.
      Trade();               //     Engage in worksite interactions
                             //       and record worksite payoffs.
      UpdateExp();           //     Update expected utilities
    }                        //       using newly recorded payoffs.
                             //   Environmental Step:
    AssessFitness();         //     Assess agent fitness scores.
    Output();                //     Output agent information.
                             //   Evolution Step:
    EvolveGen();             //     Separately evolve the worksite strategies of pure
                             //       workers, worker-employers, and pure employers.
  }
  Return 0;
}
```

Table 1: Logical Flow of the TNG Labor Market Framework

with internal behavioral rules. The agents can therefore engage in anticipatory behavior. Moreover, using stored agent addresses together with internalized communication protocols, they can communicate with each other at event-triggered times, a feature not present in standard economic models.

Second, as seen in Table 1, the TNG labor market framework is modular in design. This means that experimentation with alternative specifications for market structure, search and matching among workers and employers, worksite interactions, expectation formation and updating, and evolution of worksite strategies can easily be undertaken — much like changing a lightbulb in a multi-bulb lamp — as long as the interfaces (inputs and outputs) for the modules implementing these specifications remain unchanged. Moreover, each of these modules can potentially be grounded in agent-initiated actions in the sense that the module is implemented via behavioral rules internal to the agents. Finally, the transitory and longer-run implications of each alternative

```
class Agent
{
    Internal State Information:
        My physiological attributes;
        My endowments;
        My beliefs/preferences;
        Addresses for other agents;
        Additional data about other agents.
    Internal Behavioral Rules:
        Rules for communicating with other agents;
        Rules for gathering, processing, and updating information;
        Rules for determining my contractual partners;
        Rules for conducting my worksite interactions;
        Rules for updating my beliefs/preferences;
        Rules for calculating my welfare;
        Rules for altering my rules.
};
```

Table 2: Schematic Description of an Agent

module specification can be studied at three different levels: individual characteristics of workers and employers; interactions among workers and employers (network formation); and social welfare as measured by descriptive statistics such as average agent welfare and unemployment rates.

A brief description will now be given for the particular module specifications used in all experiments reported below. See McFadzean and Tesfatsion (1997) for a more careful description.

If an employer accepts a work offer from a worker in any given trade cycle, the worker and employer are said to be *matched* for that trade cycle. Each match constitutes a mutually agreed upon contract stating that the worker shall be employed 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 worker and employer engage in a worksite interaction modelled as a two-person prisoner's dilemma game reflecting the basic efficiency wage hypothesis that worker effort levels are affected by overall working conditions (e.g., wage levels, respectful treatment, safety considerations, ...). The worker 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 payoffs is assumed to be the same for each match in each trade cycle: namely, as seen in Table 3, a cooperator whose

| | | Employer | |
|---|---|---|---|
| | | c | d |
| Worker | c | (C,C) | (L,H) |
| | d | (H,L) | (D,D) |

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

contractual partner defects receives the lowest possible payoff $L$ (the sucker payoff); a defector whose contractual partner also defects receives a payoff $D$; a cooperator whose contractual partner also cooperates receives a payoff $C$; and a defector whose contractual partner cooperates receives the highest possible payoff $H$ (the temptation payoff). The payoffs are assumed to be measured in utility terms and to be normalized about 0, so that $L < D < 0 < C < H$. They are also assumed to satisfy the usual regularity condition $(L + H)/2 < C$ guaranteeing that mutual cooperation dominates alternating cooperation and defection on average.

Matches between workers and employers are determined using a modified version of the well-studied "deferred acceptance mechanism" originally designed by Gale and Shapley (1962).[5] Under this modified mechanism, hereafter referred to as the *deferred choice and refusal* (DCR) mechanism, each worker submits up to $wq$ work offers to employers he ranks as most preferable on the basis of expected payoff and who he judges to be tolerable in the sense that their expected payoff 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 payoff and he places them on a waiting list; all other work offers are refused. Workers 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.

[5]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.

A worker incurs a transactions 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.[6] An agent who neither submits nor accepts work offers during a trade cycle receives a single *wallflower payoff 0* for the entire trade cycle. The refusal and wallflower payoffs are each assumed to be measured in utility terms.

Agents use a simple learning algorithm to update their expected utilities on the basis of new payoff information. Each agent $v$ assigns an exogenously given initial expected utility $U^o$ to each potential contractual 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 $z$ and dividing this sum by one plus the number of interactions with $z$.

The worksite behavior of each agent is governed by 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 actions and utility payoffs 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 and/or his initial expected utility assessment of the agent. Each agent thus keeps separate track of the particular state he is in with regard to each of his potential worksite partners.[7]

At the end of each trade cycle loop, the worksite strategies of pure workers, pure employers, and worker-employers are separately evolved by means of standardly specified genetic algorithms

[6] This is equivalent to assuming: (a) each worker incurs a transactions cost for each work offer he makes; and (b) the worksite payoffs in Table 3 are each increased by the amount of this transactions cost, so that a worker who succeeds in having a work offer accepted is able to recoup the transactions cost he incurred by making this offer.

[7] The implications of this type of informational assumption in search and match contexts has been explored by Rubinstein and Wolinsky (1990). They show that, even in the absence of trading frictions, noncompetitive equilibria can arise when agent behavior depends on specific information such as the identity of trading partners and personal histories with trading partners.

involving recombination, mutation, and elitism operations.[8] This evolution is meant to reflect the formation and transmission of new ideas rather than biological reproduction. Specifically, if a worksite strategy successfully results in high fitness for an agent of a particular type, where fitness is measured by average payoff, 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. Given the extent of information currently allowed to agents during the evolution step — i.e., knowledge of the complete strategies of all other agents of the same type, whether expressed in interactions or not — the evolution step is more appropriately interpreted as an iterative stochastic search algorithm for determining potential equilibrium strategy configurations rather than as a cultural transmission mechanism per se. The resulting welfare outcomes will be used in subsequent work as benchmarks against which to assess the effectiveness of more realistically modelled cultural transmission mechanisms.

## III. Descriptive Statistics

In this section care is taken to explain the measures that have been constructed to aid in the experimental determination of correlations between treatment factors and contractual network formation, and between contractual network formation and the types of worksite behaviors and social welfare outcomes that these contractual networks support. Contractual 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, social welfare

[8] More precisely, for each agent type (pure workers, pure employers, or worker-employers), 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) *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) worksite strategies in the existing collection of strategies are retained for the new collection of strategies and the least fit worksite strategies in the existing strategy collection are replaced with offspring strategies. See McFadzean and Tesfatsion (1997) for a more detailed discussion of this use of genetic algorithms in the TNG, and see Goldberg (1989) and Mitchell and Forrest (1994) for a general discussion of genetic algorithm design and use.

refers to the overall utility levels achieved by each agent type (pure workers, pure employers, and worker-employers) from repeated worksite interactions in a possibly changing network of contractual partners.

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

As noted in the introduction, all of the TNG labor market experiments reported in this study are implemented using version 105b of the TNG source code developed by McFadzean and Tesfatsion (1997). Let $s$ denote a seed value for the pseudo-random number generator incorporated in this code, and let $e$ denote a *potential TNG economy*, i.e., an economy characterized structurally by the TNG source code together with all of the user-specified TNG parameter values apart from $s$. The *realized TNG economy* generated from $e$, given the seed value $s$, is denoted by $(s, e)$.

Worksite strategies are represented as finite state machines,[9] 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 *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.

Two agents $v$ and $z$ are said to exhibit a *persistent relationhip* during a given trade cycle loop $T$ of a realized TNG economy (s,e) if the following two conditions hold: (a) their worksite histories with each other during the course of $T$ take the form $H_v$:$P_v$ and $H_z$:$P_z$ with nonnull $P_v$ and $P_z$; and (b) accepted work offers between $v$ and $z$ do not permanently cease during $T$ either by choice (a

[9] A *finite state machine (FSM)* 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 together with current system inputs. For the application at hand, the inputs are the actions selected by a worker and an employer engaged in a worksite interaction. See McFadzean and Tesfatsion (1997) for a more detailed discussion and illustration of the FSM representation of agent strategies used in the TNG source code.

permanent switch away to strictly preferred contractual partners) or by refusal (one agent becoming intolerable to the other because his expected utility drops below zero).

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

Let $V^o(e)$ denote a *base contractual pattern* that partially or fully specifies a potential pattern of contractual relationships among the agents $V(e)$ in the potential TNG economy $e$. For example, $V^o(e)$ could designate that each worker directs work offers to at least two employers. Let $K^o(e)$ denote the *base contractual network class* consisting of all contractual networks $K(e)$ whose edges conform to the base contractual pattern $V^o(e)$. Also, let $K(s,e)$ denote the contractual network depicting the actual pattern of contractual relationships among the agents $V(e)$ in the final generation of the realized TNG economy $(s,e)$. The reduced form contractual network $K^p(s,e)$ derived from $K(s,e)$ by eliminating all edges of $K(s,e)$ that correspond to non-persistent relationships is referred to as the *persistent contractual network* for $(s,e)$.

The *distance* $D^o(s,e)$ between the persistent contractual network $K^p(s,e)$ and the base contractual network class $K^o(e)$ for a realized TNG economy $(s,e)$ is then defined to be the number of vertices (agents) in $K^p(s,e)$ whose arrow patterns (persistent relationships) fail to conform to the base contractual pattern $V^o(e)$. This distance measure provides a rough way to classify the different types of persistent contractual networks observed to arise for a given value of $e$ as the seed value $s$ is varied.

### B. *Classification of Worksite Behaviors and Welfare Outcomes*

An agent in the final generation of a realized TNG economy $(s,e)$ is referred to as an

*unprovoked defector (UD)* if he engages in at least one defection against another agent who has not previously defected against him. The vector giving the separate UD percentages for pure workers, pure employers, and worker-employers in the final generation of agents for $(s, e)$ is referred to as the *UD profile* for $(s, e)$. The UD profile measures the extent to which the different types of agents behave agressively in worksite interactions with contractual partners who are either strangers or who so far have been consistently cooperative.[10]

Also, an agent in the final generation of a realized TNG economy $(s, e)$ is referred to as a *persistent wallflower (PWF)* if he constitutes an isolated vertex of the persistent contractual network $K^p(s, e)$. Alternatively, the agent is referred to as a *repeat defector (RD)* if he establishes at least one persistent relationship for which the persistent portion $P$ of his worksite history $H{:}P$ includes a defection $d$. If, instead, the agent 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 a *persistent cooperator (PC)*. By construction, each agent in the final generation of $(s, e)$ belongs to one and only one of these categories; i.e., he must be either a PWF, or an RD, or a PC.

The vectors giving the separate PWF, RD, and PC percentages for pure workers, pure employers, and worker-employers in the final generation of $(s, e)$ are referred to as the *PWF profile*, the *RD profile*, and the *PC profile* for $(s, e)$, respectively. The PWF profile measures the extent to which the different agent types in this final generation fail to establish any persistent relationships. Since the DCR matching mechanism ensures that each worker in the TNG labor market attempts to make at least one work offer in the initial cycle of each trade cycle loop in which he participates, the PWF percentage for pure workers gives their persistent unemployment rate. Also, the PWF percentage for pure employers gives their persistent vacancy rate. On the other hand, the PWF percentage for worker-employers measures the extent to which they are both persistently unemployed (as workers) and persistently vacant (as employers). The RD profile measures the extent to which the

[10] The importance of stance towards strangers in determining subsequent outcomes in path dependent contexts such as the TNG labor market framework has been stressed by Orbell and Dawes (1993).

different agent types in the final generation of $(s, e)$ establish persistent relationships characterized by predacious behavior (recurrent or continuous defection). Finally, the PC profile measures the extent to which the different agent types in this final generation establish persistent relationships characterized by fully cooperative behavior.

The vector that separately gives the average fitness score for pure workers, pure employers, and worker-employers, respectively, in the final generation of agents for a realized TNG economy $(s, e)$ is referred to as the *FIT profile* for $(s, e)$. The FIT profile constitutes a measure of social welfare.

## IV. Experimental Design

The labor market experiments reported in the following section focus on three simple labor market structures: endogenous-type markets comprising 24 worker-employers; two-sided markets comprising 12 pure workers and 12 pure employers; and partially-fluid markets comprising 8 pure workers, 8 pure employers, and 8 worker-employers. For each market structure, four different configurations for the worker offer quota $wq$ and employer acceptance quota $eq$ are examined: high excess job capacity ($eq >> wq$); zero excess job capacity ($eq = wq = 1$); tight job capacity ($eq = 1$ and $wq = 2$); and extremely tight job capacity ($eq << wq$). Thus, for each market structure, the focus is on ex ante capacity effects as measured by the ratio of potential work offers to potential job openings, holding fixed the market concentration levels as measured by relative numbers of workers and employers.

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 agent type and quota values for a two-sided market experiment with high excess job capacity. The parameter values in Table 4, together with the TNG source code, constitute a potential TNG economy $e$ in the sense defined in Section III.A.

For each tested $e$, twenty TNG economies $(s, e)$ were experimentally generated using twenty arbitrarily selected seed values $s$ for the TNG pseudo-random number generator.[11] The persistent

[11] 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,

```
// 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.
    WallflowerPayoff = +0.0     // Payoff W 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^o.
    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.
// PARAMETER VALUES VARIED ACROSS EXPERIMENTS
    PureWorkers = 12            // Number of pure workers.
    PureEmployers = 12          // Number of pure employers.
    WorkerEmployers = 0         // Number of worker-employers.
    WorkerQuota = 1             // Worker offer quota wq.
    EmployerQuota = 12          // Employer acceptance quota eq.
```

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

contractual network $K^p(s,e)$ for each run $s$ was determined and graphically depicted, and the mean and standard deviation for the UD (unprovoked defector), PWF (persistent wallflower), RD (repeat defector), PC (persistent cooperator), and FIT (fitness) profiles were determined and recorded.[12]

A base contractual pattern $V^o(e)$ was then specified for each tested $e$. Although the choice of this base pattern is simply a normalization determining a 0 point for the distance measure $D^o$, and hence intrinsically arbitrary, the degree of specificity of this base pattern governs the dispersion of the resulting distance values and the extent to which these distance values display useful correlations with worksite behaviors as measured by the UD, PWF, RD, PC, and FIT profiles. In practice, then, the choice of the base contractual pattern was fine-tuned so that the resulting distance values provided a meaningful informative classification of network types. 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

425, and 212. The final fourteen values were determined by random throws of two and three die. The TNG source code implementation of the TNG 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.

[12]By construction, as explained in Section III.B, the PWF, RD, and PC profiles must sum to the total percentages of agent types. Consequently, only the PWF and PC profiles are reported below in the table presentations of experimental findings.

$D^o(s, e)$ 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 potential economy $e$ and the minimum, maximum, and average fitness scores for the agents in each of the 100 generations were graphically generated for each realised economy $(s, e)$.

## V. Experimental Findings

### A. *Overview*

A general discussion of the principal experimental findings is presented in Section I. In this section, additional clarification of this general discussion is first provided by making use of the more precise terminology and measures introduced in the intervening Sections II through IV. A more detailed discussion of experimental findings is then given on a case by case basis.

In two-sided and partially-fluid economies $e$, unemployment rates as measured by persistent wallflower (PWF) percentages for pure workers tend to increase as job capacity is incrementally tightened whereas vacancy rates as measured by PWF percentages for pure employers tend to decline. Nevertheless, countervailing forces at least partially offset these basic market power effects.

Specifically, agents who are favored with greater ex ante market power tend to express more aggressive and predacious worksite behavior, in the sense that they engage in more UD (unprovoked defection) and RD (repeat defection) in their worksite interactions. This can result in unemployment due to firings, i.e., unemployment due to employer decisions to refuse all further work offers from aggressive and/or predacious workers; and it can also result in vacancies due to quits as exploited workers direct their future work offers elsewhere. Also, for economies $e$ characterized ex ante by extremely tight job capacity, the predominant outcome is complete coordination failure (unemployment rates and vacancy rates of 100 percent) despite the ex ante market power advantage held by pure employers. The difficulty is that the high transactions costs (negative refusal payoffs)

accumulated by workers in their attempts to find job openings tend to prevent the formation of any persistent relationships between workers and employers.

A second interesting finding observed for many (if not all) of the twenty realized economies $(s, e)$ generated for each tested economy $e$ is the remarkable stability exhibited by the average fitness scores for different agent types from one generation to the next, with stability often setting in as early as generation 10. This observed stability in average fitness scores occurs despite repeated changes induced in the underlying worksite strategies of agents by repeated applications of genetic algorithm operations. Runs $(s, e)$ for which instabilities in average fitness scores were detected are noted in the discussion of specific experimental findings, below.

A third interesting finding observed for the twenty realized economies $(s, e)$ generated for each tested economy $e$ is the existence of multiple distinct forms of persistent contractual networks $K^p(s, e)$. More precisely, the distance values $D^o(s, e)$ for these persistent contractual networks tend to cluster around a small but multiple number of isolated distance values. Endogenous-type economies tend to have from one to two distinct clusters for each capacity treatment; two-sided markets tend to have from two to three distinct clusters for each capacity treatment; and partially-fluid markets tend to have three distinct clusters for each capacity treatment. In general, each of these distinct clusters supports a distinct pattern of worksite behaviors and welfare outcomes for the different agent types. That is, the mean distance value for each distance cluster tends to be strongly correlated with the mean UD, PWF, RD, PC, and FIT profiles calculated for the cluster. Consequently, for the TNG labor market framework at hand, there does not appear to be a unimodal central-tendency network distribution in the sense defined by Banks and Carley (1994).

A fourth interesting finding is that the welfare outcomes supported by the various forms of persistent contractual networks $K^p(s, e)$ that arise for each tested two-sided and partially-fluid economy $e$ tend to strongly reflect the differential market power of workers and employers as implied by the specification of the treatment factors for $e$. More precisely, in conditions of either high excess job capacity (which favors workers) or zero excess job capacity or tight job capacity (which favor

employers), the most predominant network forms that arise and persist tend to support welfare outcomes under which the advantaged agent type attains not only a markedly higher welfare level than other agent types but also a markedly higher welfare level than this advantaged agent type attains under other network forms observed to arise for this $e$. This correlation between market power and welfare outcomes breaks down, however, when job capacity is extremely tight. In the latter case the most predominant persistent network form that arises and persists is a completely disconnected network consisting of workers and employers who have all devolved to wallflowers and who all attain approximately zero welfare levels.

Finally, a fifth interesting finding is that the efficiency and stability criteria conventionally used to evaluate the performance of matching mechanisms in steady-state or static market contexts turn out to be highly incomplete indicators of performance for the TNG labor market.

As detailed in Section II, the DCR mechanism used to match workers and employers in all experiments reported in this study is a version of the well-known matching mechanism developed by Gale and Shapley (1962). The matching outcomes generated via the DCR mechanism in each trade cycle have been shown in Tesfatsion (1995, 1997a) to have the usual local optimality properties associated with Gale-Shapley type matching mechanisms: namely, pairwise stability; and Pareto optimality from the vantage point of workers, the agents who actively make offers. Nevertheless, the long-run outcomes observed in these evolutionary labor market experiments include autarkic economies in which all agents are persistent wallflowers, exploitive economies in which employers persistently defect against cooperate workers or workers persistently defect against cooperative employers, and fully harmonious economies in which all agents are persistent cooperators. Moreover, due to transactions costs (negative $R$ payoffs) and opportunity costs (0 wallflower payoffs), social welfare can still be low even if all active agents are persistent cooperators. These evolutionary outcomes are systematically related to market structure and to ex ante capacity constraints as represented by the worker offer quota $wq$ and the employer acceptance quota $eq$.

B. *Findings for Endogenous-Type Labor Markets*

Consider an endogenous-type labor market economy $e$ comprising 24 worker-employers with a worker offer quota $wq = 1$ and an employer acceptance quota $eq = 24$. These quota values indicate that $e$ has a high excess job capacity in the sense that the total number of work offers that employers can accept in each trade cycle far exceeds the maximum number of work offers that workers can make. As depicted in Figure 1(a), the base contractual pattern $V^o(e)$ for this economy $e$ is as follows: Each worker-employer directs work offers to other worker-employers without latching.[13]

**— Insert Figure 1 About Here —**

For this $e$, as detailed in Table 5(a),[14] 90% of the runs $(s, e)$ were observed to lie in the distance cluster 0–3. More precisely, in 18 of the 20 runs for this $e$, at most 3 of the 24 worker-employers in the final agent generation deviated from the base contractual pattern $V^o(e)$. Moreover, the mean UD profile for distance cluster 0–3 was 3%, meaning that on average only 3% of the agents in the final generation of each run in this distance cluster engaged in aggressive worksite behavior. The mean PWF profile for this distance cluster was 1%, i.e., on average, only 1% of the agents in the final generation of each run in this distance cluster were both unemployed (as workers) and vacant (as employers). The mean PC profile for this distance cluster was 96%, i.e., on average, 96% of the agents in the final generation of each run in this distance cluster ended up engaging in persistently cooperative behavior. Finally, the mean FIT profile for this distance cluster was 1.36, meaning that the agents in the final generation of each run in this distance cluster ended up with an average utility level (per interaction) that was very close to the mutual cooperation payoff level of 1.40.

[13] A worker is said to be *latched* to an employer $z$ if he works for $z$ continuously (in each successive trade cycle) rather than intermittently (randomly or recurrently). In the directed graph representations for base contractual patterns in Figures 1, 3, and 4, below, latched persistent relationships are depicted as straight edges and intermittent persistent relationships are depicted as zig-zag edges.

[14] In Tables 5, 6, and 7, below, the standard deviations for the UD (unprovoked defector), PWF (persistent wallflower), and PC (persistent cooperator) profiles are measured in percentages; they appear in parentheses beneath the mean values for these profiles and are rounded off to the nearest integer value. Also, the standard deviations for the FIT (fitness) profiles appear in parentheses below the mean FIT profiles and are rounded off to two decimal places. The calculation of these standard deviations is not applicable (NA) for distance clusters encompassing only one run, i.e., for distance clusters encompassing only 5% of the total sample of 20 realized economies.

**— Insert Table 5 About Here —**

A rough stability check was conducted for each of the 18 realized economies $(s, e)$ in distance cluster 0–3 for this high excess job capacity economy $e$ to check whether the information recorded in Table 5(a) for the final (fiftieth) generations appeared to be informative for other generations as well. Specifically, holding all other parameter values fixed, the number of generations was increased to 100 and the minimum, average, and maximum fitness scores attained by the agents in each of these 100 generations were recorded and graphically depicted. Figure 2 depicts the stability results obtained for the realized economy $(413, e)$ with distance value 0; these results are typical of the stability results obtained for all economies in distance cluster 0–3. The average fitness scores are seen to fluctuate closely around the mutual cooperation payoff level, 1.40, over generations 10 through 100.

**— Insert Figure 2 About Here —**

As seen in Table 5(a), two outlier runs also occurred for this high excess job capacity $e$ at distance values 11 and 23. The outlier run at distance 11 is characterized by a high degree of latching and a high degree of RD behavior. The outlier run at distance 23 is even more interesting. A wallflower crash occurs in generation 18, in the sense that almost all agents become persistent wallflowers. Specifically, UD and RD behavior are so prevalent among generation 18 agents that most quickly become intolerable to all other agents as worksite partners; indeed, only three worksite interactions take place in each of the three final trade cycles. By generation 50 this outlier run is still in an unsettled state. As seen in Table 5(a), 96% of the agents engage in UD behavior, although 88% ultimately end up in latched PC relationships. The stability check for this realized economy indicates, however, that the economy fully recovers from the wallflower crash by generation 64; UD behavior is rare among generation 64 agents and most succeed in establishing PC relationships. Moreover, this recovery is sustained through generation 100.

All of these observations would appear to have a simple structural explanation. In endogenous-

type economies, all agents evolve together in the evolution step. Hence, any worksite strategies garnering below-average fitness scores are soon eliminated and replaced with variants of more successful strategies. A strong evolutionary inducement thus exists towards uniform expressed worksite behavior and, in particular, towards mutual cooperation, which is the uniform expressed worksite behavior that generates the highest agent fitness scores. The only issue, then, is the extent to which the ex ante capacity constraints result in transactions costs (negative refusal payoffs) for workers trying to find tolerable job openings. In the case of high excess job capacity, workers face zero structural risk of refusal from employers due to ex ante capacity constraints; employers only refuse workers if they engage in an intolerable number of worksite defections.

This situation changes dramatically, however, as job capacity is incrementally tightened. As depicted in Figure 1 and detailed in Table 5, the base contractual pattern changes from random dispersion of work offers to disjoint doubly-latched pairings of workers and employers, and average agent fitness scores monotonically decrease. It is not aggression or predation in the form of UD or RD behavior that results in lower fitness scores for the agents in these tighter job capacity cases but rather the ever larger accumulation of transaction costs (negative refusal payoffs) that the agents incur in their attempts to find tolerable worksite partners.

### C. *Findings for Two-Sided Labor Markets*

Next consider the case of a two-sided labor market economy $e$ comprising 12 pure workers and 12 pure employers with a worker offer quota $wq = 1$ and an employer acceptance quota $eq = 12$, implying that excess job capacity is high. Indeed, the structural risk to workers of having their work offers refused by employers on the basis of limited ex ante acceptance capacity is zero. In contrast, employers are forced to remain vacant unless workers happen to direct work offers their way, implying that the employers face a substantial structural risk of incurring wallflower payoffs. The economy $e$ thus represents a "workers' market." As depicted in Figure 3(a), the base contractual pattern $V^o(e)$ for this economy $e$ is as follows: Each worker is latched to at least one

employer, and no employer is vacant.

**— Insert Figure 3 About Here —**

As seen in Table 6(a), 75% of the runs $(s, e)$ for this high excess job capacity $e$ were observed to lie in the distance cluster 3–9. In this distance cluster, the very low mean FIT value of 0.35 for employers is due to two factors: a high vacancy rate (high mean PWF percentage) due to high excess job capacity; and aggressive and persistently predacious behavior (high mean UD and low mean PC percentages) by workers that induces retaliatory RD behavior in some employers.

The persistent contractual networks that arise for the runs $(s, e)$ in distance cluster 3–9 reveal the following typical scenario: RD workers latch on to a selected subset of employers and drive down their fitness scores to small positive values, leaving other employers vacant with approximately zero fitness scores. (Indeed, the distance $D^o(s, e)$ for each run $(s, e)$ in distance cluster 3–9 is essentially a count of the number of persistently vacant employers, i.e., the number of employers who are PWFs.) This scenario ensures that the worksite strategies of the exploited employers are advantaged in the evolution step relative to the worksite strategies of the employers who are vacant. Since (pure) workers and (pure) employers evolve separately, the worksite strategies of the exploited employers tend to reproduce into the next generation. In this way the workers breed and maintain a subset of largely cooperative employers that they repeatedly exploit to their benefit.

**— Insert Table 6 About Here —**

Table 6(a) also shows that the remaining 25% of the runs for this $e$ lie in a second distance cluster 23–24. The mean FIT value of 1.02 achieved by employers in this second distance cluster is higher than that achieved in distance cluster 3–9 due to the higher mean percentage of PC behavior exhibited by both workers and employers. This mean FIT value is nevertheless substantially below the mutual cooperation payoff level, 1.40, due to the 5% vacancy rate among employers, a structual consequence of high excess job capacity that is independent of how cooperatively the employers behave in their worksite interactions. The typical contractual pattern exhibited in this distance

cluster is PC workers randomly directing work offers among employers without latching. Note that the mean FIT value 1.39 achieved by workers is very close to the mutual cooperation payoff level.

When excess job capacity is reduced to zero, the typical contractual network dramatically changes. As depicted in Figure 3(b) and detailed in Table 6(b), about 80% of the workers now form persistent relationships with employers in the form of disjoint latched pairings. The reason for the latching is that workers who fail to latch tend to accumulate large transactions costs (large numbers of negative refusal payoffs) and so become relatively disadvanted in the evolution step relative to those who latch. Nevertheless, even workers who succeed in latching on to one employer typically accumulate 2 or 3 refusal payoffs from a wide range of employers on the way to attaining this coordinated state, and these transactions costs tend to lower the mean FIT value of workers relative to employers.

The stability checks conducted for this zero excess job capacity case reveal that many of the realized economies exhibit unsettled average fitness score behavior over generations 1 through 100 in the form of persistent drifting, bubbling, or regime shifts. The reason for this appears to be that contractual networks are particularly vulnerable to initially cooperative mutant invaders when excess job capacity is zero since the networks form in response to transactions costs and yet support largely PC or even $c$:$c$ worksite behavior.

As job capacity keeps tightening, workers have an increasingly difficult time forming any persistent relationships with employers, a finding indicated in Figure 3 by the decreasing size of worker boxes relative to employer boxes as one moves from part (a) to part (d). This increased coordination failure is detailed in Table 6. Note, in particular, the growing mean percentage of workers who become unemployed (PWFs) as job capacity successively tightens.

### D. *Findings for Partially-Fluid Labor Markets*

Finally, consider a partially-fluid labor market economy $e$ comprising 8 pure workers (pw), 8 pure employers (pe), and 8 worker-employers (we) with a worker offer quota $wq = 1$ and an

employer acceptance quota $eq = 16$, implying that excess job capacity is high. As depicted in Figure 4(a), the base contractual pattern $V^o(e)$ for this economy $e$ is as follows: Each worker directs work offers to employers without latching, and no pure employer is vacant.

**— Insert Figure 4 About Here —**

As seem in Table 7(a), the runs $(s, e)$ for this $e$ are divided about equally into three distance clusters. In the first distance cluster, although all agents exhibit a high degree of PC behavior, and few become persistent wallflowers, pure employers nevertheless tend to accumulate large numbers of wallflower payoffs. Consequently, pure employers have a mean FIT value, 1.03, that is low relative to the mean FIT value of 1.38 for pure workers and 1.38 for worker-employers. In the remaining two distance clusters, there is a substantial increase in latching behavior. Also, agents (particularly workers) engage in increased UD and RD behavior; and there is an increase in unemployment among pure workers and vacancy among pure employers that results in relatively low mean FIT values for these two agent types.

**— Insert Table 7 About Here —**

The stability checks for the 20 runs for this high excess job capacity economy $e$ reveal unsettled average fitness score behavior over generations 1 through 100 in the form of a wallflower collapse (1 run), bubbles (2 runs), regime shifts (6 runs), and persistent drifting (4 runs). It was at first conjectured that this observed instability might be due to the small population size of 8 for each agent type. Surprisingly, however, when the experiments were re-run with an increased population size of 12 for each agent type, keeping all other parameter values fixed, the resulting distance values, worksite behaviors, and social welfare outcomes closely resembled those obtained for the smaller population size.

It therefore appears that the instabilities observed in average fitness scores for these individual runs may instead be due to the fluid role played by worker-employers. In particular, the ability of worker-employers to function either as workers or as employers permits them to crowd out

the pure workers or the pure employers, causing them to degenerate into PWFs. In addition, worker-employers have the unique ability to form a self-sufficient network of contractual relationships without the participation of either pure workers or pure employers. Indeed, the persistent contractual networks for the second distance cluster in Table 7(a) are characterized by degeneracies of this type.

As job capacity is incrementally tightened, the risk to pure employers of high wallflower payoff accumulation recedes and is replaced by the risk to pure workers of high refusal payoff accumulation. As seen in Table 7, the increasingly favorable structural setting for pure employers tends to encourage increased UD behavior by pure employers and to discourage UD behavior by pure workers. Consequently, there is an increased tendency for the flexible worker-employer agents to behave as pure employers, in the sense that they continue to receive work offers but they ultimately stop making any work offers themselves.

This tendency is seen in the changing nature of the base contractual patterns depicted in Figure 4, which give the most predominant types of contractual networks that form as job capacity is decreased from high excess to tight. In particular, as seen in Figure 4(c) and Table 7(c), in 75% of the runs for the tight job capacity case nearly all of the worker-employers behave as pure employers in their persistent relationships. When job capacity becomes extremely tight, however, Figure 4(d) and Table 7(d) show that complete coordination failure occurs in 75% of the runs, in the sense that all agents in these runs degenerate into persistent wallflowers.

Finally, comparing part (a) of Figure 4 with parts (b), (c), and (d), note the extraordinarily strong disciplinary role played by ex ante capacity constraints in the determination of evolutionary outcomes for partially-fluid labor market economies. For example, as one moves from high excess job capacity in part (a) to tight job capacity in part (b), the economy moves from diffusive work offers to disjoint latched triads consisting of one pure worker, one worker-employer, and one pure employer, with welfare outcomes shifting decidedly in favor of the pure employers.

Indeed, as indicated in Figures 1 and 3 as well, ex ante capacity constraints play a strong

coordinating role in all of the previously reported experimental findings. As suggested by Gode and Sunder (1993), when attempting to understand the cause of perceived regularities in market outcomes, it is important to carefully separate the effects of institutional constraints per se from the effects of the cognitive functioning of the agents participating in the market. It will be interesting to determine, in future studies, the extent to which the network, behavior, and welfare outcomes determined for the experiments at hand are retained under similar ex ante capacity constraints as the modelling for agent cognition is varied.

## VI. Concluding Remarks

In order for the TNG labor market framework to become a useful tool for illuminating real-world labor market processes, it must be better adapted to specific labor market contexts using natural data, survey data, and human-subject laboratory data.

For example, as stressed in the survey study by Bewley (1998), economists need to take better account in their labor market models of the importance of social and psychological factors in determining worksite interactions and outcomes. The current TNG labor market modelling of worksite interactions as two-person games played between autonomous workers and employers with internal cognitive structure is only a first small step in this direction. In future studies, signalling among agents (e.g., wage bids and offers) needs to be introduced, and capacity constraints must be endogenized so that they are determined as a function of past decisions and events. Also, the role of government regulations (e.g., minimum wage laws) must be considered. Finally, the labor market framework needs to be imbedded in a more complete modelling of a decentralized market economy.

Nevertheless, it is hoped that the findings presented in the previous section illustrate how an agent-based computational approach can facilitate the rigorous and routine experimental study of non-steady-state labor market dynamics.

## REFERENCES


**Aghion, Philippe, and Howitt, Peter.**, "Growth and Unemployment." *Review of Economic Studies*, 1994, Vol. 61, pp. 477–494.

**Banks, David, and Carley, Kathleen.** "Metric Inference for Social Networks." *Journal of Classification*, 1994, Vol. 11, pp. 121–149.

**Bewley, Truman.** *Listening to Business: A Study of Wage Rigidity*. Book Manuscript, Cowles Foundation, Yale University, April 1998.

**Diamond, Peter.** "Aggregate Demand Management in Search Equilibrium." *Journal of Political Economy*, 1982, Vol. 90, pp. 881–895.

**Diamond, Peter.** *Search Equilibrium Approach to the Micro Foundations of Macroeconomics.*, Cambridge, MA: The MIT Press, 1984.

**Gale, David, and Shapley, Lloyd.** "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 1962, Vol. 69, pp. 9–15.

**Gode, Dhananjay K., and Sunder, Shyam.** "Allocative Efficiency in Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality." *Journal of Political Economy*, 1993, Vol. 101, pp. 119–137.

**Goldberg, David.** *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley, 1989.

**Hosios, Arthur.** "On the Efficiency of Matching and Related Models of Search and Unemployment," *Review of Economic Studies*, 1990, Vol. 57, pp. 279-298.

**Jackson, Matthew O., and Wolinsky, Asher.** "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 1996, Vol. 71, pp. 44–74.

**Jovanovic, Boyan.** "Job Matching and the Theory of Turnover." *Journal of Political Economy*, 1979, Vol. 87, pp. 972–990.

**Kirman, Alan, Oddou, C., and Weber, S.** "Stochastic Communication and Coali-

tion Formation." *Econometrica*, 1986, Vol. 54, pp. 129–138.

**MacLeod, W. Bentley, and Malcomson, James M.** "Motivation and Markets." *American Economic Review*, June 1998, Vol. 88(3), pp. 388–411.

**McFadzean, David.** *SimBioSys: A class framework for evolutionary simulations.* Master's Thesis, Department of Computer Science, University of Calgary, Alberta, Canada, 1995.

**McFadzean, David, and Tesfatsion, Leigh.** "A C++ Platform for the Evolution of Trade Networks." Economic Report No. 39, Iowa State University, September 1997, to appear in *Computational Economics*.

**Mitchell, Melanie, and Forrest, Stephanie.** "Genetic Algorithms and Artificial Life," *Artificial Life*, 1994, Vol. 1, pp. 267–289.

**Montgomery, James D.** "Social Networks and Labor Market Outcomes: Toward an Economic Analysis." *American Economic Review*, December 1991, Vol. 81 (No.5), pp. 1408–1418.

**Mortensen, Dale T., and Pissarides, Christopher A.** "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, 1994, Vol. 61, pp. 397–415.

**Myerson, Roger.** "Graphs and Cooperation in Games." *Mathematics and Operations Research*, 1977, Vol. 2, pp. 225–229.

**Orbell, John. M., and Dawes, Robyn M.** "Social Welfare, Cooperators' Advantage, and the Option of Not Playing the Game." *American Sociological Review*, December 1993, Vol. 58, pp. 787–800.

**Pissarides, Christopher A.** "Short-Run Dynamics of Unemployment, Vacancies, and Real Wages." *American Economic Review*, 1985, Vol. 75, pp. 676–690.

**Pissarides, Christopher A.** *Equilibrium Unemployment Theory.* Oxford, UK: Oxford University Press, 1990.

**Roth, Alvin, and Sotomayor, Marilda A. O.** *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge, UK: Cambridge University Press, 1990.

**Rubinstein, Ariel, and Wolinsky, Asher.** "Equilibrium in a Market with Sequential Bargaining." *Econometrica*, 1985, Vol. 53, pp. 1133–1150.

**Rubinstein, Ariel, and Wolinsky, Asher.** "Decentralized Trading, Strategic Behavior, and the Walrasian Outcome." *Review of Economic Studies*, 1990, Vol. 57, pp. 63–78.

**Tesfatsion, Leigh.** "A Trade Network Game with Endogenous Partner Selection." ISU Economic Report No. 36, April 1995.

**Tesfatsion, Leigh.** "A Trade Network Game with Endogenous Partner Selection," in Hans Amman, Berc Rustem, and Andrew Whinston, eds., *Computational Approaches to Economic Problems*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1997a, pp. 249–269.

**Tesfatsion, Leigh.** "How Economists Can Get Alife," in W. Brian Arthur, Steven Durlauf, and David Lane, eds., *The Economy as an Evolving Complex System, II*. Proceedings Volume XXVII, Santa Fe Institute Studies in the Sciences of Complexity, Reading, MA: Addison-Wesley, 1997b, pp. 533–564.

**Tesfatsion, Leigh.** "Gale-Shapley Matching in an Evolutionary Trade Network Game." Economic Report No. 43, Iowa State University, April 1998a.

**Tesfatsion, Leigh.** "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*. Proceedings of the Seventh Annual Conference on Evolutionary Programming, Berlin: Springer-Verlag, 1998b, pp. 15–24.

**van den Nouweland, Anne.** *Games and Graphs in Economic Situations*. Ph.D.

disseration, Tilburg University, 1993.

**Watts, Duncan J., and Steven H. Strogatz.** "Collective Dynamics of 'Small-World' Networks," *Nature*, June 1998, Vol. 393, pp. 440–442.

**Wolinsky, Asher.** "Matching, Search and Bargaining." *Journal of Economic Theory*, 1987, Vol. 42, pp. 311-333.

Copyright © 1998 Leigh Tesfatsion. All Rights Reserved.

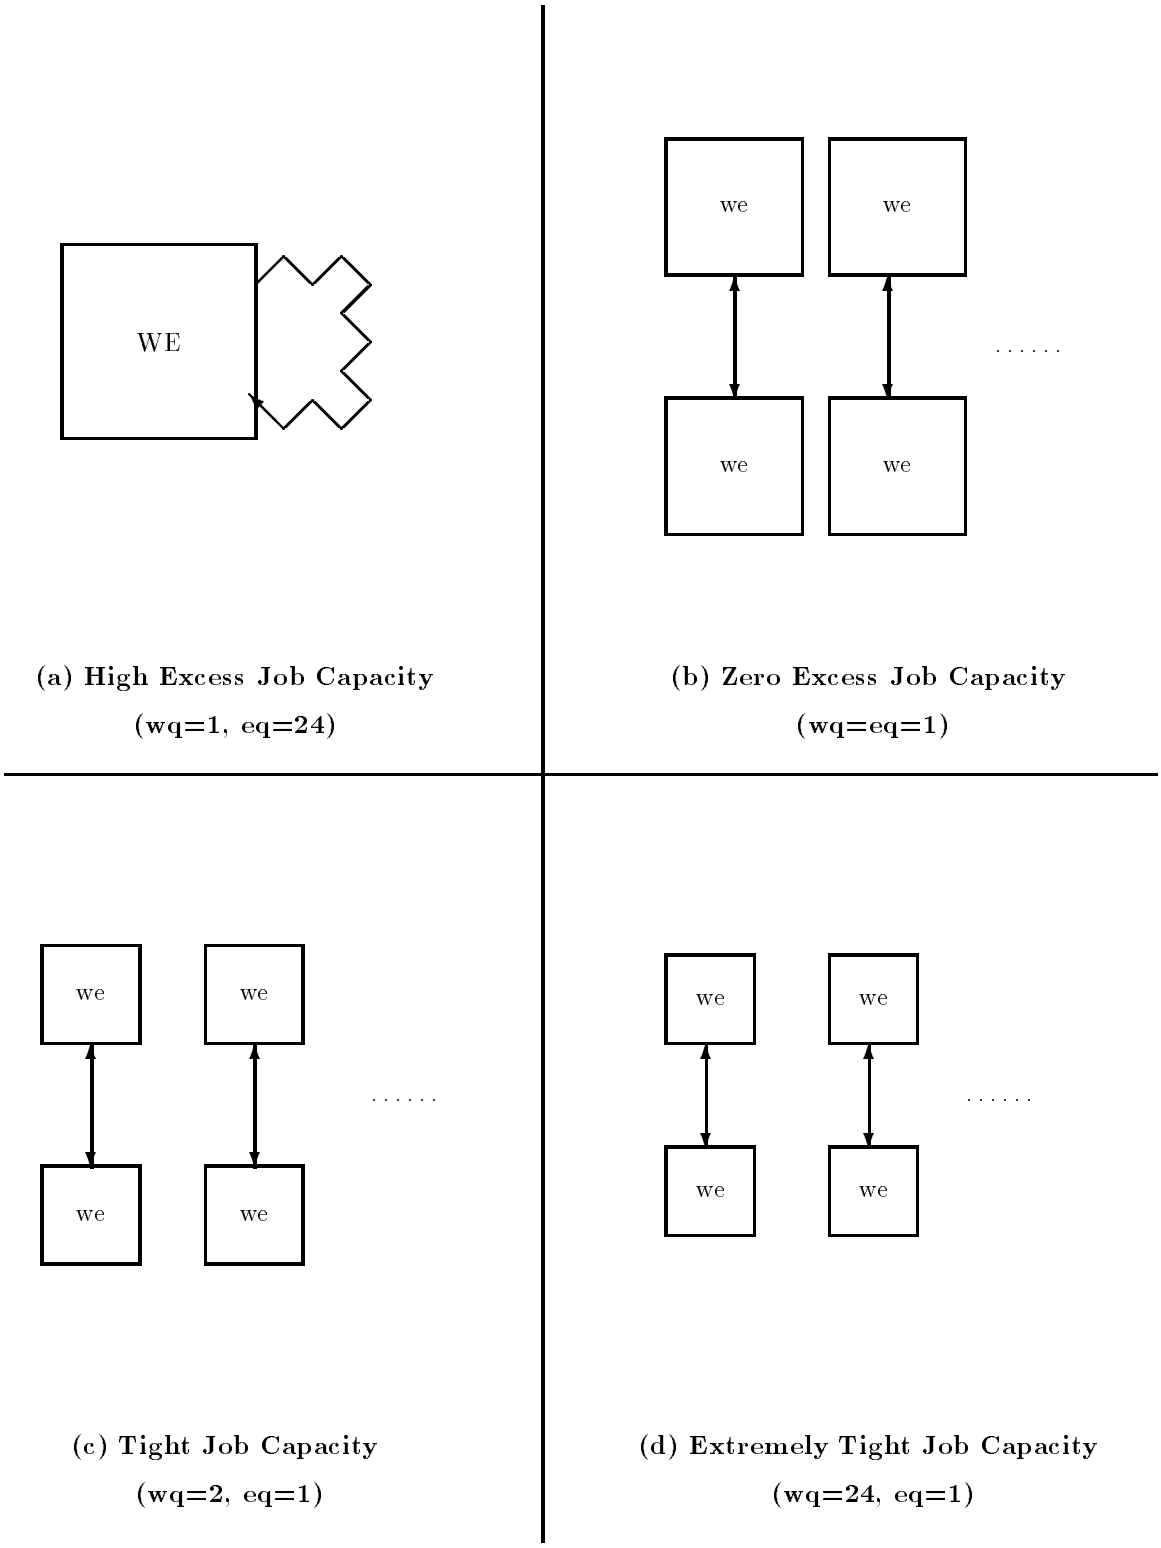


**(a) High Excess Job Capacity**
**(wq=1, eq=24)**

**(b) Zero Excess Job Capacity**
**(wq=eq=1)**

**(c) Tight Job Capacity**
**(wq=2, eq=1)**

**(d) Extremely Tight Job Capacity**
**(wq=24, eq=1)**

**Figure 1: Base Contractual Patterns for Endogenous-Type Labor Markets with Different Ex Ante Capacities.** A relatively larger box for the worker-employers (WE) under a particular capacity specification indicates that the worker-employers achieve a relatively higher mean FIT value under this capacity specification in realized economies whose contractual networks approximate the base contractual pattern. Straight directed edges indicate continuous persistent relationships (latching) and zig-zag directed edges indicate intermittent persistent relationships.

| $D^o$ Cluster | % Runs | Mean UD | Mean PWF | Mean PC | Mean FIT |
|---|---|---|---|---|---|
| 0–3 | 90% | 3%<br>(3%) | 1%<br>(2%) | 96%<br>(4%) | 1.36<br>(.05) |
| 11 | 5% | 8%<br>(NA) | 0%<br>(NA) | 38%<br>(NA) | 1.34<br>(NA) |
| 23 | 5% | 96%<br>(NA) | 0%<br>(NA) | 88%<br>(NA) | 1.13<br>(NA) |

Table 5(a): High Excess Job Capacity (wq=1, eq=24)

| $D^o$ Cluster | % Runs | Mean UD | Mean PWF | Mean PC | Mean FIT |
|---|---|---|---|---|---|
| 0–4 | 100% | 8%<br>(22%) | 2%<br>(3%) | 96%<br>(5%) | 1.19<br>(.06) |

Table 5(b): Zero Excess Job Capacity (wq=eq=1)

| $D^o$ Cluster | % Runs | Mean UD | Mean PWF | Mean PC | Mean FIT |
|---|---|---|---|---|---|
| 2–11 | 85% | 1%<br>(2%) | 1%<br>(2%) | 98%<br>(4%) | 0.95<br>(.04) |
| 24 | 15% | 100%<br>(0%) | 100%<br>(0%) | 0%<br>(0%) | -0.17<br>(.00) |

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

| $D^o$ Cluster | % Runs | Mean UD | Mean PWF | Mean PC | Mean FIT |
|---|---|---|---|---|---|
| 2–12 | 70% | 3%<br>(5%) | 3%<br>(7%) | 94%<br>(9%) | 0.89<br>(.07) |
| 24 | 30% | 100%<br>(0%) | 100%<br>(0%) | 0%<br>(0%) | -0.16<br>(.00) |

Table 5(d): Extremely Tight Job Capacity (wq=24, eq=1)

**Table 5. Experimental Findings for Endogeneous-Type Labor Markets with Different Ex Ante Capacities**

**THIS FIGURE IS AVAILABLE FROM THE AUTHOR UPON REQUEST**

**Figure 2: Stability Check for an Endogenous-Type Economy with High Excess Job Capacity.** The maximum, minimum, and average fitness scores are graphed for agent generations 1 through 100. By generation 25 the average fitness scores closely fluctuate around 1.40, the mutual cooperation payoff level.

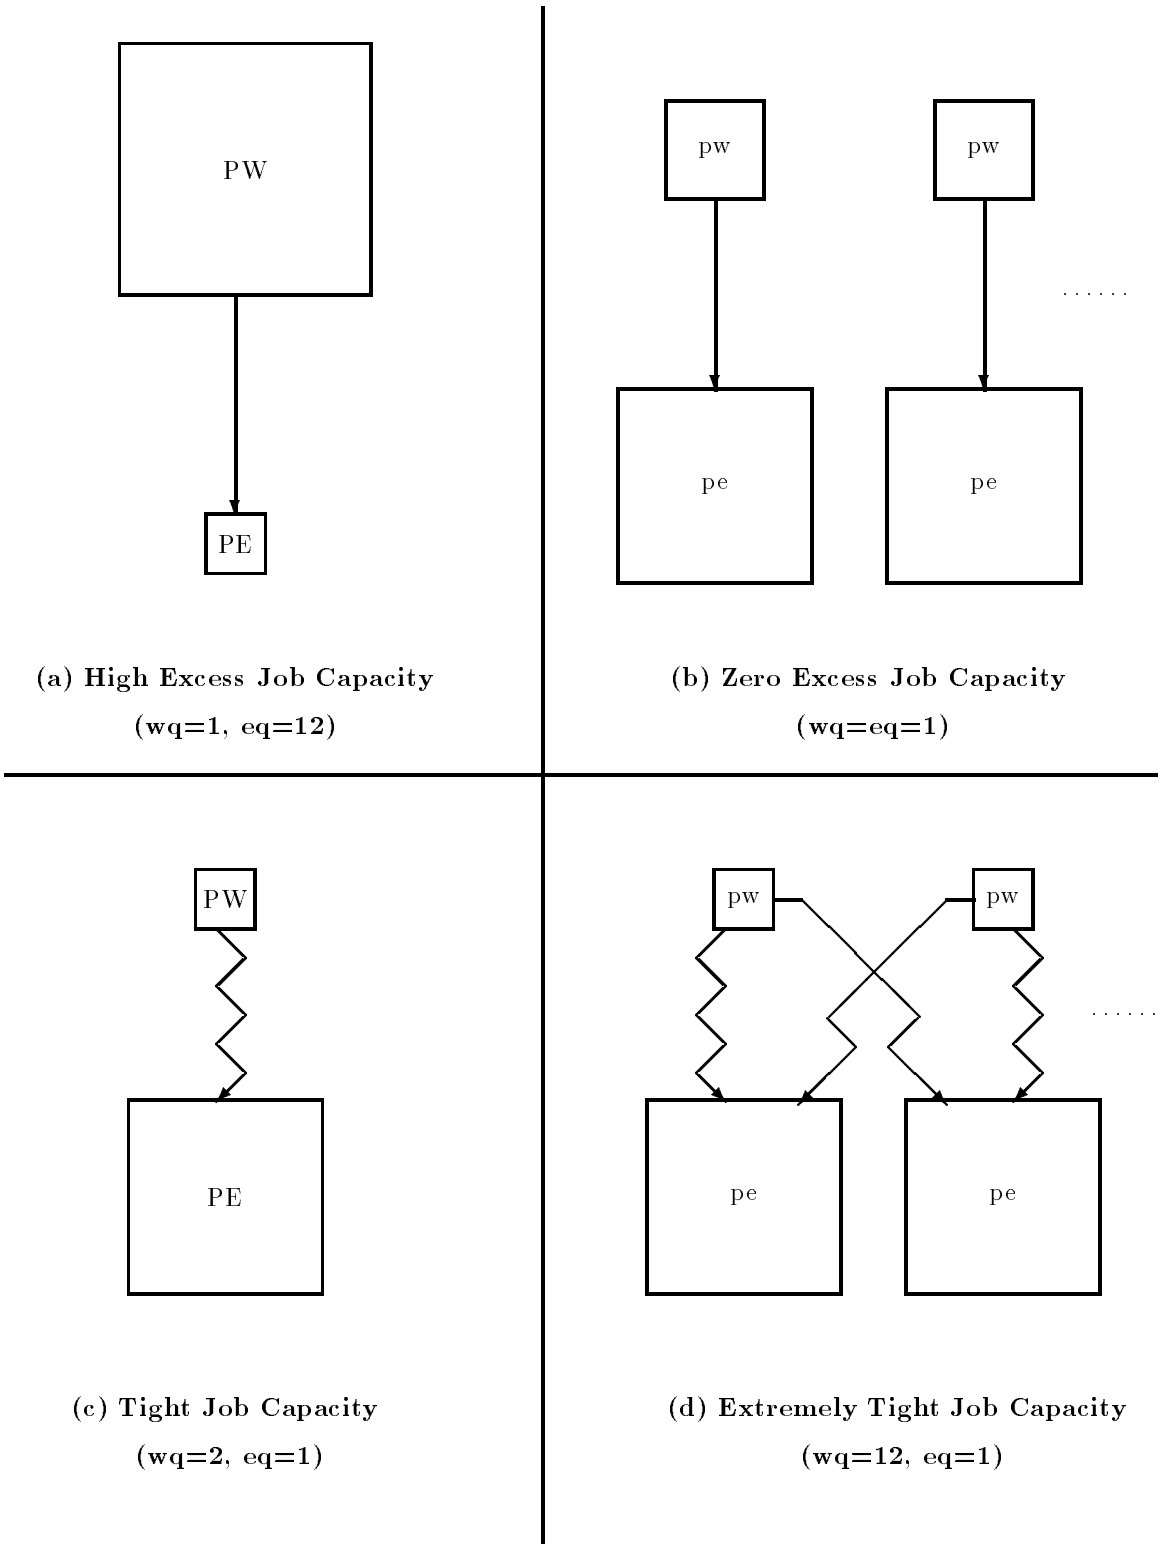


**Figure 3: Base Contractual Patterns for Two-Sided Labor Markets with Different Ex Ante Capacities.** A relatively larger box for an agent type — pure workers (PW) or pure employers (PE) — under a particular capacity specification indicates that this agent type achieves a relatively higher mean FIT value under this capacity specification in the realized economies whose contractual networks approximate the base contractual pattern. Straight directed edges indicate continuous persistent relationships (latching) and zig-zag directed edges indicate intermittent persistent relationships.

| $D^o$ Cluster | % Runs | Mean UD | | Mean PWF | | Mean PC | | Mean FIT | |
|---|---|---|---|---|---|---|---|---|---|
| | | pw | pe | pw | pe | pw | pe | pw | pe |
| 3–9 | 75% | 97% (5%) | 16% (34%) | 2% (3%) | 40% (12%) | 3% (5%) | 39% (28%) | 1.74 (.27) | 0.35 (.14) |
| 23–24 | 25% | 2% (3%) | 5% (7%) | 2% (3%) | 5% (7%) | 98% (3%) | 95% (7%) | 1.39 (.02) | 1.02 (.03) |

Table 6(a): High Excess Job Capacity (wq=1, eq=12)

| $D^o$ Cluster | % Runs | Mean UD | | Mean PWF | | Mean PC | | Mean FIT | |
|---|---|---|---|---|---|---|---|---|---|
| | | pw | pe | pw | pe | pw | pe | pw | pe |
| 0–2 | 80% | 15% (32%) | 22% (38%) | 1% (3%) | 1% (3%) | 94% (6%) | 86% (25%) | 1.07 (.20) | 1.34 (.21) |
| 4 | 5% | 100% (NA) | 100% (NA) | 17% (NA) | 17% (NA) | 0% (NA) | 0% (NA) | 0.62 (NA) | 0.29 (NA) |
| 24 | 15% | 0% (0%) | 22% (20%) | 0% (0%) | 8% (0%) | 89% (16%) | 78% (20%) | 0.24 (.08) | 1.42 (.05) |

Table 6(b): Zero Excess Job Capacity (wq=eq=1)

| $D^o$ Cluster | % Runs | Mean UD | | Mean PWF | | Mean PC | | Mean FIT | |
|---|---|---|---|---|---|---|---|---|---|
| | | pw | pe | pw | pe | pw | pe | pw | pe |
| 0–7 | 55% | 2% (3%) | 5% (9%) | 19% (10%) | 4% (7%) | 81% (10%) | 96% (6%) | 0.30 (.05) | 1.35 (.09) |
| 24 | 45% | 100% (0%) | 90% (28%) | 82% (26%) | 77% (34%) | 3% (8%) | 5% (13%) | 0.04 (.20) | 0.22 (.39) |

Table 6(c): Tight Job Capacity (wq=2, eq=1)

| $D^o$ Cluster | % Runs | Mean UD | | Mean PWF | | Mean PC | | Mean FIT | |
|---|---|---|---|---|---|---|---|---|---|
| | | pw | pe | pw | pe | pw | pe | pw | pe |
| 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 6(d): Extremely Tight Job Capacity (wq=12, eq=1)

**Table 6. Experimental Findings for Two-Sided Labor Markets with Different Ex Ante Capacities**

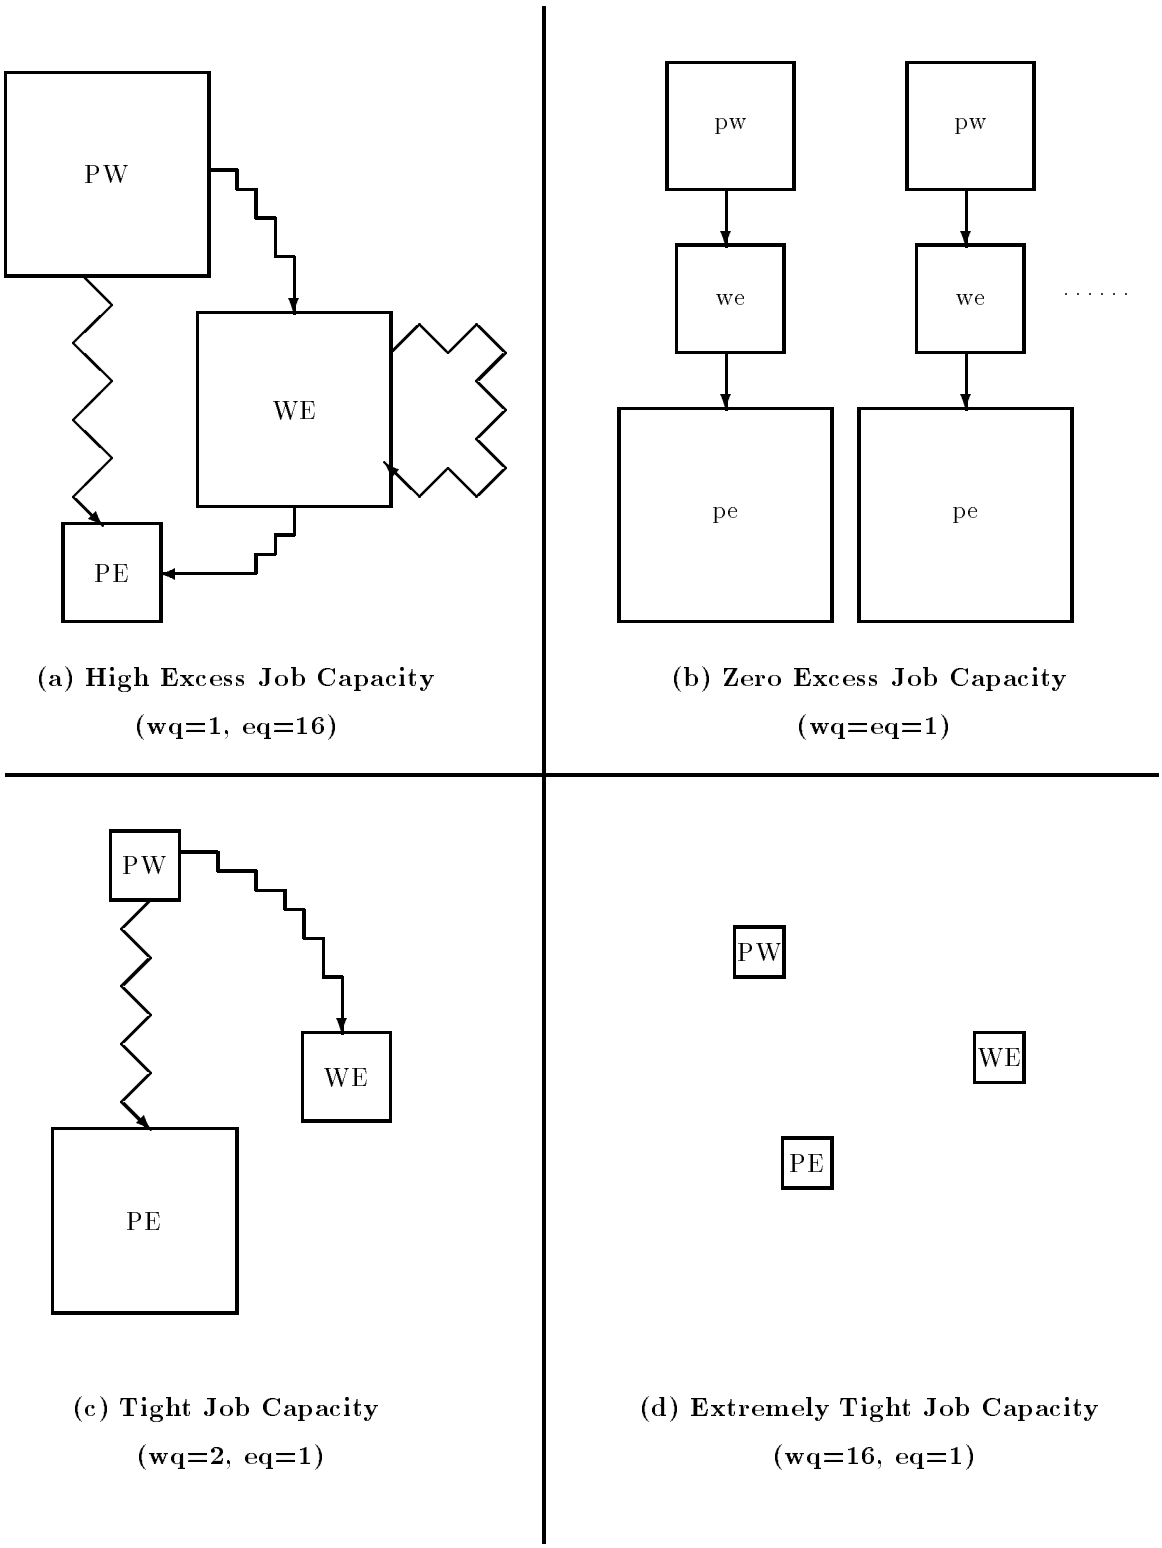


**Figure 4: Base Contractual Patterns for Partially-Fluid Labor Markets with Different Ex Ante Capacities.** A relatively larger box for an agent type — pure workers (PW), worker-employers (WE), or pure employers (PE) — under a particular capacity specification indicates that this agent type achieves a relatively higher mean FIT value under this capacity specification in the realized economies whose contractual networks approximate the base contractual pattern. Straight directed edges indicate continuous persistent relationships (latching) and zig-zag directed edges indicate intermittent persistent relationships.

| $D^o$ Clst. | % of Runs | Mean UD | | | Mean PWF | | | Mean PC | | | Mean FIT | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | pw | pe | we | pw | pe | we | pw | pe | we | pw | pe | we |
| 0–2 | 30% | 2% (5%) | 2% (5%) | 6% (10%) | 2% (5%) | 2% (5%) | 0% (0%) | 98% (5%) | 98% (5%) | 81% (32%) | 1.38 (.01) | 1.03 (.03) | 1.38 (.03) |
| 6–9 | 35% | 25% (40%) | 41% (48%) | 39% (41%) | 14% (35%) | 41% (48%) | 0% (0%) | 75% (40%) | 48% (41%) | 64% (38%) | 1.16 (.11) | 0.73 (.24) | 1.25 (.09) |
| 16–21 | 35% | 98% (4%) | 23% (38%) | 98% (4%) | 16% (35%) | 30% (30%) | 2% (4%) | 18% (34%) | 40% (37%) | 21% (35%) | 1.13 (.63) | 0.55 (.25) | 1.43 (.30) |

Table 7(a): High Excess Job Capacity (wq=1, eq=16)

| $D^o$ Clst. | % of Runs | Mean UD | | | Mean PWF | | | Mean PC | | | Mean FIT | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | pw | pe | we | pw | pe | we | pw | pe | we | pw | pe | we |
| 0–6 | 80% | 6% (24%) | 20% (10%) | 9% (24%) | 0% (0%) | 1% (3%) | 2% (4%) | 98% (4%) | 92% (24%) | 97% (5%) | 1.16 (.06) | 1.42 (.14) | 1.11 (.13) |
| 16–24 | 20% | 100% (0%) | 28% (42%) | 47% (47%) | 28% (42%) | 25% (31%) | 6% (11%) | 0% (0%) | 19% (14%) | 53% (47%) | 0.20 (.17) | 0.87 (.33) | 0.87 (.22) |

Table 7(b): Zero Excess Job Capacity (wq=eq=1)

| $D^o$ Clst. | % of Runs | Mean UD | | | Mean PWF | | | Mean PC | | | Mean FIT | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | pw | pe | we | pw | pe | we | pw | pe | we | pw | pe | we |
| 0–7 | 75% | 2% (4%) | 19% (38%) | 17% (34%) | 6% (8%) | 4% (7%) | 1% (3%) | 75% (28%) | 80% (38%) | 83% (34%) | 0.63 (.19) | 1.30 (.19) | 0.90 (.07) |
| 10 | 10% | 0% (0%) | 50% (50%) | 50% (50%) | 19% (6%) | 44% (44%) | 50% (50%) | 81% (6%) | 56% (44%) | 50% (50%) | 0.32 (.10) | 0.85 (.53) | 0.46 (.49) |
| 24 | 15% | 100% (0%) | 100% (0%) | 100% (0%) | 100% (0%) | 100% (0%) | 100% (0%) | 0% (0%) | 0% (0%) | 0% (0%) | -0.13 (.00) | -0.02 (.00) | -0.13 (.00) |

Table 7(c): Tight Job Capacity (wq=2, eq=1)

| $D^o$ Clst. | % of Runs | Mean UD | | | Mean PWF | | | Mean PC | | | Mean FIT | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | pw | pe | we | pw | pe | we | pw | pe | we | pw | pe | we |
| 0 | 75% | 100% (0%) | 100% (0%) | 100% (0%) | 100% (0%) | 100% (0%) | 100% (0%) | 0% (0%) | 0% (0%) | 0% (0%) | -0.13 (.00) | -0.01 (.00) | -0.13 (.00) |
| 16 | 5% | 0% (NA) | 0% (NA) | 100% (NA) | 25% (NA) | 0% (NA) | 25% (NA) | 38% (NA) | 88% (NA) | 0% (NA) | 0.52 (NA) | 1.33 (NA) | 0.57 (NA) |
| 23–24 | 20% | 0% (0%) | 6% (11%) | 3% (5%) | 0% (0%) | 3% (5%) | 0% (0%) | 97% (5%) | 97% (5%) | 97% (5%) | 0.39 (.07) | 1.36 (.07) | 0.79 (.04) |

Table 7(d): Extremely Tight Job Capacity (wq=16, eq=1)

**Table 7. Experimental Findings for Partially-Fluid Labor Markets with Different Ex Ante Capacities**