

Structure, behavior, and market power in an evolutionary labor market with adaptive search[☆]

Leigh Tesfatsion*

Department of Economics, Iowa State University, Heady Hall 375, Ames, IA 50011-1070, USA

Accepted 22 February 2000

Abstract

This study uses an agent-based computational labor market framework to experimentally study the relationship between job capacity, job concentration, and market power. Job capacity is measured by the ratio of potential job openings to potential work offers, and job concentration is measured by the ratio of work suppliers to employers. For each experimental treatment, work suppliers and employers repeatedly seek preferred worksite partners based on continually updated expected utility, engage in efficiency-wage worksite interactions modelled as prisoner's dilemma games, and evolve their worksite behaviors over time. The main finding is that job capacity consistently trumps job concentration when it comes to predicting the relative ability of work suppliers and employers to exercise market power. © 2001 Elsevier Science B.V. All rights reserved.

JEL classification: C6; C7; J6; L1

Keywords: Labor market dynamics; Market power; Job capacity; Job concentration; Adaptive search; Networks; Endogenous interactions; Agent-based computational economics; Evolutionary game

1. Introduction

Market power refers to the ability of sellers or buyers to exert a perceptible control over market outcomes that enables them to attain higher individual

[☆]The author is grateful to M. Kilkenny, P. Orazem, M. Pingle, H.C. Quirnbach, T. Wier, and two anonymous reviewers for helpful comments.

* Tel.: + 1-515-294-0138; fax: + 1-515-294-0221.

E-mail address: tesfatsi@iastate.edu (L. Tesfatsion).

welfare levels than they would achieve under competitive market conditions. Understanding the relationship between market structure, market behavior, and market power in markets with multiple agents engaged in repeated strategic interactions has been a major focus of analytical, empirical, and human-subject experimental researchers in industrial organization since the early 1970s. To date, however, definitive conclusions have been difficult to obtain.

For example, Tirole (1988, Part II) presents a unified theoretical treatment of oligopoly decision-making in terms of noncooperative game theory. He focuses on the choice of price, capacity, product positioning, research and development, and other strategic variables. Only equilibrium behavior is considered, however; and a common finding for his games with incomplete information is that multiple equilibria exist with widely differing characteristics. The question then arises: given agents with incomplete information and limited computational capabilities, under what conditions would these agents learn to coordinate on one type of equilibrium versus another, and what would be the resulting dynamic implications for market power?

In a survey of empirical work on market power in industrial organization, Bresnahan (1989, pp. 1051–1055) summarizes his overall findings as follows: ‘although the (new empirical industrial organization) has had a great deal to say about measuring market power, it has had very little, as yet, to say about the causes of market power’. Holt (1995, Section VII) notes that, although the nonmonopolized double auction is widely used in experimental research with human subjects, whether market power has any efficiency effects in this context remains an open issue. For posted-offer auctions, Holt points out that capacity constraints and some forms of transactions costs have reliably produced supra-competitive prices in experiments performed by himself and several other researchers; but so far the number of experiments along these lines has been small.

This lack of definitive results reflects the complex nature of market power in actual real-world markets. Given this complexity, it would seem useful to complement these previous approaches to the study of market power with controlled computational experiments.

This paper investigates the evolution of market power in the context of a computational labor market framework with strategically interacting work suppliers and employers.¹ As will be clarified in Section 2, the labor market

¹ A preliminary version of this labor market framework was presented 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. A version of the framework with a different fitness measure is also used in Tesfatsion (1999) to examine hysteresis (path dependency) in labor markets. The framework is an example of agent-based computational economics (ACE), the computational study of economies modelled as evolving systems of autonomous interacting agents. For various ACE-related resources, including surveys, readings, software, and pointers to researchers and research groups, see the ACE web site at <http://www.econ.iastate.edu/tesfatsi/ace.htm>.

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 market power in relation to market structure and market behavior over time by specifying these exogenous aspects in relatively simple terms.

Thus, as implemented for this study, the labor market framework comprises multiple work suppliers and employers who 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. All work suppliers are assumed to have the same size wq , where wq is the maximum number of potential work offers that each work supplier can make. Similarly, all employers are assumed to have the same size eq , where eq is the maximum number of job openings that each employer can provide. Moreover, there is no entry into, or exit from, the labor market; the number NW of work suppliers and the number NE of employers are both held fixed during the course of each experimental run.

Market power for work suppliers is measured by the degree to which their average attained welfare level deviates from the average welfare level that they would obtain in a competitive (full employment) market outcome under the assumption of mutually cooperative worksite behavior. Market power for employers is similarly defined.

Intuitively, it seems reasonable to postulate that the extent to which market power accrues to work suppliers or employers in the labor market framework depends in part on *job capacity*, as measured by the ratio $(NE \cdot eq)/(NW \cdot wq)$ of total potential job openings to total potential work offers, and on *job concentration* as measured by the ratio NW/NE of the total number of work suppliers to the total number of employers. Job capacity measures the total potential availability of job openings relative to work offers, whereas job concentration measures the extent to which control over job openings is concentrated among relatively few employers.² By construction, for any given level of job concentration, job capacity varies inversely with the size wq of each work supplier and directly with the size eq of each employer.

The joint implications of job capacity and job concentration for the exercise of market power are not easy to predict a priori. For example, consider the overall effect on welfare and market power when the labor market framework

² The potential importance of distinguishing between capacity and concentration effects has been emphasized by Holt (1995, p. 396) in the context of posted-offer auction experiments: 'By reassigning units of capacity from one seller to another, there is no change in either market supply or excess supply at supra-competitive prices, yet market power can change'.

comprises 12 work suppliers, each able to make one work offer, together with six employers who each have four job openings. The excess supply of job openings favors monopoly power by work suppliers, but the concentration of job openings in relatively few hands favors monopsony power by employers.

The experimental design of this study consists of the systematic variation, from high to low, of both job capacity and job concentration. For each given specification of these two market structure conditions, 20 different runs are generated using 20 different pseudo-random number seed values.³ In examining the resulting run histories, particular attention is focused on the experimental determination of correlations between market structure and the formation of networks among work suppliers and employers, and between network formations and the types of worksite behaviors, welfare outcomes, and market power outcomes that these networks support.

The primary objective of this study is to test the following four hypotheses regarding the impact of job capacity and job concentration on the ability of work suppliers and employers to exercise market power:

- H1 (*Relative market power hypothesis for job capacity*). (a) The ability of work suppliers to exercise market power is less than the ability of employers to exercise market power in conditions of tight or balanced job capacity; and (b) the ability of work suppliers to exercise market power is greater than the ability of employers to exercise market power in conditions of excess job capacity.
- H2 (*Relative market power hypothesis for job concentration*). (a) The ability of work suppliers to exercise market power is less than the ability of employers to exercise market power in conditions of high or balanced job concentration; and (b) the ability of work suppliers to exercise market power is greater than the ability of employers to exercise market power in conditions of low job concentration.
- H3 (*Job capacity sensitivity hypothesis*). (a) The ability of work suppliers to exercise market power increases as job capacity increases, all else equal; and (b) the ability of employers to exercise market power decreases as job capacity increases, all else equal.
- H4 (*Job concentration sensitivity hypothesis*). (a) The ability of work suppliers to exercise market power increases as job concentration decreases, all else equal; and (b) the ability of employers to exercise market power decreases as job concentration decreases, all else equal.

³ All experiments reported in this study are implemented using version 105c 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 <http://www.econ.iastate.edu/tesfatsi/accocode.htm>, along with extensive user instructions.

Hypotheses H3 and H4 can hold simultaneously without contradiction. However, hypotheses H1(b) and H2(a) yield contradictory predictions in conditions of excess job capacity and high or balanced job concentration, and hypotheses H1(a) and H2(b) yield contradictory predictions in conditions of tight or balanced job capacity and low job concentration.

The main finding of this study at the aggregate data level is that job capacity is the dominant factor affecting the ability of work suppliers and employers to exercise market power. Aggregate market power outcomes strongly support the job capacity hypotheses H1 and H3(a) and weakly support the job capacity hypothesis H3(b), but they provide little support for either of the job concentration hypotheses H2 or H4. Surprisingly, controlling for job capacity, job concentration has only small unsystematic effects on attained market power levels.

To better understand these aggregate market power findings, the complicated nonlinear relations linking market structure, market behavior, and market power are carefully examined at a more disaggregated level. This examination reveals even stronger support for the aggregate data finding that job capacity is the key variable determining the relative market power of work suppliers and employers. For example, in contrast to aggregate data indications, it is shown that the job capacity hypothesis H3(b) is strongly supported if a small number of sample economies are omitted for which complete coordination failure occurs and all agents perform extremely poorly. The disaggregated data also highlight the importance of job search costs and behavioral flexibility (provocability) for the realistic assessment of market power opportunities *ex ante* and the accurate measurement of market power *ex post*.

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, welfare outcomes, and market power outcomes. The experimental design of the study is explained in Section 4, and experimental findings are discussed in Section 5. Concluding remarks are given in Section 6.

2. Labor market framework

The labor market framework outlined below differs in several essential respects from standard labor market models (e.g., Ehrenberg and Smith, 1997). First, it is a dynamic process model defined algorithmically in terms of the internal states and behavioral rules characterizing work suppliers and employers rather than by the usual system of demand, supply, and equilibrium equations. The only equations that explicitly appear 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.

Table 1
General form of the internal structure of an agent

```

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.
};

```

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 a researcher but not disturbed, is apt.

The labor market framework comprises NW work suppliers who make work offers and NE employers who receive work offers, where NW and NE can be any positive integers. 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.⁴ Although highly simplified, these parametric specifications will be seen in Section 4, to permit the study of labor markets operating under a variety of job capacity and job concentration conditions.

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

⁴ 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.

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 rules⁵ 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.

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 and initial expected utility assessments regarding his potential worksite partners. The work suppliers and employers then enter into a *trade cycle loop* during which they repeatedly search for preferred 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 employers each separately evolve (structurally modify) their worksite behavioral rules based on the past utility outcomes secured with these rules, and a new generation then commences.

The particular module specifications used in all experiments reported below will now be described in roughly the order depicted in Table 2.⁶

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).⁷ Under the terms of this auction, hereafter referred to as the *deferred choice and refusal* (DCR) mechanism, each work supplier first submits work offers to a maximum of wq employers he ranks as most preferable on the basis of expected utility and whom he judges to be tolerable in the sense that their expected utility is not negative. Each employer then selects up to eq of the work offers he has received to date that he finds

⁵ 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.

⁶ McFadzean and Tesfatsion (1999) provide a detailed discussion of these module specifications. The source code implementations for these module specifications – TNG version 105c – can be downloaded as freeware; see footnote 3.

⁷ 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.

Table 2
Logical flow of the labor market framework

```

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

                                // Environment Step:
        AssessFitness();       // Work suppliers and employers
                                // assess their utility levels.

                                // Evolution Step:
        EvolveGen();          // Work suppliers and employers separately
                                // evolve their worksite strategies, and
                                // a new generation cycle begins.
    }

    Return 0;
}

```

tolerable and most preferable on the basis of expected utility, and he places these selected work offers on a waiting list; all other work offers are refused. Work suppliers who have work offers refused then redirect these work offers to any tolerable preferred employers who have not yet refused them, and the process

repeats. Once an employer stops receiving new work offers, he accepts all work offers currently on his waiting list.

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.⁸ A work supplier or employer who neither submits nor accepts 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, work suppliers and employers can each shirk on the worksite, to the detriment of the other, and can possibly improve their own welfare by doing so. Work suppliers can reduce their disutility of work in the short run by not working as hard as their employers expect, and employers can enhance their profit in the short run by not providing benefits their work suppliers expect to receive. Offsetting these incentives are factors that discourage shirking. Employers can punish shirking work suppliers by firing them (i.e., by refusing all future work offers), and work suppliers can punish shirking employers by quitting (i.e., by redirecting future work offers elsewhere).

These various possibilities are captured by having each matched work supplier and employer engage in a worksite interaction modelled as a two-person prisoner's dilemma game. The work supplier can either cooperate (exert high work effort) or defect (shirk). 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).

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 $L < D < 0 < C < H$.

⁸ This modelling for job search costs is equivalent to assuming the following two conditions: (i) each work supplier must pay a job search cost for each work offer he makes to an employer; and (ii) any work supplier whose work offer is accepted is able to recoup the job search costs he incurred in making this work offer through his subsequent worksite payoff.

Table 3
Payoff matrix for the worksite prisoner's dilemma game

		Employer	
		<i>c</i>	<i>d</i>
Work supplier	<i>c</i>	(<i>C</i> , <i>C</i>)	(<i>L</i> , <i>H</i>)
	<i>d</i>	(<i>H</i> , <i>L</i>)	(<i>D</i> , <i>D</i>)

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). The 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^0 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^0 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 personality of each agent, as expressed in his worksite interactions, is governed by a *worksite strategy* that is maintained throughout the course of each trade cycle loop. These worksite strategies are represented as finite-memory pure strategies for playing a prisoner's dilemma game with an arbitrary partner an indefinite number of times. At the commencement of each trade cycle loop, agents have no information about the worksite strategies of other agents; each agent can only learn about these strategies by engaging other agents in repeated worksite interactions and observing the actions and utility outcomes that ensue. Each agent keeps separate track of his interaction history with each potential worksite partner, and each agent's choice of an action in a current worksite interaction with another agent is determined on the basis of his own past interactions with this other agent plus his initial expected utility assessment of the agent.⁹ This means, in particular,

⁹ The implication 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.

that an agent may end up revealing different aspects of his personality to different worksite partners due to differences in their interaction histories. For example, a work supplier may develop a mutually cooperative relationship with one employer while at the same time he is shirking on the job with a second employer.

At the end of each trade cycle loop, the *utility (fitness)* of each work supplier and employer is measured by normalized total net payoff, that is, by total net payoff divided by the fixed number of trade cycles constituting each trade cycle loop. For employers, total net payoff is measured by total net worksite payoffs; for work suppliers, total net payoff is measured by total net worksite payoffs plus the (negative) sum of any incurred refusal payoffs.

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.¹⁰ This evolution is meant to reflect the 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 welfare outcomes will be used in subsequent work as a benchmark against which to assess the performance of more realistically modelled social learning mechanisms.

¹⁰ 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) *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 introduced into the strategy collection 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 strategy collection are retained for the new strategy collection and the least fit worksite 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 general discussion of genetic algorithm design and use in economics.

3. Descriptive statistics

Path dependency arises in the labor market framework because tie-breaking rules and the inductive learning of new worksite strategies in the evolution step both rely in part on randomization. When each particular experimental treatment is repeated for a range of different pseudo-random number seed values, a distribution of network, behavioral, welfare, and market power outcomes is generated. Consequently, the mapping between treatment factors and outcomes is one-to-many and 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, welfare outcomes, and market power outcomes that these networks support. Networks depict who is working for whom, and with what regularity. Worksite behavior refers to the specific actions undertaken by work suppliers and employers in their worksite interactions. Welfare refers to the utility levels attained by work suppliers and employers as a result of job search and worksite interactions. Finally, market power refers to the ability of work suppliers or employers to secure for themselves an average utility level that exceeds the level they would attain under competitive market conditions.

3.1. Classification of networks by distance

First introduced is a distance measure on networks that permits the classification of these network into alternative types. This distance measure calculates the extent to which an observed pattern of relationships among work suppliers and employers deviates from an idealized pattern compatible with a competitive (full employment) market outcome.

All labor market experiments reported in this study are implemented using the Trade Network Game (TNG) source code developed by McFadzean and Tesfatsion (1999). Let s denote a seed value for the pseudo-random number generator incorporated in this TNG source code, and let E denote a *potential economy*, i.e., an economy characterized structurally by the TNG source code together with specific values for all TNG source code parameters¹¹ 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,¹² hence the actions undertaken by any agent v in repeated worksite interactions with

¹¹A complete annotated listing of these user-supplied TNG source code parameters is given in Section 4.

¹²A *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 work supplier 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.

another agent z must eventually cycle.¹³ 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 nonrepeated 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 an employer e are said to exhibit a *persistent relationship* during a given trade cycle loop of a sample economy (s, E) if the following two conditions hold: (a) their worksite histories with each other during the course of this loop 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 this loop either by choice (a permanent switch away to a strictly preferred partner) or by refusal (one agent becomes intolerable to the other because of too many defections).

A work supplier or employer who fails to form any persistent relationships during a given trade cycle loop is called *persistently inactive*. A work supplier in a persistent relationship with an employer in a given trade cycle loop is said to be *latched* to the employer if he works for the employer continuously (in each successive trade cycle) rather than recurrently (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 (contracts) 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 agents 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) . By construction, any persistently inactive agent constitutes an isolated vertex of $K^p(s, E)$.

For each potential economy E , let $V^0(E)$ denote the following partially specified relationship pattern: *Each work supplier directs work offers to employers without latching, and no work supplier or employer is persistently inactive*. Let

¹³ For the particular parameter values used in this study, the maximum possible cycle length is 1024. In the experiments reported below in Section 5, however, the observed cycle lengths tend to lie between 1 and 7, and the maximum observed cycle length is only 23.

$K^0(E)$ denote the collection of all networks whose edges conform to the pattern $V^0(E)$. The network class $K^0(E)$ is capable of supporting competitive market conditions for E , in a sense that will be clarified in Section 3.3. Consequently, $V^0(E)$ and $K^0(E)$ will henceforth be referred to as the *competitive relationship pattern* and the *competitive network class* for E , respectively.

The distance $D^0(s, E)$ between the persistent network $K^p(s, E)$ and the competitive network class $K^0(E)$ for a sample economy (s, E) is defined to be the number of vertices (agents) in $K^p(s, E)$ whose edges (persistent relationships) fail to conform to the competitive relationship pattern $V^0(E)$. By construction, then, the distance $D^0(s, E)$ indicates the extent to which $K^p(s, E)$ deviates from the ‘null hypothesis’ $K^0(E)$ of a competitive market network. More generally, as will be seen in Section 5, the distance measure D^0 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.

3.2. 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 called *aggressive* if he engages in at least one defection against another agent who has not previously defected against him. The 1×2 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 the work suppliers and employers behave opportunistically in worksite interactions with partners who are strangers¹⁴ or who so far have been consistently cooperative.

As noted in Section 3.1, a work supplier or employer in the final generation of (s, E) is called *persistently inactive* if he fails to establish any persistent relationships and hence constitutes an isolated vertex of the persistent network $K^p(s, E)$. The 1×2 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 percentage for work suppliers constitutes their persistent unemployment rate, and 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 . The 1×2 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 *repeat defector profile* for (s, E) . The repeat defector profile measures the extent

¹⁴ The importance of stance toward strangers and first impressions for determining subsequent outcomes in socioeconomic contexts is stressed by Rabin and Schrag (1999).

to which the work suppliers and employers engage in recurrent or continuous defections.

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

3.3. Market power measurement

The current study adopts the standard industrial organization approach to the measurement of market power: namely, market power is measured by the degree to which the actual welfare levels attained by work suppliers and employers compare against an idealized competitive yardstick. This competitive yardstick requires absence of strategic behavior, symmetric treatment of equals, and full employment.

Specifically, given any potential economy E , *competitive market conditions* are said to hold for E if the following four conditions are satisfied: (i) Work suppliers and employers behave cooperatively in all of their worksite interactions; (ii) each work supplier has the same number of accepted work offers as any other work supplier over the course of each complete trade cycle loop; (iii) each employer has the same number of vacant job openings as any other employer over the course of each complete trade cycle loop; and (iv) full employment is obtained in each trade cycle, in the sense that the ratio of accepted work offers to total potential work offers is as high as possible given the particular job capacity level specified for E .

These competitive market conditions are idealized conditions that may or may not be attained in any actual sample economy (s, E) . In the absence of job search costs, however, any network supporting these four competitive market conditions would have to be in the competitive network class $K^0(E)$. This follows because condition (i) would then imply that work suppliers indifferently direct their work offers among employers, resulting in recurrent rather than latched relationships, and the symmetric treatment conditions (ii) and (iii) and

the full employment condition (iv) imply that no work supplier or employer is persistently inactive.

As will be seen in Section 5, the utility profile $U^*(E) = (U_w^*(E), U_e^*(E))$ that work suppliers and employers would obtain in any trade cycle loop under competitive market conditions is straightforward to calculate for each tested potential economy E . For any actual sample economy (s, E) corresponding to E , the *market power* of work suppliers and employers in the final generation of (s, E) is measured by the extent to which their realized utility profile $U(s, E)$ deviates from $U^*(E)$. Specifically, the market power of work suppliers is measured in percentage terms by

$$MP_w(s, E) = \frac{U_w(s, E) - U_w^*(E)}{U_w^*(E)} \times 100 \quad (1)$$

and the market power of employers is measured in percentage terms by

$$MP_e(s, E) = \frac{U_e(s, E) - U_e^*(E)}{U_e^*(E)} \times 100. \quad (2)$$

The vector $MP(s, E) = (MP_w(s, E), MP_e(s, E))$ is hereafter referred to as the *market power profile* for (s, E) . As will be clarified in Section 5, the market power profile must be interpreted with care; the competitive yardstick used in its construction ignores the fact that organizational costs typically have to be incurred to reach and sustain any market state, whether competitive or not.

4. Experimental design

The experiments reported in Section 5 are for two-sided labor markets comprising NW work suppliers and NE employers. Each work supplier has the same work offer quota, wq , and each employer has the same acceptance quota, eq . The experimental design focuses on the independent variation of two factors: job concentration as measured by $JCON = NW/NE$; and job capacity as measured by $JCAP = (NE \cdot eq)/(NW \cdot wq)$. As shown in Table 4, three settings are tested for each factor — low, balanced, and high — resulting in a 3×3 design matrix comprising a total of nine tested potential economies E .

All remaining parameters are maintained at fixed values throughout all experiments. Table 5 lists these fixed parameter values along with specific NW , NE , wq , and eq values yielding a $JCON$ value equal to 2 and a $JCAP$ value equal to 1.¹⁵

¹⁵ In a previous study focusing on iterated prisoner's dilemma with choice and refusal of partners by Ashlock et al. (1996), the prisoner's dilemma payoffs were set at the commonly used levels $L = 0$, $D = 1$, $C = 3$, and $H = 5$ and the minimum tolerance level and inactivity payoff were both set at 1.6. For comparison purposes, the present study uses these same specifications except that they are normalized by subtracting 1.6 from each value so that the minimum tolerance level and inactivity payoff are now both set at zero.

Table 4

Two-factor experimental design: job concentration ($JCON = NW/NE$) versus job capacity ($JCAP = [NE \cdot eq]/[NW \cdot wq]$)

	Tight job capacity $JCAP = 1/2$		Balanced job capacity $JCAP = 1$		Excess job capacity $JCAP = 2$	
High job concentration $JCON = 2$	$NW = 12$ $wq = 1$	$NE = 6$ $eq = 1$	$NW = 12$ $wq = 1$	$NE = 6$ $eq = 2$	$NW = 12$ $wq = 1$	$NE = 6$ $eq = 4$
Balanced job concentration $JCON = 1$	$NW = 12$ $wq = 2$	$NE = 12$ $eq = 1$	$NW = 12$ $wq = 1$	$NE = 12$ $eq = 1$	$NW = 12$ $wq = 1$	$NE = 12$ $eq = 2$
Low job concentration $JCON = 1/2$	$NW = 6$ $wq = 4$	$NE = 12$ $eq = 1$	$NW = 6$ $wq = 2$	$NE = 12$ $eq = 1$	$NW = 6$ $wq = 1$	$NE = 12$ $eq = 1$

For each tested potential economy E , 20 sample economies (s, E) were experimentally generated using 20 arbitrarily selected seed values s for the pseudo-random number generator included in the TNG source code.¹⁶ For each run s , the persistent network $K^p(s, E)$ was determined and graphically depicted, and the four behavioral profiles (aggressive, p-inactive, repeat defector, p-nice), the utility profile, and the market power profile were determined and recorded.

The distance $D^0(s, E)$ of $K^p(s, E)$ from the competitive network class $K^0(E)$ was then determined for each run s , and a histogram for the distance values $D^0(s, e)$ was constructed giving the percentage of runs s corresponding to each possible distance value.

5. Experimental findings

As detailed in Section 1, the primary objective of this study is to test hypotheses H1–H4 regarding the ability of work suppliers and employers to exercise market power under various job capacity and job concentration conditions. Hypotheses H1 and H3 roughly state that work suppliers do better (and

¹⁶ These 20 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 14 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.

Table 5

Parameter values for a potential economy E with $JCON = 2$ and $JCAP = 1$

// PARAMETER VALUES HELD FIXED ACROSS EXPERIMENTS	
GMax = 50	// Total number of generations.
IMax = 150	// Number of trade cycles per trade cycle loop.
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^e .
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	
WorkSuppliers = 12	// Number of work suppliers NW .
Employers = 6	// Number of employers NE .
WorkQuota = 1	// Work offer quota wq .
EmployerQuota = 2	// Employer acceptance quota eq .

relatively better than employers) in terms of exercising market power as job capacity is increased, and hypotheses H2 and H4 roughly state that employers do better (and relatively better than work suppliers) in terms of exercising market power as job concentration is increased. Aggregate market power findings are reported first, followed by a report on market power findings at a more disaggregated level.

5.1. Aggregate market power findings

Recall the market power measures $MP_w(s, E)$ and $MP_e(s, E)$ for work suppliers and employers defined by relations (1) and (2) in Section 3.3 for any sample economy (s, E) . By construction, these measures are positively valued if and only if the actual utility levels $U_w(s, E)$ and $U_e(s, E)$ attained by work suppliers and employers — which include an accounting for organizational costs — exceed idealized competitive utility levels $U_w^*(E)$ and $U_e^*(E)$ for which no such accounting is taken.

More precisely, as will be seen below, the competitive utility levels ignore three types of organizational costs that can significantly affect the actual utility levels attained by work suppliers and employers in any sample economy: (i) Search and inactivity *sunk* costs incurred during the process of establishing a persistent network of relationships; (ii) search and inactivity *variable* costs

Table 6

Aggregate market power outcomes. Means and standard deviations for the market power measures $MP_w(s, E)$ and $MP_e(s, E)$ for work suppliers w and employers e across all runs s for each of the nine tested potential economies E

	Tight job capacity $JCAP = 1/2$		Balanced job capacity $JCAP = 1$		Excess job capacity $JCAP = 2$	
	w	e	w	e	w	e
High job concentration $JCON = 2$	– 54.7 (34.2)	– 23.0 (45.9)	– 18.7 (21.1)	– 10.3 (19.1)	+ 4.3 (21.4)	– 44.7 (31.2)
Balanced job concentration $JCON = 1$	– 55.6 (32.3)	– 25.7 (42.2)	– 17.6 (19.8)	– 11.5 (17.3)	+ 17.4 (20.9)	– 39.9 (27.7)
Low job concentration $JCON = 1/2$	– 62.8 (26.4)	– 30.1 (45.1)	– 20.5 (21.9)	– 10.4 (16.4)	– 0.3 (18.8)	– 28.5 (29.1)

incurred in the process of maintaining a persistent network of relationships; and (iii) utility losses (negative payoffs) incurred when worksite partners defect. Even if the competitive utility levels are ultimately attained in a steady-state sense, the sunk costs associated with attaining this competitive state may result in actual utility levels that are below the competitive utility levels.

Given these considerations, it is not surprising that the experimentally determined values for the market power measures (1) and (2), aggregated across the 20 sample economies for each of the nine tested potential economies, were found to be negatively valued in all but two cases. These aggregate market power outcomes, reported in Table 6, suggest that it is not possible in general to infer the relative ability of work suppliers and employers to exercise market power under alternative structural conditions simply by comparing the signs of their attained market power measures. Rather, for such a determination, attention must also be paid to the relative magnitudes of these attained market power measures.

The key implication of the aggregate market power outcomes reported in Table 6 is that job capacity is the dominant factor determining relative market power. Specifically, these outcomes provide strong support for the job capacity hypotheses H1 and H3(a) and weak support for the job capacity hypothesis H3(b). Regarding the latter, H3(b) fails to hold as job capacity is increased from tight to balanced but does hold as job capacity is further increased from balanced to excess. On the other hand, these outcomes do not support the job concentration hypotheses H2 and H4.

The following sections seek a better understanding of these market power findings by examining the experimental data at a more disaggregated level.

5.2. Disaggregated market power findings for high job concentration

Detailed experimental findings for high job concentration ($JCON = 2$) and job capacity $JCAP$ varying from tight to excess are reported in Table 7. For reasons clarified below, these findings generally support the job capacity hypotheses H1 and H3, and this support is particularly strong when attention is restricted to dominant distance clusters. On the other hand, these findings strongly contradict the job concentration hypothesis H2(a) under conditions of excess job capacity.

Consider, first, the potential economy E characterized by high job concentration ($JCON = 2$) and tight job capacity ($JCAP = 1/2$). As shown in Table 4, for this E there are 12 work suppliers who each control one work offer and six employers who each control one job opening, hence job openings are concentrated in the hands of relatively few employers. Moreover, there are two potential work offers for each potential job opening, hence work suppliers face a high structural risk of unemployment. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (0.70, 1.40)$.¹⁷

According to hypotheses H1(a) and H2(a), work suppliers should be disadvantaged relative to employers in this potential economy E regarding their ability to exercise market power. Experimental findings for the 20 sample economies (s, E) run for this E are reported in Table 7(a). As seen, these sample economies can be partitioned into three distinct distance (D^0) clusters supporting three distinct types of behavioral and utility outcomes.¹⁸ In each distance

¹⁷This determination of $U^*(E)$ results from the following considerations. Given the structural conditions defining the potential economy E with $JCON = 2$ and $JCAP = 1/2$ (see Table 4), together with the competitive market conditions set out in Section 3.3, each of the 12 work suppliers engages in one worksite interaction in 75 of the 150 trade cycles constituting a trade cycle loop, each of the six employers engages in one worksite interaction in each of the 150 trade cycles constituting a trade cycle loop, and each worksite interaction is characterized by mutually cooperative behavior with mutual cooperation payoffs of (1.40, 1.40) for the work supplier and employer. Consequently, recalling that the utility level attained by an agent over the course of any trade cycle loop is defined to be his average (per trade cycle) payoff, the utility level attained by each work supplier over the course of a trade cycle loop is equal to $0.70 = 1.40/2$, and the utility level attained by each employer over the course of a trade cycle loop is equal to 1.40.

¹⁸For each distance cluster, the mean and standard deviation are reported for each component of the three behavioral profiles (aggressive, p-inactive, p-nice), the utility profile, and the market power profile across all of the sample runs s that lie in this distance cluster. By construction (see Section 3.2), the components of the p-inactive, repeat defector, and p-nice profiles corresponding to each agent type must sum to 100% since each agent is a member of one and only one of these three behavioral categories. Since the repeat defector profiles can thus be inferred from the reported p-inactive and p-nice profiles, they are omitted to conserve space. A similar remark holds for all remaining distance cluster results reported in Table 7 and in subsequent tables.

Table 7
Experimental findings for high job concentration ($JCON = 2$) and varying job capacity levels

D^0	% of runs	Aggressive		p-Inactive		p-Nice		Utility		MPOWER		
		w (%)	e (%)	w (%)	e (%)	w (%)	e (%)	w	e	w (%)	e (%)	
(a) <i>Tight job capacity</i> ($JCAP = \frac{1}{2}$)												
0-7	25	0	7	40	0	55	97	0.48	1.35	-32	-2	
		(0)	(8)	(3)	(0)	(9)	(7)	(0.04)	(0.07)	(6)	(4)	
10-13	45	6	56	51	4	42	46	0.41	1.55	-41	+10	
		(7)	(50)	(5)	(7)	(14)	(46)	(0.19)	(0.27)	(27)	(19)	
15-18	30	88	92	90	81	1	3	0.04	0.14	-94	-89	
		(28)	(19)	(11)	(22)	(3)	(6)	(0.15)	(0.16)	(22)	(10)	
(b) <i>Balanced job capacity</i> ($JCAP = 1$)												
0-9	15	6	67	17	0	28	22	0.69	2.06	-51	-27	
		(8)	(47)	(24)	(0)	(39)	(31)	(0.51)	(0.64)	(36)	(23)	
12	80	25	17	1	0	90	86	1.29	2.72	-8	-3	
		(42)	(33)	(2)	(0)	(25)	(25)	(0.19)	(0.14)	(13)	(5)	
15	5	100	17	75	50	0	0	0.13	0.79	-91	-72	
		NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	
(c) <i>Excess job capacity</i> ($JCAP = 2$)												
0-5	25	25	0	2	0	75	70	1.09	2.39	-8	-15	
		(38)	(0)	(3)	(0)	(38)	(37)	(0.43)	(0.61)	(11)	(22)	
12-14	75	93	52	5	10	19	36	1.52	1.27	+9	-55	
		(25)	(46)	(6)	(10)	(35)	(40)	(0.31)	(0.76)	(22)	(27)	

cluster, however, the mean market power level attained by work suppliers is lower than the mean market power level attained by employers as predicted by hypotheses H1(a) and H2(a).

The distance cluster 0–7 in Table 7(a) includes 25% of the 20 sample economies run for this potential economy E . By definition of the distance measure D^0 , these are the sample economies whose persistent networks are closest to the competitive network class $K^0(E)$. The persistent networks that evolve for this distance cluster are largely recurrent in form, and the worksite behaviors supported by these persistent networks are largely p-nice (persistently cooperative).

The frequent p-nice behavior exhibited by employed work suppliers in distance cluster 0–7 is consistent with efficiency-wage theory, which predicts that high job search costs should discourage shirking (worksite defections). Given tight job capacity ($JCAP = 1/2$), at most half of all potential work offers can be accepted at any given time. Any work supplier fired for shirking can thus expect to incur high job search costs (negative refusal payoffs) as he attempts to secure new employment.

Indeed, as seen in Table 7(a), 40% of the work suppliers in distance cluster 0–7 are ultimately p-inactive (unemployed). These work suppliers become so discouraged by the high job search costs they incur in their attempts to secure employment that the expected utility they assign to each employer eventually becomes negative, discouraging any further work offers. In addition to the opportunity costs incurred by unemployed work suppliers, the recurrent form of the persistent networks in this distance cluster implies that employed work suppliers incur variable job search costs in the process of maintaining these networks. The mean utility level 0.48 attained by work suppliers in this ‘competitive’ distance cluster is therefore markedly lower than the competitive utility level 0.70, resulting in a low mean market power level of -32% for work suppliers.

In contrast, the mean utility level 1.35 attained by employers in distance cluster 0–7 is close to their competitive utility level 1.40, hence their mean market power is close to 0. The question then arises whether employers could do even better if they took advantage of the weak structural position of work suppliers by more frequent defections.

The next distance cluster 10–13 reported in Table 7(a), comprising 45% of the sample economies for this E , reveals that employers can indeed do better on average by engaging in more aggressive behavior in their initial worksite relationships and more repeat defection behavior in their persistent worksite relationships. As depicted in Fig. 1(b), the typical persistent network that evolves in this dominant distance cluster consists of six disjoint pairs of work suppliers and employers in persistent latched relationships together with six p-inactive work suppliers. The mean utility and market power levels attained by work suppliers in this dominant distance cluster are smaller than the levels they

attained in distance cluster 0–7; the costs of increased employer defection rates and an increased unemployment rate outweigh on average the benefits of reduced variable search costs due to the latched nature of the persistent networks. On the other hand, the mean utility and market power levels attained by employers are markedly higher.

Nevertheless, the final distance cluster 15–18 reported in Table 7(a), comprising 30% of the sample economies for this E , cautions that too determined an effort by employers to exercise market power through aggression and repeat defections can be self-defeating. Work suppliers interacting with strongly aggressive and predacious employers tend to evolve worksite behavioral rules that are aggressive and predacious in turn, whatever their structural situation. Indeed, in distance cluster 15–18, worksite interactions come to be dominated by mutual defection behavior, which leads to quits and firings. Note that 90% of the work suppliers and 81% of the employers in this distance cluster ultimately become p-inactive. This widespread coordination failure dramatically decreases the mean utility and market power levels attained by work suppliers and employers alike.

In short, the results reported in Table 7(a) provide strong support for hypotheses H1(a) and H2(a); employers are advantaged relative to work suppliers with regard to their ability to exercise market power in conditions of high job concentration and tight job capacity. Yet a closer examination of the data reveals that the mobility and behavioral flexibility (provocability) of the work suppliers provides them with some degree of protection against exploitive behavior by employers.

Consider, next, a potential economy E for which job concentration remains high ($JCON = 2$) but job capacity increases from tight to balanced ($JCAP = 1$), so that the total number of potential job openings is just equal to the total number of potential work offers. Specifically, referring to Table 4, 12 work suppliers each control one work offer and six employers each control two job openings. Thus, work suppliers no longer face a structural risk of unemployment due to capacity constraints. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (1.40, 2.80)$.

Hypotheses H1(a) and H2(a) predict that work suppliers will be disadvantaged relative to employers in this potential economy E . Work suppliers must incur job search costs despite balanced job capacity and, in addition, jobs are concentrated in the hands of relatively few employers. Indeed, in each of the three distance clusters reported in Table 7(b) for this E , the mean market power level attained by work suppliers is lower than the mean market power level attained by employers.

In particular, the Table 7(b) findings reveal that 80% of the sample economies observed for this E lie in a distance cluster $D^0 = 12$. As depicted in Fig. 1(c), the

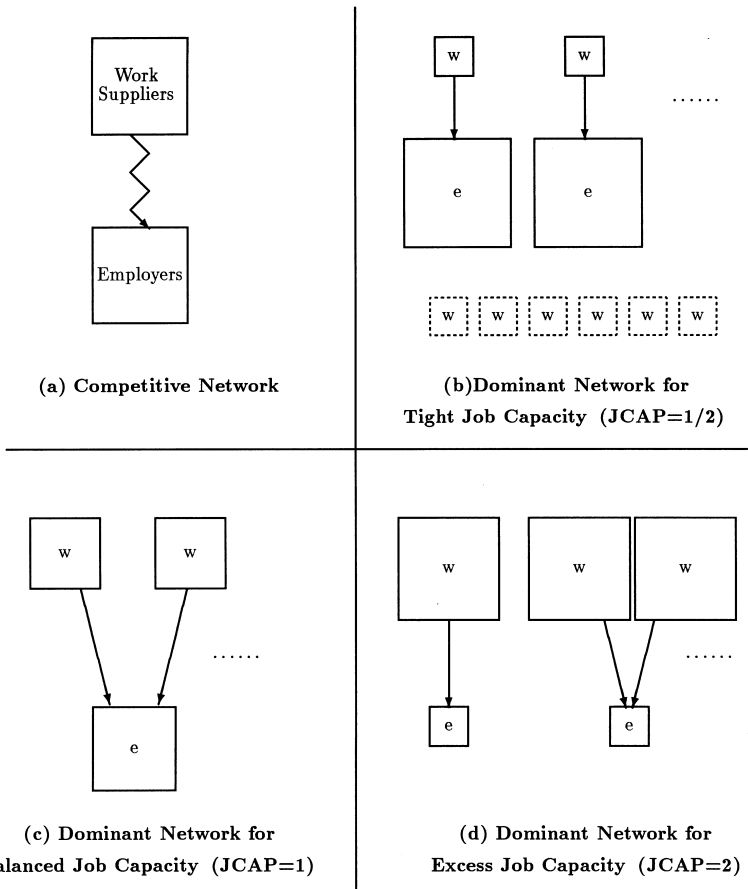


Fig. 1. Competitive and dominant persistent networks for high job concentration ($JCON = 2$) and varying job capacity levels. A relatively larger solid-line box for work suppliers w or employers e indicates a relatively higher attained mean market power level. Dash-line boxes denote p-inactive agents. Straight-line directed arrows denote latched (continuous) relationships and zig-zag directed arrows denote recurrent (random or periodic) relationships.

persistent networks that evolve for this dominant distance cluster all take the form of disjoint triads comprising two work suppliers latched to a single employer, with almost all agents expressing p-nice worksite behaviors. The latched nature of these persistent networks means that work suppliers avoid variable job search costs, although they still incur sunk job search costs in the process of forming these networks.

Finally, consider what happens when job concentration remains high ($JCON = 2$) but job capacity further increases from balanced to excess

($JCAP = 2$). Referring to Table 4, for this potential economy E 12 work suppliers each control one work offer and six employers each control four job openings. Consequently, there are two job openings for every potential work offer, implying that employers face a high structural risk of vacancy. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (1.40, 2.80)$.

As seen in Table 7(c), 25% of the sample economies for this E lie in a distance cluster 0–5 close to the competitive network class ($D^0 = 0$). The persistent networks that evolve for these sample economies are largely recurrent in form, and the large majority of work suppliers and employers are p-nice. Nevertheless, 25% of the work suppliers are aggressive in their initial interactions with employers, and this gives them an edge over employers in terms of their attained market power. Thus, hypothesis H1(b) is supported and hypothesis H2(a) is contradicted in this distance cluster.

Table 7(c) also shows that the remaining 75% of the sample economies observed for this E lie in a distance cluster 12–14. As depicted in Fig. 1(d), the typical persistent network that evolves in this dominant distance cluster consists of disjoint components comprising from one to four work suppliers latched to a single employer who engage in frequent defections against this employer. Note that 93% of the work suppliers in this dominant distance cluster are aggressive and 76% engage in repeat defections. While some employers respond by evolving aggressive and/or repeat defection behavior themselves, work suppliers still do markedly better than employers on average in terms of exercising market power, attaining a mean market power level of +9% in comparison with only –55% for employers. Consequently, in this dominant distance cluster, hypothesis H1(b) is strongly supported and hypothesis H2(a) is strongly contradicted.

Recall from Section 5.1 that the aggregate market power outcomes reported in Table 6 indicate a violation of the job capacity sensitivity hypothesis H3 under conditions of high job concentration. Specifically, in contradiction to the prediction of H3(b), employers on average do not experience a decline in their ability to exercise market power as job capacity increases from tight to balanced. The disaggregated findings reported in Table 7 show that this is a misleading conclusion. Omitting the relatively small number of sample economies in Table 7(a) for which widespread coordination failure occurs and both agent types do extremely poorly — i.e., distance cluster 15–18 in Table 7(a) — hypothesis H3 is actually strongly supported as job capacity increases from tight to balanced. This pooling problem is indicated in Table 6 by the high standard deviations for the mean market power levels attained by work suppliers and employers under conditions of high job concentration and tight job capacity.

5.3. Disaggregated market power findings for balanced job concentration

Detailed experimental findings for balanced job concentration ($JCON = 1$) and job capacity $JCAP$ varying from tight to excess are reported in Table 8. For reasons explained below, these findings generally support the job capacity hypotheses H1 and H3, and this support is particularly strong when attention is restricted to dominant distance clusters. In contrast, these findings strongly contradict the relative job concentration hypothesis H2(a) under conditions of excess job capacity and the job concentration sensitivity hypothesis H4(a) under conditions of balanced job capacity.

Consider the potential economy E characterized by balanced job concentration ($JCON = 1$) and tight job capacity ($JCAP = 1/2$). Referring to Table 4, for this E there are 12 work suppliers who each control two work offers and 12 employers who each control one job opening. Thus, although job concentration is balanced, two potential work offers exist for each potential job opening, implying that work suppliers face a high structural risk of unemployment. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (1.40, 1.40)$.

As indicated in Table 8(a), the sample economies observed for this E lie in three distinct distance clusters. In each cluster, employers attain a markedly higher mean market power level than work suppliers. Consequently, hypotheses H1(a) and H2(a) are strongly supported.

In particular, Table 8(a) shows that 55% of the sample economies observed for this E lie in a distance cluster 1–7. As depicted in Fig. 2(b), the typical persistent network that evolves in this dominant distance cluster is a recurrent network that supports p-nice worksite behaviors by both work suppliers and employers. Nevertheless, 23% of the work suppliers in this dominant distance cluster become p-inactive (unemployed) primarily due to bad luck. More precisely, these work suppliers incur high job search costs by chance, as employers are forced to refuse work offers from cooperative work suppliers because of job capacity limitations. These high job search costs ultimately lead the work suppliers to associate a negative expected utility with each potential employer, which discourages them from making any further work offers.

The question again arises whether employers could take better advantage of the relatively weak structural situation of work suppliers facing tight job capacity by engaging in more aggressive and predacious worksite behavior. The answer, again, is a qualified yes.

The next distance cluster 12–14 reported in Table 8(a) comprises 20% of the remaining sample economies for this E . Employers in this secondary distance cluster are much more aggressive in their initial worksite relationships and less cooperative in their persistent worksite relationships, as indicated by their 75% aggressive percentage and their 50% p-nice percentage. Work suppliers retaliate

Table 8
Experimental findings for balanced job concentration ($JCON = 1$) and varying job capacity levels

D^0	% of runs	Aggressive		p-Inactive		p-Nice		Utility		MPOWER		
		w (%)	e (%)	w (%)	e (%)	w (%)	e (%)	w	e	w (%)	e (%)	
(a) Tight job capacity ($JCAP = \frac{1}{2}$)												
1–7	55	2	5	23	2	74	95	0.87	1.38	–38	–1	
		(4)	(1)	(9)	(5)	(11)	(12)	(0.13)	(0.04)	(9)	(3)	
12–14	20	10	7.5	38	4	48	50	0.88	1.46	–37	+4	
		(18)	(43)	(9)	(7)	(25)	(42)	(0.21)	(0.22)	(15)	(15)	
24	25	100	100	100	100	0	0	–0.12	–0.02	–109	–101	
		(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	
(b) Balanced job capacity ($JCAP = 1$)												
1	5	0	0	8	0	100	100	0.45	1.37	–68	–2	
		NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	
12	70	12	31	0	0	85	85	1.23	1.33	–12	–5	
		(26)	(44)	(0)	(0)	(26)	(26)	(0.19)	(0.14)	(13)	(10)	
13–15	25	63	33	13	15	48	67	1.08	0.97	–23	–31	
		(45)	(40)	(4)	(3)	(40)	(35)	(0.30)	(0.29)	(21)	(21)	
(c) Excess job capacity ($JCAP = 2$)												
0	5	17	0	0	0	83	100	1.42	1.30	1	–7	
		NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	
4–6	25	2	0	0	0	98	98	1.37	1.38	–2	–1	
		(3)	(0)	(0)	(0)	(0)	(0)	(0.00)	(0.04)	(0)	(2)	
14–17	70	100	29	2	29	10	60	1.74	0.61	+26	–56	
		(0)	(45)	(3)	(9)	(26)	(32)	(0.32)	(0.27)	(20)	(19)	

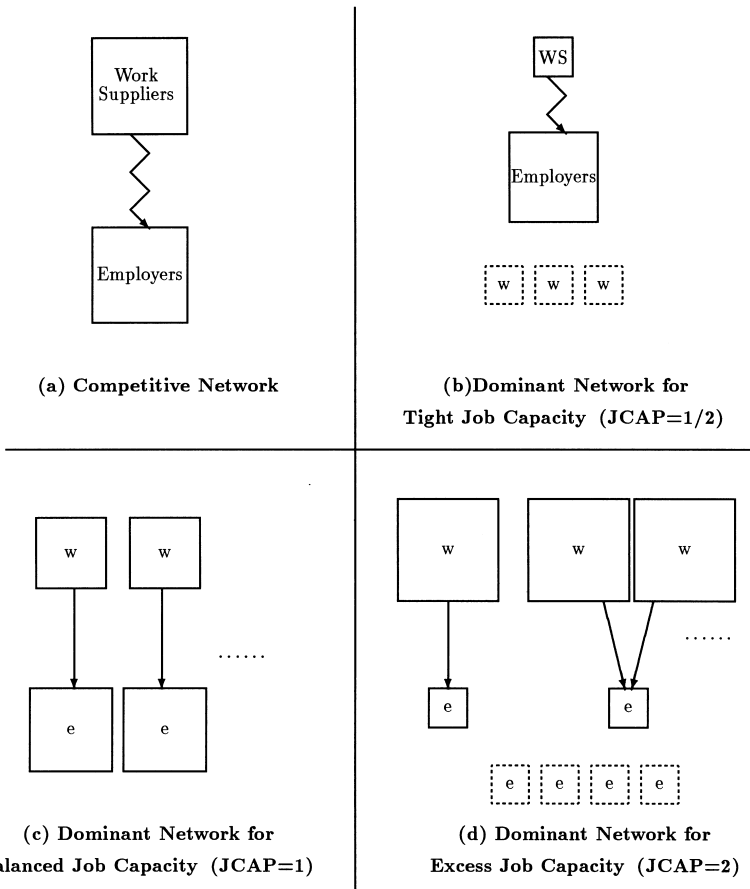


Fig. 2. Competitive and dominant persistent networks for balanced job concentration ($JCON = 1$) and varying job capacity levels. A relatively larger solid-line box for work suppliers w or employers e indicates a relatively higher attained mean market power level. Dash-line boxes denote p-inactive agents. Straight-line directed arrows denote latched (continuous) relationships and zig-zag directed arrows denote recurrent (random or periodic) relationships.

by being less cooperative in turn, which results in increased firings and an increased unemployment (p-inactivity) rate of 38%. The persistent networks that evolve are now latched rather than recurrent in form — each work supplier is either p-inactive or latched to one or two employers — which reduces variable job search costs for employed work suppliers. For work suppliers, these changes balance out, leaving them with mean utility and market power levels that differ little from the levels they attained in the dominant distance cluster 1–7. In

contrast, employers manage to secure a small increase in their mean utility and market power levels.

On the other hand, the final distance cluster $D^0 = 24$ in Table 8(a), comprising the remaining 25% of the sample economies for this E , illustrates the risk faced by employers when they attempt to exploit their structural advantage too vigorously. In this distance cluster, work suppliers retaliate strongly against initial and repeated defections by employers, resulting in the evolution of such aggressive and predacious worksite strategies on the part of all agents that no persistent relationships form. This complete coordination failure is evidenced by the 100% p-inactivity (unemployment and vacancy) rates reported for work suppliers and employers.

Consider, next, a potential economy E for which job concentration remains balanced ($JCON = 1$) but job capacity increases from tight to balanced ($JCAP = 1$), implying that neither agent type is disadvantaged by job concentration or job capacity conditions per se. Referring to Table 4, for this E there are 12 work suppliers who each control one work offer and 12 employers who each control one job opening. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (1.40, 1.40)$.

Table 8(b) shows that 70% of the sample economies observed for this E lie in a distance cluster $D^0 = 12$. As depicted in Fig. 2(c), the persistent networks that evolve in this dominant distance cluster all consist of disjoint latched pairings of work suppliers and employers exhibiting largely p-nice worksite behavior. In addition, a single sample economy was observed at distance $D^0 = 1$ with a completely recurrent persistent network and 100% p-nice behavior. Work suppliers do relatively poorly in this outlier sample economy because of high job search costs. The remaining 25% of the sample economies observed for this E lie in a distance cluster 13–15 with persistent networks consisting of mixtures of p-inactive agents and latched pairings of work suppliers and employers predominantly exhibiting repeat defection behavior.

Hypotheses H1(a) and H2(a) predict for this E that work suppliers should be disadvantaged relative to employers in terms of their ability to exercise market power; for, despite balanced job concentration and job capacity conditions, work suppliers bear the brunt of job search costs in the form of refusal payoffs. The findings in Table 8(b) show that this prediction is accurate for the dominant distance cluster $D^0 = 12$ containing 70% of the sample economies, and also for the single sample economy in distance cluster $D^0 = 1$. However, in the remaining distance cluster 13–15 containing 25% of the sample economies, work suppliers attain a slightly higher mean market power measure than employers.

Finally, consider a potential economy E for which job concentration remains balanced ($JCON = 1$) but job capacity increases from balanced to excess ($JCAP = 2$). As indicated in Table 4, 12 work suppliers each control one work

offer and 12 employers each control two work offers, implying that employers face a high structural risk of vacancy. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (1.40, 1.40)$.

Hypothesis H1(b) predicts that work suppliers should do better than employers in this E , whereas hypothesis H2(a) predicts that work suppliers should do worse. The findings reported in Table 8(c) generally support H1(b), and this support is strong for the sample economies in the dominant distance cluster. Consequently, H2(a) is strongly contradicted.

Specifically, Table 8(c) shows that 70% of the sample economies observed for this E lie in a distance cluster 14–17. As depicted in Fig. 2(d), the typical persistent network that evolves in this dominant distance cluster is a mixture of p-inactive employers and disjoint latched groupings in which one or two work suppliers are latched to a single employer. Work suppliers are predominantly repeat defectors whereas most employers are p-nice. The mean market power level attained by work suppliers is +26% whereas the mean power level attained by employers is only –56%.

The remaining sample economies for this E lie in secondary distance clusters 4–6 and $D^0 = 0$. The persistent networks that evolve in distance cluster 4–6 are mixtures of recurrent and latched relationships, and work suppliers and employers exhibit largely p-nice worksite behaviors and attain similar mean market power levels; the slightly lower level attained by work suppliers is primarily due to job search costs. For the single sample economy in the ‘competitive’ distance cluster $D^0 = 0$, the persistent network is entirely recurrent and all employers are p-nice. However, aggressive and repeat defector behavior on the part of a small number of work suppliers results in work suppliers attaining a higher mean market power level than employers, despite job search costs.

According to the job capacity sensitivity hypothesis H3, work suppliers should be increasingly better off and employers increasingly worse off as job capacity varies from tight to excess, given a fixed balanced job concentration level. Omitting the single sample economy observed to lie in distance cluster $D^0 = 1$ in Table 8(b) — for which the persistent network is recurrent, all behavior is p-nice, and workers attain a low mean market power level of –68% due to high job search costs — hypothesis H3(a) is strongly supported. Also, omitting outlier sample economies in Table 8(a) for which coordination failure is widespread and both agent types do exceedingly poorly — i.e., distance cluster $D^0 = 24$ in Table 8(a) — hypothesis H3(b) is strongly supported as well.

Recall that the aggregate market power outcomes reported in Table 6 indicate that hypothesis H3(b) is violated as one moves from tight to balanced job capacity, keeping job concentration balanced, for employers are not worse off on average. The disaggregated data outcomes in Table 8 now show this to be a misleading inference arising from inappropriate data pooling. This pooling

problem is indicated in Table 6 by the high standard deviations for the mean market power levels attained by work suppliers and employers under conditions of balanced job concentration and tight job capacity.

Finally, the job concentration sensitivity hypothesis H4 predicts that work suppliers should be better off and employers worse off operating under conditions of *balanced* job concentration, as reported in Table 8, than work suppliers and employers operating under conditions of *high* job concentration as reported in Table 7, for any given job capacity level. The findings reported in Tables 7 and 8 provide weak support for hypothesis H4 under conditions of tight or excess job capacity, and for hypothesis H4(b) under conditions of balanced job capacity, in the sense that these hypotheses hold when attention is restricted to dominant distance clusters. On the other hand, H4(a) is not even weakly supported under conditions of balanced job capacity. Specifically, restricting attention to the dominant distance clusters $D^0 = 12$ in Tables 7(b) and 8(b), work suppliers are actually worse off as job concentration is decreased from high to balanced. The latter outcome is due to the higher frequencies of repeat defection behavior on the part of both work suppliers and employers in the balanced job concentration case.

5.4. Disaggregated market power findings for low job concentration

Detailed experimental findings for low job concentration ($JCON = 1/2$) and job capacity $JCAP$ varying from tight to excess are reported in Table 9. For reasons explained below, these findings generally support the job capacity hypotheses H1 and H3, and the support for H1 and H3(a) is particularly strong when attention is restricted to dominant distance clusters. In contrast, these findings strongly contradict the relative job concentration hypothesis H2(b) under conditions of tight or balanced job capacity and the job concentration sensitivity hypothesis H4 under all job capacity conditions.

Consider, first, the potential economy E characterized by low job concentration ($JCON = 1/2$) and tight job capacity ($JCAP = 1/2$). Referring to Table 4, for this E there are 6 work suppliers who each control four work offers and 12 employers who each control one job opening. Work offers are thus concentrated in the hands of relatively few work suppliers; but there are two potential work offers for each potential job opening, implying that work suppliers face a high structural risk of unemployment. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (2.8, 1.40)$.

Table 9(a) shows that 50% of the sample economies for this potential economy E lie in distance cluster 1–4. As depicted in Fig. 3(b), the typical persistent network that evolves in this dominant distance cluster is a recurrent network that supports p-nice worksite behaviors by both work suppliers and employers.

Table 9
Experimental findings for low job concentration ($JCON = \frac{1}{2}$) and varying job capacity levels

D^0	% of runs	Aggressive		p-Inactive		p-Nice		Utility		MPOWER		
		w (%)	e (%)	w (%)	e (%)	w (%)	e (%)	w	e	w (%)	e (%)	
(a) Tight job capacity ($JCAP = \frac{1}{2}$)												
0-4	50	2	2	7	0	90	98	1.67	1.40	-40	0	
		(5)	(5)	(8)	(0)	(15)	(5)	(0.12)	(0.01)	(4)	(0)	
6-16	35	45	82	43	45	29	11	0.61	0.81	-78	-42	
		(48)	(35)	(28)	(28)	(32)	(23)	(0.54)	(0.64)	(19)	(45)	
18	15	100	100	100	100	0	0	-0.12	-0.01	-101	-103	
		(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(4)	(5)	
(b) Balanced job capacity ($JCAP = 1$)												
1	5	0	17	0	8	83	83	1.11	1.39	-64	-1	
		NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	
6	80	20	33	0	0	58	60	2.40	1.26	-14	-10	
		(39)	(45)	(0)	(0)	(46)	(47)	(0.51)	(0.24)	(18)	(17)	
7-8	15	39	36	6	11	61	28	1.73	1.21	-38	-14	
		(44)	(45)	(8)	(4)	(44)	(39)	(0.49)	(0.18)	(18)	(13)	
(c) Excess job capacity ($JCAP = 2$)												
0-4	40	2	0	0	0	100	100	1.36	0.70	-3	-1	
		(6)	(0)	(0)	(0)	(0)	(0)	(0.01)	(0.01)	(1)	(2)	
12	45	100	22	0	50	19	15	1.49	0.39	7	-45	
		(0)	(42)	(0)	(0)	(0)	(0)	(0.26)	(0.17)	(18)	(24)	
13	15	100	67	17	58	11	8	1.15	0.34	-16	-51	
		(0)	(47)	(0)	(0)	(16)	(7)	(0.43)	(0.17)	(31)	(24)	

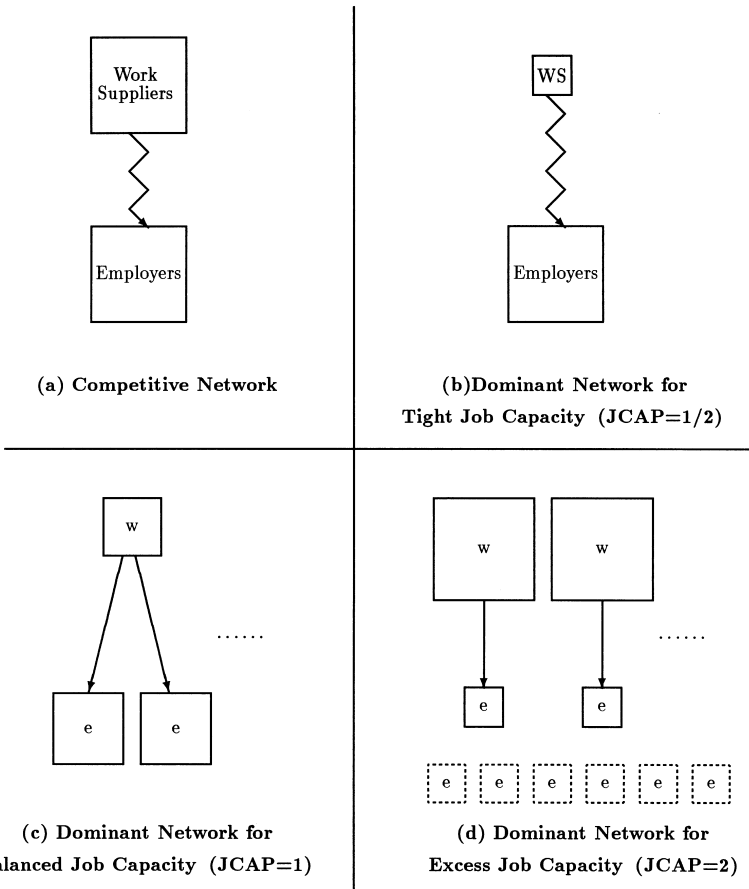


Fig. 3. Competitive and dominant persistent networks for low job concentration ($JCON = 1/2$) and varying job capacity levels. A relatively larger solid-line box for work suppliers w or employers e indicates a relatively higher attained mean market power level. Dash-line boxes denote p-inactive agents. Straight-line directed arrows denote latched (continuous) relationships and zig-zag directed arrows denote recurrent (random or periodic) relationships.

Nevertheless, 7% of the work suppliers in this dominant distance cluster become p-inactive (unemployed) due to high job search costs incurred by chance, and variable job search costs are high as well. The result is that the -40% mean market power level attained by work suppliers is much lower than the 0% mean market power level attained by employers, which supports hypothesis H1(a) and contradicts hypothesis H2(b).

The distance cluster 6–16 in Table 9(a) comprises 35% of the remaining sample economies for this E . Work suppliers and employers in this secondary

distance cluster are much more aggressive in their initial worksite relationships and far less cooperative in their persistent worksite relationships, as indicated by the aggressive and p-nice profiles reported for this distance cluster in Table 9(a). The persistent networks that evolve are now mixtures of p-inactive agents and latched groupings in which each work supplier is either p-inactive or in a latched relationship with one through four employers. The latched nature of these networks reduces variable job search costs for employed work suppliers. Nevertheless, due to increased defection frequencies leading to mutual defection worksite behavior, quits, and firings, both work suppliers and employers in this secondary distance cluster are much worse off than the work suppliers and employers in the dominant distance cluster. The -78% mean market power level attained by work suppliers is lower than the -42% mean market power level attained by employers, implying that hypothesis H1(a) is supported and hypothesis H2(b) is contradicted.

The final distance cluster $D^0 = 18$ in Table 9(a) comprises the remaining 15% of the sample economies for this E . In this distance cluster, work suppliers and employers both evolve such aggressive and predacious worksite strategies that no persistent relationships form. This complete coordination failure is evidenced by the 100% p-inactivity (unemployment and vacancy) rates reported for work suppliers and employers. Although work suppliers attain a slightly higher mean market power level than employers, both levels are so exceedingly low that this small difference is inconsequential.

Consider, next, a potential economy E for which job concentration remains low ($JCON = 1/2$) but job capacity increases from tight to balanced ($JCAP = 1$), implying that work suppliers no longer face a structural risk of unemployment. Referring to Table 4, for this E there are 6 work suppliers who each control two work offers and 12 employers who each control one job opening. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (2.80, 1.40)$.

As Table 9(b) indicates, an outlier sample economy was observed for this potential economy E in distance cluster $D^0 = 1$. The persistent network for this outlier is recurrent. Consequently, although worksite behaviors are largely p-nice, the variable job search costs incurred by work suppliers in the process of maintaining the persistent network are so high that work suppliers end up with a mean market power level of only -68% . Since employers attain a -1% mean market power level, hypothesis H1(a) is supported and hypothesis H2(b) is contradicted.

Table 9(b) shows that 80% of the sample economies observed for this E lie in a distance cluster $D^0 = 6$. As depicted in Fig. 3(c), the persistent network that consistently evolves for each sample economy in this dominant distance cluster is a collection of disjoint latched groupings each consisting of one work supplier

latched to two employers. Employers tend to be more aggressive than work suppliers in their initial worksite relationships, and a significant portion of both work suppliers and employers end up exhibiting repeat defection behavior in their persistent worksite relationships. The -14% mean market power level attained by work suppliers is lower than the -10% mean market power level attained by employers, implying that hypothesis H1(a) is supported and hypothesis H2(b) is contradicted.

Table 9(b) also shows that the remaining 15% of the sample economies observed for this E lie in a distance cluster 7–8. The persistent networks that evolve in this secondary distance cluster are mixtures of p-inactive agents and disjoint latched groupings each consisting of one work supplier latched to two employers. Both agent types are more aggressive in this secondary distance cluster than in the dominant distance cluster, and employers in particular engage in much more frequent repeat defection worksite behavior. This in turn encourages more frequent quits by work suppliers, resulting in increased p-inactivity rates for both work suppliers and employers that decreases the mean market power levels attained by both agent types in comparison with the dominant distance cluster. Since the -38% mean market power level attained by work suppliers is lower than the -14% mean market power level attained by employers, hypothesis H1(a) is supported and hypothesis H2(b) is contradicted.

Finally, consider a potential economy E for which job concentration remains low ($JCON = 1/2$) but job capacity increases from balanced to excess ($JCAP = 2$). For this E , as indicated in Table 4, 6 work suppliers each control one work offer and 12 employers each control one job opening, implying that employers face a high structural risk of vacancy. The utility levels that would be attained by work suppliers and employers over the course of a trade cycle loop for this E under competitive market conditions are given by the competitive utility profile $U^*(E) = (1.40, 0.70)$.

Table 9(c) shows that 45% of the sample economies observed for this E lie in a distance cluster $D^0 = 12$. As depicted in Fig. 3(d), the persistent network that evolves in each sample economy in this dominant distance cluster consists of six p-inactive employers and six disjoint latched pairings comprising a work supplier and an employer. Work suppliers are 100% aggressive, and most work suppliers and employers exhibit repeat defection worksite behaviors. The mean market power level attained by work suppliers is $+7\%$ whereas the mean power level attained by employers is only -45% , hence hypotheses H1(b) and H2(b) are both supported.

The remaining sample economies for this E lie in secondary distance clusters 0–4 and $D^0 = 13$. The persistent networks that evolve in the distance cluster 0–4 are mixtures of recurrent and latched relationships. Work suppliers and employers are 100% p-nice and attain similar mean market power levels; the slightly lower level attained by work suppliers is due to job search costs incurred by chance. In distance cluster $D^0 = 13$, the persistent network that evolves in

each sample economy consists of one p-inactive work supplier, seven p-inactive employers, and five disjoint latched pairings comprising a work supplier and an employer. Work suppliers are 100% aggressive and employers are 67% aggressive, and most agents exhibit repeat defection worksite behaviors. The – 16% mean market power level attained by work suppliers is higher than the – 51% mean market power level attained by employers, hence hypotheses H1(b) and H2(b) are supported.

According to the job capacity sensitivity hypothesis H3, work suppliers should be increasingly better off and employers increasingly worse off as job capacity varies from tight to excess, given a low balanced job concentration level. Hypothesis H3(a) is strongly supported by the findings in Table 9, and hypothesis H3(b) is generally supported. As in the previous two sections, the apparently major violation of hypothesis H3(b) indicated by the aggregate market power outcomes for employers reported in Table 6 as one moves from tight to balanced job capacity under conditions of low job concentration arises from a pooling problem: the inclusion in Table 9(a) of a relatively small distance cluster $D^0 = 18$ in which coordination failure is widespread and both agent types attain exceedingly low mean market power levels. Here, however, the violation is not completely eliminated by removing this outlier distance cluster.

Finally, the job concentration sensitivity hypothesis H4 predicts that work suppliers should be better off and employers worse off operating under conditions of *low* job concentration, as reported in Table 9, than work suppliers and employers operating under conditions of *balanced* job concentration as reported in Table 8, for any given job capacity level. The findings reported in Tables 8 and 9 do not support hypothesis H4, even if attention is restricted to dominant distance clusters. Indeed, for each given job capacity level, work suppliers in the dominant distance clusters actually do *worse* as job concentration is decreased from balanced to low whereas employers generally do better.

6. Concluding remarks

As shown in previous sections, the aggregate market power findings of this study indicate that job capacity generally has the hypothesized H1 and H3 effects: all else equal, increased job capacity increases the market power of work suppliers and reduces the market power of employers both in absolute and relative terms. Disaggregated market power findings reveal that these effects are particularly strong and clear when attention is focused on the most commonly observed network formations for each treatment setting. In contrast, neither aggregated nor disaggregated market power findings provide much support for the job concentration hypotheses H2 and H4. On the contrary, controlling for job capacity, the effects of job concentration on the ability of work suppliers and employers to exercise market power are surprisingly small and unsystematic.

Hypotheses H2 and H4 seem a priori intuitive on the grounds that concentrating work offers in fewer work supplier hands should provide work suppliers with an increased opportunity to exercise monopoly power, and concentrating job openings in fewer employer hands should provide employers with an increased opportunity to exercise monopsonist power. On the other hand, it may be that too much concentration lessens the ability of work suppliers or employers as a whole to adapt their worksite strategies in a flexible manner in response to the worksite strategies used by their worksite partners. Consequently, there may be too little diversity in the pool of worksite strategies used by the concentrated agent type for evolutionary selection pressures to efficiently act upon. To test this inflexibility hypothesis, it will be necessary to introduce the absolute numbers of work suppliers and employers as treatment factors in addition to their concentration ratio. It will also be necessary to examine alternative learning algorithms calibrated more carefully to the learning behavior observed in real-world labor markets and in human-subject labor market experiments.

Further work is needed to test the robustness of the findings of this study to variations in the scope and range of other parameter specifications as well. As preliminary as these findings may be, however, they do caution against the common practice of confounding capacity and concentration effects in market power studies by letting these two factors vary together in an uncontrolled manner.

A potentially important by-product of the current study is a better understanding of the fundamental role played by organizational costs in sculpting and sustaining network formations and hence in determining the welfare and market power levels attainable by work suppliers and employers. For example, when job capacity is tight, high job search costs can result in widespread worker discouragement and exit from the labor force that ends up hurting employers as well as work suppliers. Consequently, although tight job capacity tends to favor employers, it also increases the risk of network coordination failure. Organizational costs are a key focus of researchers in transactions costs economics — see, for example, Williamson and Marsten (1999). Nevertheless, the significant problems posed by network organizational costs for the definition and measurement of market power have not received much attention to date.

A second potentially important by-product of the current study is a better appreciation of the role played by behavioral flexibility in protecting agents against the exercise of market power by other agents. In the labor market framework, work suppliers and employers repeatedly choose and refuse their worksite partners and evolve their worksite behaviors over time. Thus, even when jobs are in excess supply, attempts by work suppliers to exert market power by repeatedly defecting in worksite interactions can provoke retaliatory defections by employers as well as firings (refusals of all future work offers from

offending work suppliers). Similarly, even when jobs are tight, attempts by employers to exert market power by repeatedly defecting in worksite interactions can provoke retaliatory defections by work suppliers as well as quits (work suppliers redirecting future work offers elsewhere). Consequently, although structural asymmetries tend to favor one agent type over another, attempts by the favored agent type to exploit this advantage increase the risk of behavioral coordination failure.

In summary, agent-based computational frameworks such as the labor market framework used in this study permit the systematic experimental investigation of behavioral and network formation processes that appear critical for understanding the relation between structure, behavior, and market power in real-world labor markets.

References

- Ashlock, D., Smucker, M.D., Stanley, E.A., Tesfatsion, L., 1996. Preferential partner selection in an evolutionary study of prisoner's dilemma. *BioSystems* 37, 99–125.
- Bresnahan, T.F., 1989. Empirical studies of industries with market power. In: Schmalensee, R., Willig, R.D. (Eds.), *Handbook of Industrial Organization*, Vol. II. Elsevier, Amsterdam, pp. 1011–1057.
- Ehrenberg, R.G., Smith, R.S., 1997. *Modern Labor Economics*, 6th Edition. Addison-Wesley, Reading, MA.
- Gale, D., Shapley, L., 1962. College admissions and the stability of marriage. *American Mathematical Monthly* 69, 9–15.
- Holt, C., 1995. Industrial organization: a survey of laboratory research. In: Kagel, J.H., Roth, A.E. (Eds.), *Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, pp. 349–443.
- McFadzean, D., 1995. SimBioSys: a class framework for evolutionary simulations. Master's Thesis, Computer Science Department, University of Calgary, Alberta.
- McFadzean, D., Tesfatsion, L., 1999. A C++ platform for the evolution of trade networks. *Computational Economics* 14, 109–134.
- Rabin, M., Schrag, J., 1999. First impressions matter: a model of confirmatory bias. *Quarterly Journal of Economics* 114, 37–82.
- Roth, A., Sotomayor, M.A.O., 1990. *Two-sided Matching: A Study in Game-theoretic Modeling and Analysis*. Cambridge University Press, Cambridge, UK.
- Rubinstein, A., Wolinsky, A., 1990. Decentralized trading, strategic behavior, and the Walrasian outcome. *Review of Economic Studies* 57, 63–78.
- Sargent, T.J., 1993. *Bounded Rationality in Macroeconomics*. Clarendon Press, Oxford.
- Tesfatsion, L., 1997a. A trade network game with endogenous partner selection. In: Amman, H., Rustem, B., Whinston, A. (Eds.), *Computational Approaches to Economic Problems*. Kluwer Academic Publishers, The Netherlands, pp. 249–269.
- Tesfatsion, L., 1997b. How economists can get alive. In: Arthur, W.B., Durlauf, S., Lane, D. (Eds.), *The Economy as an Evolving Complex System, II*, Volume XXVII, Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley, Reading, MA, pp. 533–564.
- Tesfatsion, L., 1998. Preferential partner selection in evolutionary labor markets: a study in agent-based computational economics. In: Porto, V.W., Saravanan, N., Waagen, D., Eiben, A.E. (Eds.), *Evolutionary programming*, Vol. VII. Springer, Berlin, pp. 15–24.

- Tesfatsion, L., 1999. Hysteresis in an evolutionary labor market with adaptive search, Economic Report 50, Iowa State University. In: Chen, S.H. (Ed.), *Evolutionary Computation in Economics and Finance*. Springer, Berlin, to appear.
- Tirole, J., 1988. *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.
- Williamson, O.E., Marsten, S.E., 1999. *The Economics of Transactions Costs*. Edward Elgar Publishing, Northampton, MA.