2000

Concentration, capacity and market power in an evolutionary labor market

Leigh Tesfatsion
Iowa State University, tesfatsi@iastate.edu

Concentration, Capacity, and Market Power in an Evolutionary Labor Market

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

9 April 2000

Abstract

This paper reports on an experimental study of the relationship between job capacity, job concentration, and market power in the context of an agent-based computational model of a labor market. 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 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.

1 Introduction

Market power refers to the ability of buyers and sellers to exert a perceptible control over market outcomes that enables them to attain higher individual welfare levels than they would achieve under competitive market conditions. Economists have used analytical modelling (Tirole, 1988), empirical studies (Bresnahan, 1989), and human-subject laboratory experiments (Holt, 1995) in an attempt to understand the relationship between market structure, market behavior, and market power in markets with multiple agents engaged in repeated strategic interactions. In view of the complex nature of this relationship, however, definitive conclusions have been difficult to obtain.

This paper investigates the evolution of market power in the context of a computational labor market framework with strategically interacting agents.[1] Multiple work suppliers and employers repeatedly participate in costly searches for worksite partners on the basis of continually updated expected utility, engage in 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

[1] The present paper provides a summary overview of a more detailed experimental study (Tesfatsion, 1999) in press. All experiments reported in this study are implemented using version 105c of the Trade Network Game (TNG) source code developed by (Tesfatsion, 1997) and (McFadzean and Tesfatsion, 1999), which in turn is supported by SimBioSys, a general C++ class framework for evolutionary simulations developed by (McFadzean, 1995). Source code for both the TNG and SimBioSys can be downloaded as freeware at the current author's web site, along with extensive user instructions.

$(NE \cdot eq)/(NW \cdot wq)$ 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.

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 comprises twelve 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. The primary objective of the 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 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 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: 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: 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.

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. In (Tesfatsion, 1999) a more detailed examination of the experimental data is undertaken at a disaggregated level to determine correlations among market structure, network formations, worksite behaviors, and welfare outcomes. 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.

2 Labor Market Framework

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.[2]

Work suppliers and employers are modelled as autonomous interacting agents with internalized social norms, internally stored state information, and internal behavioral rules. Each agent, whether a work supplier or an employer, has this same general internal structure. However, work suppliers differ from employers in terms of their specific market protocols, fixed attributes, and initial endowments; and all agents can acquire different state information and evolve different worksite behavioral rules over time on the basis of their past experi-

[2] 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.

ences. In particular, 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 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 1, 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 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.

Brief descriptions of the implementations used for the agent activities appearing in Table 1 will now be given. More detailed descriptions can be found in (Tesfatsion, 1999).

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). 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 who 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 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 accept 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. 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, 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 are assumed to be measured in utility terms and to be normalized about the inactivity payoff 0 so that $L < D < 0 < C < H$. Thus, a work supplier or employer that ends up either as a sucker with payoff L or in a mutual defection relation with payoff D receives negative utility, a worse outcome than inactivity (unemployment or vacancy). 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.

```
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();               //    Con¯gure 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;
}
```

Table 1: Logical Flow of the Labor Market Framework.

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

The 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. This means, in particular, 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 biased in favor of more fit agents. 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.

3 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 2, 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 3 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.

For each tested potential economy E, twenty sample economies (s, E) were experimentally generated using twenty arbitrarily selected seed values s for the pseudo-random number generator included in the TNG source code.[3] For each run s, the "market power profile" was then determined and recorded, as explained in the next section.

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

4 Measurement of Market Power

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 obtains 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).

As shown in (Tesfatsion, 1999), 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

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

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

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

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

	Tight Job Capacity JCAP=1/2	Balanced Job Capacity JCAP=1	Excess Job Capacity JCAP=2
High Job Concentration JCON=2	NW=12 NE=6 wq=1 eq=1	NW=12 NE=6 wq=1 eq=2	NW=12 NE=6 wq=1 eq=4
Balanced Job Concentration JCON=1	NW=12 NE=12 wq=2 eq=1	NW=12 NE=12 wq=1 eq=1	NW=12 NE=12 wq=1 eq=2
Low Job Concentration JCON=1/2	NW=6 NE=12 wq=4 eq=1	NW=6 NE=12 wq=2 eq=1	NW=6 NE=12 wq=1 eq=1

Table 2: Two-Factor Experimental Design

// 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^o.
Elite = 67	// GA elite percentage for each agent type.
MutationRate = .005	// GA mutation rate (bit toggle probability).
FsmStates = 16	// Number of internal FSM states.
FsmMemory = 1	// FSM memory (in bits) for past move recall.
// PARAMETER VALUES VARIED ACROSS EXPERIMENTS	
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.

Table 3: Illustrative Parameter Values for a Potential Economy E with JCON=2 and JCAP=1.

reach and sustain any market state, whether competitive or not.

5 Experimental Results

As detailed in Section 1, the primary objective of this study is to test hypotheses H1 through 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 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 below; a detailed report on market power findings at a more disaggregated level is given in (Tesfatsion, 1999).

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 4 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 detailed in (Tesfatsion, 1999), 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 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 twenty 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 4, 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 4 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, the outcomes reported in Table 4 do not support the job concentration hypotheses H2 and H4.

The standard deviations reported in Table 4 tend to be large. The disaggregated findings reported in (Tesfatsion, 1999) show that these large standard deviations arise from a pooling problem. The sample runs generated for each cell of the experimental design in Table 2 fail to exhibit a central tendency; rather, they cluster around two or three distinct types of outcomes, suggesting the existence of multiple basins of attraction. Generally one cluster within each cell is dominant, in the sense that most sample runs lie within this cluster, and the job capacity hypotheses H1 and H3 receive particularly strong support when attention is restricted to these dominant clusters. Averaging market power outcomes across the multiple clusters for each cell weakens the indicated support for hypotheses H1 and H3 and increases measured standard deviations.

6 Concluding Remarks

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. In contrast, these findings do not support the job concentration hypotheses H2 and H4. To 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 "genetic diversity" in the pool of worksite strategies used by the concen-

	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)
	w	e	w	e	w	e
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)
	w	e	w	e	w	e
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)

Table 4: 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.

trated 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, and to examine alternative learning algorithms calibrated more carefully to observations of learning behavior 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 in economics of confounding capacity and concentration effects in market power studies by letting these two factors vary together in an uncontrolled manner.

Bibliography

Bresnahan, T. F. (1989), "Empirical Studies of Industries with Market Power," in: R. Schmalensee and R. D. Willig, eds., *Handbook of Industrial Organization, Vol. II*, Elsevier Science Publishers B.V., Amsterdam, 1011–1057.

Gale, D. and L. Shapley (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: J. H. Kagel and A. E. Roth, eds., *Handbook of Experimental Economics*, Princeton University Press, Princeton, N.J., 349–443.

McFadzean, D. (1995), *SimBioSys: A Class Framework for Evolutionary Simulations*, Master's Thesis, Computer Science Department, University of Calgary, Alberta.

McFadzean, D. and L. Tesfatsion (1999), "A C++ Platform for the Evolution of Trade Networks," *Computational Economics* 14, 109–134.

Tesfatsion, L. (1997), "A Trade Network Game with Endogenous Partner Selection," in: H. Amman, B. Rustem, and A. Whinston, eds., *Computational Approaches to Economic Problems*, Kluwer Academic Publishers, The Netherlands, 249–269.

Tesfatsion, L. (1999), "Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search," Economic Report 51, Iowa State University, to appear in the *Journal of Economic Dynamics and Control*.

Tirole, J. (1988), The theory of industrial organization (MIT Press, Cambridge, MA).