ISU Economic Report 45
Revised 15 November 1998

Teaching Agent-Based Computational Economics to Graduate Students

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

Abstract: Agent-based computational economics (ACE) is roughly defined as the computational study of economies modelled as evolving decentralized systems of autonomous interacting agents. A key focus of ACE research is understanding how global regularities arise from the bottom up, through the repeated local interactions of autonomous agents channeled through socio-economic institutions, rather than from top down coordination mechanisms such as imposed market clearing constraints or an assumption of single representative agents. This paper discusses how ACE materials have been introduced into graduate-level courses in macroeconomic theory over the past several years, using an ACE labor market framework for concrete illustration.

1 Introduction

The newly developing field of agent-based computational economics (ACE) is roughly defined by its practioners as the computational study of economies modelled as evolving decentralized systems of autonomous interacting agents. A principal concern of ACE researchers is to understand the apparently spontaneous formation of global regularities in economic processes, such as the unplanned coordination of trading activities in decentralized market economies that economists associate with Adam Smith's invisible hand. The challenge is to explain how these global regularities arise through repeated local interactions of autonomous agents channeled through socio-economic institutions rather than through fictitious coordination mechanisms such as imposed market clearing constraints or an assumption of single representative agents.

The study of evolutionary economics is by no means new, of course. Even before Darwin, attempts were made to apply evolutionary ideas to socio-economic behavior (Richards [47]). Although this early work is now largely ignored by economists, economic textbooks still typically include at least some mention of the ideas of Joseph Schumpeter [53] regarding the evolution of economic institutions.

Moreover, Schumpeter's work, together with the seminal work by Armen Alchian [2] on uncertainty and evolution in economic systems, appears to have strongly influenced the subsequent well-known work by Nelson and Winter [45] and various of their collaborators on evolutionary theories of economic change. In addition, one has the work of W. Brian Arthur on economies incorporating positive feedbacks, the work by Richard Day on dynamic economies characterized by complex phase transitions, the work by John Foster on an evolutionary approach to macroeconomics, Ron Heiner's work on the origins of predictable behavior, Jack Hirshleifer's work on evolutionary models in economics and law, and Ulrich Witt's work on economic natural selection. These and numerous other interesting studies on evolutionary economics are reviewed by Witt [63] and Nelson [44]. More recently, as detailed in Friedman [19], Fudenberg and Levine [20], Hofbauer and Sigmund [25], and Samuelson [51], a number of researchers have been focusing on the potential economic applicability of evolutionary game theory with replicator dynamics in which game strategies distributed over a fixed number of strategy types reproduce over time in direct proportion to their relative fitness.

Exploiting the recent advent of more powerful computational tools, such as object-oriented programming, ACE researchers have been able to extend this earlier work on evolutionary economics in four key ways. First, agents in ACE frameworks are typically modelled as heterogeneous entities that determine their interactions with other agents and with their environment on the basis of internalized data and behavioral rules. These agents thus tend to have a great deal more internal cognitive structure and autonomy than conventionally modelled economic agents. Second, a broader range of agent interactions is typically permitted in ACE frameworks, with predatory and cooperative associations taking center stage along with price and quantity relationships. Third, the evolutionary process is generally represented as natural selection pressures acting directly on agent characteristics rather than as population-level laws of motion. These natural selection pressures result in the continual creation of new modes of agent behavior and an ever-changing network of agent interactions. Fourth, ACE frameworks are computer implemented as virtual economic worlds that grow themselves along a real time-line, much like a culture dish develops in a laboratory. In principle, once initial conditions are set, all subsequent events in these virtual economic worlds are initiated and driven by agent-agent and agent-environment interactions; no further outside interventions by the modeler (e.g., off-line fixed point calculations) are permitted.

In brief, then, ACE is a blend of concepts and tools from evolutionary economics, cognitive science, and computer science.[1] It represents a methodological approach that may ultimately permit two important developments: (a) the rigorous testing, refinement, and extension of theories developed in the earlier literature on evolutionary economics that were found to be analytically intractable; and (b) the rigorous formulation and testing of conceptually integrated socio-economic theories compatible with theory and data from many different relevant fields currently separated by artificial disciplinary boundaries. Examples of studies focusing on key ACE-related issues include: Anderson et al. [3]; Arifovic [4]; Arthur [5]; Arthur et al. [6]; Axelrod [8]; Bell [9]; Birchenhall [10]; Bullard and Duffy [11]; De Vany [13]; Duong [16]; Durlauf [17]; Epstein and Axtell [18]; Gode and Sunder [22]; Guriev and Shakhova [24]; Holland [26]; Holland and Miller [27]; Kirman [29, 30, 31]; Kollman et al. [32]; Lane [33]; Mailath et al. [37]; Marimon et al. [38]; Marks [39]; McFadzean and Tesfatsion [41]; Miller [42]; Routledge [49]; Rust et al. [50]; Sargent [52]; Shubik [54]; Tesfatsion [57, 59, 60, 61]; Vriend [62]; and Young [64, 65].

How might ACE be taught to graduate students in a typical department of economics today? In keeping with the subject matter of ACE, as well as the newness of the methodological approach

[1] As explained more carefully in Tesfatsion [58], ACE can be viewed as a specialization to economics of the basic artificial life (alife) paradigm. For interesting introductions to alife, see Levy [35], Lindgren and Nordahl [36], and Sigmund [55].

for economic study, I do not believe that any definitive top-down answer to this question can or should be given; let 550 flowers bloom.[2] Rather, I will simply suggest one way this teaching can be initiated by outlining how I have introduced ACE materials in macroeconomics theory courses for Ph.D. and Masters level students at Iowa State University over the past several years.[3]

Briefly, in keeping with the importance of coordination issues for macroeconomics, I devote a third of the course to this topic. I start by presenting to students a bare-bones version of the basic economic paradigm of coordination success, Walrasian general equilibrium. As detailed in Katzner [28], this paradigm represents a precisely formulated set of conditions under which feasible allocations of goods and services can be supported by price systems in decentralized market economies characterized by price-taking consumers and firms and private ownership of capital and labor. Its defining structural characteristic is that direct agent-agent interactions do not take place. Rather, all agent-agent interactions are mediated by an implicit clearing house colorfully referred to as the Walrasian Auctioneer.[4]

More precisely, *strategic interaction* is said to occur between two agents if the choice of a decision for at least one of the agents depends upon what he perceives or expects the decision of the other agent to be. The Walrasian general equilibrium model reflects the view that, in decentralized market economies, price systems reduce or even eliminate the need for economic agents to interact strategically. The key observation here is that values for prices and dividend payments constitute the only information conveyed to consumers and firms in the Walrasian general equilibrium model. Since prices and dividend payments are treated as parameters by these agents in their decision problems, these decision problems reduce to "control" problems. That is, the decision problem for each agent only includes decision variables fully under the agent's own control; the decision variables for other agents do not appear, implying there is no strategic interaction. In systems science

[2] Splitting the difference between Mao Tse-Tung's 100 flowers and George H. W. Bush's 1000 points of light.

[3] See www.econ.iastate.edu/classes/econ502/tesfatsion/ for a complete syllabus for one such course.

[4] This terminology appears to have been introduced by Leijonhufvud [34] in his interesting critique of the Walrasification of Keynes' general theory that occurred when attempts were made to represent this theory within an IS-LM framework.

parlance, the global allocation problem has been decomposed into a collection of individual agent allocation problems by the introduction of linking variables (prices and dividends). The equilibrium values for the linking variables are determined by calculations performed by the fictitious Walrasian Auctioneer; they do not arise from any actions of the consumers or firms within the model.

To test the robustness of the Walrasian general equilibrium model to changes in its structure, I then ask students to introduce one "simple" change into this paradigm — namely, let firms set their own prices — at which point the paradigm is seen to collapse like a house of cards. This leads naturally into a discussion of the need to consider a more comprehensive and realistic modelling of agent *interactions* in relation to macro regularities, a key focus of ACE research. I then present various illustrations of non-Walrasian modelling, including ongoing ACE research that appears to be particularly relevant for studying the self-organizing capabilities of decentralized market economies. Students interested in further ACE study are directed to the extensive resources I maintain on an ACE Web site at www.econ.iastate.edu/tesfatsi/ace.htm, including ACE surveys, an annotated syllabus of ACE-related readings, pointers to ACE software, and pointers to other ACE-related Web sites.

In the following sections, for concreteness, I discuss the specification and experimental implementation of an ACE framework for studying the formation and evolution of contractual networks in labor markets with adaptive search and worksite behavior.[5] While not a full-blown multiple-market treatment of a decentralized market economy, the labor market framework demonstrates how an ACE approach facilitates the modelling of markets from an agent-based perspective and permits the rigorous experimental study of non-steady-state dynamics. The framework builds on a series of earlier studies [7, 56, 57, 59, 60].

[5]The ACE labor market framework is implemented by means of the C++ trade network game (TNG) framework (version 105b) developed by McFadzean and Tesfatsion [41], which in turn is supported by SimBioSys, a general C++ class library for evolutionary simulations developed by McFadzean [40]. The source code for both SimBioSys and the TNG framework are available for downloading as freeware at the current author's Web site. A more detailed discussion of this ACE labor market framework in relation to the existing labor market literature can be found in Tesfatsion [61].

2 An ACE Labor Market Framework

An interesting theoretical literature stressing job search and matching in labor markets has flourished since the influential work by Diamond [14, 15] on search equilibrium. See, for example, Aghion and Howitt [1, Chapter 4]. To achieve analytical tractability, however, researchers in this literature commonly postulate an aggregate matching function that proxies the complicated process of employer recruitment, worker search, and mutual evaluation. That is, the intense *individualistic* rivalry among workers and employers that characterizes many real-world labor markets is not modelled. Moreover, again for tractability, the competition of ideas within agents is generally ignored in this literature; attention is largely focused on steady-state equilibrium behavior.

In Tesfatsion [60] it is conjectured that some of the tractability problems encountered in analytical labor market studies might be alleviated by taking an ACE approach. To explore this possibility, a simple labor market framework is developed that builds on the ACE trade network game (TNG) developed in Tesfatsion [57, 58] for studying the formation and evolution of trade networks under alternatively specified market structures. As will be clarified below, this labor market framework endogenizes many aspects of labor markets at the level of individual agents that theoretical labor market studies generally specify in a more restricted static way either as exogenously given parameters or through aggregate steady-state relationships: for example, worker preference orders over potential employers; employer preference orders over potential workers; job search and search costs; contractual matching; worker and employer worksite behaviors; worker compensations; employer earnings; quit rates; firing rates; turnover costs; and unemployment rates.

The ACE labor market framework consists of three disjoint (and possibly null) subpopulations of agents that separately evolve over time: pure workers who make work offers; pure employers who receive work offers; and worker-employers capable of both making and receiving work offers. The pure workers and worker-employers are collectively referred to as *workers*, and the worker-employers and pure employers are collectively referred to as *employers*. Each worker can have

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

Table 1: Logical Flow of the ACE Labor Market Framework

no more than wq work offers outstanding to employers at any given time, and each employer can accept no more than eq work offers from workers at any given time, where the work offer quota wq and the acceptance quota eq can be any positive integers.[6] Although highly simplified, these parametric specifications will be seen in Section 4, below, to permit the study of a variety of labor market structures operating under different ex ante capacity constraints.

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

This labor market framework facilitates the study of labor markets from an agent-based per-

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

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

Table 2: Schematic Description of an Agent

spective in two key ways. First, as depicted in Table 2, each agent is instantiated as an autonomous endogenously interacting software agent with internally stored state information and internal behavioral rules. The agents can therefore engage in anticipatory behavior. Moreover, using stored agent addresses together with internalized communication protocols, they can communicate with each other at event-triggered times, a feature not present in standard economic models.

Second, as seen in Table 1, the labor market framework is modular in design. This means that experimentation with alternative specifications for market structure, search and matching among workers and employers, worksite interactions, expectation formation and updating, and evolution of worksite strategies can easily be undertaken — much like changing a lightbulb in a multi-bulb lamp — as long as the interfaces (inputs and outputs) for the modules implementing these specifications remain unchanged. Moreover, each of these modules can potentially be grounded in agent-initiated actions in the sense that the module is implemented via behavioral rules internal to the agents. Finally, the transitory and longer-run implications of each alternative module specification can be studied at three different levels: individual characteristics of workers and employers; interactions among workers and employers (network formation); and social welfare as measured by descriptive statistics such as average agent welfare and unemployment rates.

A brief description will now be given for the particular module specifications used in all experiments reported below. See McFadzean and Tesfatsion [41] for a more careful description.

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

Specifically, each matched worker and employer engage in a worksite interaction modelled as a two-person prisoner's dilemma game reflecting the basic efficiency wage hypothesis that worker effort levels are affected by overall working conditions (e.g., wage levels, respectful treatment, safety considerations, ...). The worker can either cooperate (exert high work effort) or defect (engage in shirking). Similarly, the employer can either cooperate (provide good working conditions) or defect (provide substandard working conditions). The range of possible payoffs is assumed to be the same for each match in each trade cycle: namely, as seen in Table 3, a cooperator whose contractual partner defects receives the lowest possible payoff L (the sucker payoff); a defector whose contractual partner also defects receives a payoff D; a cooperator whose contractual partner also cooperates receives a payoff C; and a defector whose contractual partner cooperates receives the highest possible payoff H (the temptation payoff). The payoffs are assumed to be measured in utility terms and to be normalized about 0, so that $L < D < 0 < C < H$. They are also assumed to satisfy the usual regularity condition $(L + H)/2 < C$ guaranteeing that mutual cooperation dominates alternating cooperation and defection on average.

Matches between workers and employers are determined using a modified version of the well-studied "deferred acceptance mechanism" originally designed by Gale and Shapley [21].[7] Under this modified mechanism, hereafter referred to as the *deferred choice and refusal* (DCR) mechanism, each worker submits up to wq work offers to employers he ranks as most preferable on the basis

[7] See Roth and Sotomayor [48] for a careful detailed discussion of Gale-Shapley deferred acceptance matching mechanisms, including a discussion of the way in which the Association of American Medical Colleges since WWII has slowly evolved such an algorithm (the National Intern Matching Program) as a way of matching interns to hospitals in the United States.

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

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

of expected payoff and who he judges to be tolerable in the sense that their expected payoff is not negative. Similarly, each employer selects up to eq of his received offers that he finds tolerable and most preferable on the basis of expected payoff and he places them on a waiting list; all other offers are refused. Workers redirect refused offers to tolerable preferred employers who have not yet refused them, if any such employers exist. Once employers stop receiving new offers, they accept all work offers currently on their waiting lists.

A worker incurs a transactions cost in the form of a negative *refusal payoff R* whenever an employer refuses one of his offers during the matching process; the employer who does the refusing is not penalized.[8] An agent who neither submits nor accepts work offers during the matching process receives a *wallflower payoff 0*. The refusal and wallflower payoffs are each assumed to be measured in utility terms.

Agents use a simple learning algorithm to update their expected utilities on the basis of new payoff information. Each agent v assigns an exogenously given initial expected utility U^o to each potential contractual partner z with whom he has not yet interacted. Each time an interaction with z takes place, v forms an updated expected utility assessment for z by summing U^o together with all payoffs received to date from interactions with z and dividing this sum by one plus the number of interactions with z.

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

The worksite behavior of each agent is governed by a finite-memory pure strategy for playing a prisoner's dilemma game with an arbitrary partner an indefinite number of times, hereafter referred to as a *worksite strategy*. At the end of each trade cycle loop, the worksite strategies of pure workers, pure employers, and worker-employers are separately evolved by means of a standardly specified genetic algorithm involving recombination, mutation, and elitism operations.[9] This evolution is meant to reflect the formation and transmission of new ideas rather than biological reproduction. Specifically, if a worksite strategy successfully results in high fitness for an agent of a particular type, where fitness is measured by average payoff, then other agents of the same type are led to modify their own strategies to more closely resemble the successful strategy.

An important caution is in order here, however. Given the extent of information currently allowed to agents during the evolution step — i.e., knowledge of the complete strategies of all other agents of the same type, whether expressed in interactions or not — the evolution step is more appropriately interpreted as a stochastic search algorithm for determining dominant outcomes rather than as a cultural transmission mechanism per se. The resulting welfare outcomes will be used in subsequent work as benchmarks against which to assess the effectiveness of more realistically modelled cultural transmission mechanisms.

3 Descriptive Statistics

In this section care is taken to explain the ex ante and ex post measures that have been constructed to aid in the experimental determination of correlations between ex ante market structure

[9] More precisely, for each subpopulation of agents, the genetic algorithm evolves a new collection of agent worksite strategies from the existing collection of agent worksite strategies by applying the following four steps: (1) *Evaluation*, in which a fitness score is assigned to each strategy in the existing strategy collection; (2) *Recombination*, in which offspring (new ideas) are constructed by combining the genetic material (structural characteristics) of pairs of parent strategies chosen from among the most fit strategies in the existing strategy collection; (3) *Mutation*, in which additional variations (new ideas) are constructed by mutating the structural characteristics of each offspring strategy with some small probability; and (4) *Replacement*, in which the most fit (elite) worksite strategies in the existing collection of strategies are retained for the new collection of strategies and the least fit worksite strategies in the existing strategy collection are replaced with offspring strategies. See McFadzean and Tesfatsion [41] for a more detailed discussion of this use of genetic algorithms in the TNG, and see Goldberg [23] and Mitchell and Forrest [43] for a general discussion of genetic algorithm design and use.

and ex post contractual network formation, and between contractual network formation and the types of worksite behaviors and social welfare outcomes that these contractual networks support. Contractual networks depict who is working for whom, and with what regularity. Worksite behavior refers to the specific actions undertaken by a worker (or employer) in worksite interactions with any given employer (worker). Finally, social welfare measures the overall utility achieved by the workers and employers from repeated worksite interactions within the context of a possibly changing network of contractual partners.

3.1 Classification of Contractual Network Types by Distance

Let s denote a pseudo-random number generator seed value for the TNG source code used to implement the ACE labor market framework (see footnote 5), and let e denote a *potential TNG economy*, i.e., an economy characterized structurally by the TNG source code together with all of the user-specified TNG parameter values apart from s. The *realized TNG economy* generated from e, given the seed value s, is denoted by (s, e).

Since worksite strategies are represented as finite state machines,[10] the actions undertaken by any agent v in repeated worksite interactions with another agent z must eventually cycle. Consequently, these actions can be summarized in the form of a *worksite history* $H{:}P$, where the *handshake* H is a (possibly null) string of worksite actions that form a non-repeated pattern and the *persistent portion* P is a (possibly null) string of worksite actions that are cyclically repeated. For example, letting c denote cooperation and d denote defection, the worksite history $ddd{:}dc$ indicates that v defected against z in his first three worksite interactions with z and thereafter alternated between defection and cooperation.

Two agents v and z are said to exhibit a *persistent relationhip* during a given trade cycle loop T of a realized TNG economy (s,e) if the following two conditions hold: (a) their worksite histories

[10] A *finite state machine (FSM)* is a system comprising a finite collection of internal states together with a state transition function that gives the next internal state the system will enter as a function of the current state together with current system inputs. For the application at hand, the inputs are the actions selected by a worker and an employer engaged in a worksite interaction. See McFadzean and Tesfatsion [41] for a more detailed discussion and illustration of the FSM representation used in the TNG source code.

with each other during the course of T take the form $H_v{:}P_v$ and $H_z{:}P_z$ with nonnull P_v and P_z; and (b) accepted work offers between v and z do not permanently cease during T either by choice (a permanent switch away to strictly preferred contractual partners) or by refusal (one agent becoming intolerable to the other because his expected utility drops below zero).

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

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

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

3.2 Classification of Worksite Behaviors and Welfare Outcomes

An agent v in a realized TNG economy (s, e) is referred to as an *unprovoked defector (UD)* if he engages in at least one defection against another agent who has not previously defected against him. The vector giving the separate UD percentages for pure workers, pure employers, and worker-employers in the final generation of (s, e) is referred to as the *UD profile* for (s, e). The UD profile measures the extent to which the different types of agents behave agressively in worksite interactions with contractual partners who are either strangers or who so far have been consistently cooperative.[11]

Moreover, v is referred to as a *persistent wallflower (PWF)* if v constitutes an isolated node of the persistent contractual network $K^p(s, e)$. Alternatively, v is referred to as a *repeat defector (RD)* if v establishes at least one persistent relationship for which the persistent portion P of his worksite history $H{:}P$ includes a defection d. If, instead, v establishes at least one persistent relationship and his worksite history for each of his persistent relationships has the general form $H{:}c$, he is referred to as a *persistent cooperator (PC)*.

The vectors giving the separate PWF, RD, and PC percentages for pure workers, pure employers, and worker-employers in the final generation of (s, e) are referred to as the *PWF profile*, the *RD profile*, and the *PC profile* for (s, e), respectively. The PWF profile measures the extent to which the different types of agents fail to establish any persistent relationships, whereas the RD and PC profiles measure the extent to which the different types of agents establish persistent relationships characterized by predacious or fully cooperative behavior, respectively. By construction, an agent must either be a PWF, a RD, or a PC. Thus, only the PWF and PC profiles are reported in the experiments discussed below.

The vector that separately gives the average fitness score for pure workers, pure employers, and worker-employers, respectively, in the final generation of a realized TNG economy (s, e) is referred

[11] The importance of stance towards strangers in determining subsequent outcomes in path dependent contexts such as the ACE labor market framework has been stressed by Orbell and Dawes [46].

to as the *FIT profile* for (s, e). The FIT profile constitutes a measure of social welfare.

4 Some Experimental Findings

The labor market experiments reported below focus on three simple labor market structures: endogenous-type markets comprising 24 worker-employers; two-sided markets comprising 12 pure workers and 12 pure employers; and partially-fluid markets comprising 8 pure workers, 8 pure employers, and 8 worker-employers. Within each market structure, four different configurations for the worker offer quota wq and employer acceptance quota eq are examined: high excess capacity ($eq >> wq$); zero excess capacity ($eq = wq = 1$); tight capacity ($eq = 1$ and $wq = 2$); and extremely tight capacity ($eq << wq$). The genetic algorithm elite value is automatically adjusted in each experiment to maintain the elite proportion at approximately 67% for each nonzero agent type.

The values for all remaining parameters are maintained at fixed values throughout all experiments. Table 4 lists these fixed parameter values along with the specific agent type values, quota values, and elite value for a two-sided market experiment with high excess capacity. The parameter values in Table 4, together with the TNG source code, constitute a potential TNG economy e in the sense defined in Section 3.1.

For each tested e, twenty TNG economies (s, e) were experimentally generated using twenty arbitrarily selected seed values s for the TNG pseudo-random number generator.[12] The persistent contractual network $K^p(s, e)$ for each run s was determined and graphically depicted, and the mean and standard deviation for the UD (unprovoked defector), PWF (persistent wallflower), PC (persistent cooperator), and FIT (fitness) profiles were determined and recorded.

A base contractual pattern $V^o(e)$ was then specified for each tested e. Although the choice of this base pattern is simply a normalization determining a 0 point for the distance measure D^o, and hence intrinsically arbitrary, the degree of specificity of this base pattern governs the dispersion of the resulting distance values and the extent to which these distance values display useful correlations

[12] 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.

```
// PARAMETER VALUES HELD FIXED ACROSS EXPERIMENTS
    GMax = 50                  // Total number of generations.
    IMax = 150                 // Number of trade cycles in each trade cycle loop.
    MutationRate = .005        // GA bit toggle probability.
    FsmStates = 16             // Number of internal FSM states.
    FsmMemory = 1              // FSM memory (in bits) allocated to past move recall.
    RefusalPayoff = -0.5       // Payoff R received by a refused agent.
    WallflowerPayoff = +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°.
    AgentCount = 24            // Total number of agents.
// PARAMETER VALUES VARIED ACROSS EXPERIMENTS
    PureWorkers = 12           // Number of pure workers.
    PureEmployers = 12         // Number of pure employers.
    WorkerEmployers = 0        // Number of worker-employers.
    Elite = 8                  // Number of elite for each nonzero agent type.
    WorkerQuota = 1            // Worker offer quota wq.
    EmployerQuota = 12         // Employer acceptance quota eq.
```

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

with worksite behaviors as measured by the UD, PWF, PC, and FIT profiles. In practice, then, the choice of the base contractual pattern was fine-tuned so that the resulting distance values provided a meaningful informative classification of network types. Given $V^o(e)$, the distance $D^o(s,e)$ of $K^p(s,e)$ from $K^o(e)$ was recorded for each run s, and a histogram for the distance values $D^o(s,e)$ was constructed giving the percentage of runs s corresponding to each possible distance value.

Finally, as a rough stability check, the number of generations was also increased to 100 for each tested potential economy e and the minimum, maximum, and average fitness scores for the agents in each of the 100 generations were graphically generated for each realised economy (s,e).

One interesting finding observed for the tested potential economies e is the remarkable stability exhibited by the average agent fitness scores over generations 25 through 100 for many of the corresponding realized economies (s,e), with stability often setting in as early as generation 10. This observed stability in average fitness scores occurs despite the ceaseless change in the underlying worksite strategies induced by repeated application of genetic algorithm operations. Cases in which instabilities were detected in average fitness scores are noted in the discussion of specific experimental findings, below.

Another interesting finding observed for many of the tested economies e is the existence of multiple distinct types of persistent contractual networks $K^p(s,e)$, each supporting a distinct pattern of worksite behaviors. More precisely, for each e, the distance values $D^o(s,e)$ for the persistent contractual networks tend to cluster around a small number of isolated distance values, and the mean distance of each distance cluster tends to be strongly correlated with the mean UD, PWF, PC, and FIT profiles calculated for the cluster. For such economies, then, there does not appear to be any central-tendency network in the sense defined by Banks and Carley [12] but rather a number of different local basins of attraction. One possible explanation for these distinct distance clusters is that they correspond to multiple Nash equilibria for the underlying evolutionary match-and-play game in which the agents are participating. On the other hand, the distinct distance clusters could be artifacts of the relatively small sample size of 20 that was used in these experiments in order to keep the graphical determination and analysis of network formations manageable. More testing is needed here.

A third interesting finding is that the optimality criteria conventionally used to evaluate the performance of matching mechanisms in static market contexts turn out to be highly incomplete indicators of performance from an evolutionary vantage point. The static viewpoint hides the strong role played by market structure and ex ante capacity constraints in determining the types of persistent matching networks that evolve, the types of persistent interaction behaviors that these networks support, and the transactions costs and inactivity costs to agents that the achievement of these persistent networks and behaviors entails. In addition, the static viewpoint takes preference rankings over potential partners as given whereas these rankings are continually updated on the basis of past interactions in evolutionary settings. Indeed, matching networks and interaction behaviors evolve conjointly. This suggests the need for more comprehensive optimality criteria that take both facets into account.

More concretely, in all of the labor market experiments reported here, the DCR mechanism described in Section 2 is used to match workers and employers. The matching outcomes generated via the DCR mechanism have been shown (Tesfatsion [57, 58]) to have the usual optimality properties associated with Gale-Shapley type matching mechanisms: namely, pairwise stability; and Pareto optimality from the vantage point of workers, the agents who actively make offers. Nevertheless, the actual evolutionary outcomes observed in these labor market experiments include autarkic economies in which all agents are persistent wallflowers, exploitive economies in which employers persistently defect against cooperate workers or workers persistently defect against cooperative employers, and fully harmonious economies in which all agents are persistent cooperators. Moreover, due to transactions costs (negative R payoffs) and inactivity costs (0 wallflower payoffs), social welfare can still be low even if all active agents are persistent cooperators. These evolutionary outcomes are systematically related to market structure and to ex ante capacity constraints as represented by the worker offer quota wq and the employer acceptance quota eq.

A more detailed summary of the findings for each market structure will now be given.

4.1 Endogenous-Type Labor Market Experiments

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

— Insert Figure 1 About Here —

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

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

— Insert Table 5 About Here —

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

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

through 100.

— Insert Figure 2 About Here —

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

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

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

4.2 Two-Sided Labor Market Experiments

Next consider the case of a two-sided labor market economy e comprising 12 pure workers and 12 pure employers with a worker offer quota $wq = 1$ and an employer acceptance quota $eq = 12$, implying that excess capacity is high. Indeed, the structural risk to workers of having their offers refused by employers on the basis of limited acceptance capacity is zero. In contrast, the employers are forced to be inactive unless workers happen to direct work offers their way, implying that the employers face a substantial structural risk of incurring wallflower payoffs. The economy e thus represents a "workers' market." As depicted in Figure 3(a), the base contractual pattern $V^o(e)$ for this economy e is as follows: Each worker is latched to at least one employer, and no employer is a wallflower.

— Insert Figure 3 About Here —

As seen in Table 6(a), 75% of the runs (s, e) for this high capacity e were observed to lie in the distance cluster 3–9. In this distance cluster, the very low mean FIT value of 0.35 for employers is due to two factors: a high level of inactivity (high mean PWF percentage) due to high excess capacity; and aggressive and persistently predacious behavior (high mean UD and low mean PC percentages) by workers that induces retaliatory RD behavior in some employers. The persistent contractual networks resulting from the runs (s, e) in distance cluster 3–9 reveal the following typical scenario: RD workers latch on to a selected subset of employers and drive down their fitness scores to small positive values, causing the remaining employers to become inactive PWFs with fitness scores very close to 0 — indeed, the magnitudes of the distances $D^o(s, e)$ for these runs is essentially a count of the number of inactive employers. This ensures that the worksite strategies of the exploited employers are advantaged in the evolution step relative to the worksite strategies of the employers who are inactive. Since (pure) workers and (pure) employers evolve separately, the worksite strategies of the exploited employers tend to reproduce into the next generation. In this way the workers breed and maintain a subset of wimpy employers that they can repeatedly exploit to their benefit.

— Insert Table 6 About Here —

Table 6(a) also shows that the remaining 25% of the runs for this e lie in a second distance cluster 23–24. The mean FIT value of 1.02 achieved by employers in this second distance cluster is higher than that achieved in distance cluster 3–9 due to the higher mean percentage of PC behavior exhibited by both workers and employers. This mean FIT value is nevertheless substantially below the mutual cooperation payoff level, 1.40, due to the 5% inactivity level among employers, a structual consequence of high excess capacity that is independent of how cooperatively the employers behave in their worksite interactions. The typical contractual pattern exhibited in this distance cluster is PC workers randomly directing work offers among employers without latching. Note that the mean FIT value 1.39 achieved by workers is very close to the mutual cooperation payoff level.

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

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

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

4.3 Partially-Fluid Labor Market Experiments

Finally, consider a partially-fluid labor market economy e comprising 8 pure workers (pw), 8 pure employers (pe), and 8 worker-employers (we) with a worker offer quota $wq = 1$ and an employer acceptance quota $eq = 16$, implying that excess capacity is high. As depicted in Figure 4(a), the base contractual pattern $V^o(e)$ for this economy e is as follows: Each worker directs work offers to employers without latching, and no pure employer is a wallflower.

— Insert Figure 4 About Here —

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

— Insert Table 7 About Here —

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

It therefore appears that the instabilities observed in average fitness scores for these individual runs may instead be due to the fluid role played by worker-employers. In particular, the ability of worker-employers to function either as workers or as employers permits them to crowd out the pure workers or the pure employers, causing them to degenerate into PWFs. In addition, worker-employers have the unique ability to form a self-sufficient network of contractual relationships without the participation of either pure workers or pure employers. Indeed, the persistent contractual networks for the second distance cluster in Table 7(a) are characterized by degeneracies of this type.

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

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

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

Indeed, as indicated in Figures 1 and 3 as well, ex ante capacity constraints play a strong coordinating role in all of the previously reported experimental findings. As suggested by Gode and Sunder [22], when attempting to understand the cause of perceived regularities in market outcomes, it is important to carefully separate the effects of institutional constraints per se from the effects of the cognitive functioning of the agents participating in the market. It will be interesting to determine, in future studies, the extent to which the network patterns determined for the experiments at hand are retained under similar ex ante capacity conditions as the modelling for agent cognition is varied.

5 Concluding Remarks

The hallmark of the ACE approach to socio-economic modelling is a bottom up perspective, in the sense that global regularities are grounded in local agent interactions. The previous sections illustrate how the ACE approach is currently being applied to the study of evolutionary labor markets with adaptive search and worksite behavior.

As developed to date, however, this labor market framework only partially achieves the ACE goal of a bottom up perspective. The commencement of the different modules (matching, worksite interactions, evolution step) is still synchronized across all agents from the top down, and the evolution step is not yet implemented in terms of internalized agent behavioral rules. The advantage of imposing this synchronized dynamic structure with an external evolution step is that it permits some results to be obtained concerning the configuration, stability, uniqueness, and social optimality of the persistent contractual networks that form. The disadvantage is that the networks may not be robust to realistic relaxations of these top-down constraints.

In addition to achieving a more complete bottom-up modelling of a labor market, an enormous amount of work remains to be done to achieve the ultimate goal of calibrating the ACE labor market framework to specific real-world labor market contexts using natural data, survey data, and human-subject laboratory data. For example, signalling among agents (e.g., wage bids and offers) needs to be introduced, and capacity constraints must be endogenized so that they are determined as a function of past activities. Also, the role of government regulations (e.g., minimum wage laws) must be considered. Finally, there is the need to imbed labor markets in a more complete ACE modelling of a decentralized market economy.

Nevertheless, it is hoped that the preliminary results presented in the previous section suggest how an ACE approach could provide modelling foundations permitting the rigorous study of the self-organizing capabilities of decentralized market economies.

References

[1] Aghion, P., and Howitt, P., *Endogenous Growth Theory*, The MIT Press, Cambridge, MA, 1998.

[2] Alchian, A. A., "Uncertainty, Evolution, and Economic Theory," *Journal of Political Economy* 58 (1950), 211–222.

[3] Anderson, P. W., Arrow, K. J., and Pines, D. (eds.), *The Economy as an Evolving Complex System*, Proceedings Volume V, Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, Redwood City, CA, 1988.

[4] Arifovic, J., "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control* 18 (1994), 3–28.

[5] Arthur, W. B., "On Designing Economic Agents that Behave Like Human Agents," *Journal of Evolutionary Economics* 3 (1993), 1–22.

[6] Arthur, W. B., LeBaron, B., Holland, J., Palmer, R., and Taylor, J., "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market," pp. 15–44 in Arthur, W. Brian, Durlauf, S. N., and Lane, D. A. (eds.), *The Economy as an Evolving Complex System, II*, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XXVII, Addison-Wesley, Reading, MA, 1997.

[7] Ashlock, D., Smucker, M. D., Stanley, E. A., and Tesfatsion, L., "Preferential Partner Selection in an Evolutionary Study of Prisoner's Dilemma," *BioSystems* 37 (1996), 99–125.

[8] Axelrod, R., *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Princeton University Press, 1997.

[9] Bell, A. M., "Locally Interdependent Preferences in a General Equilibrium Environment," Working Paper No. 97-W02, Department of Economics and Business Administration, Vanderbilt University, 1997.

[10] Birchenhall, C. R., "Modular Technical Change and Genetic Algorithms," Working Paper No. 9508, University of Manchester, 1995.

[11] Bullard, J., and Duffy, J., "A Model of Learning and Emulation with Artificial Adaptive Agents," Economics Working Paper, University of Pittsburgh, PA, 1994.

[12] Banks, D., and Carley, K., "Metric Inference for Social Networks," *Journal of Classification* 11 (1994), 121–149.

[13] De Vany, A., "The Emergence and Evolution of Self-Organized Coalitions," pages 25–50 in Gilli, M. (ed.), *Computational Economic Systems: Models, Methods, and Econometrics*, Kluwer Scientific Publications, New York, 1996.

[14] Diamond, P., "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90 (1982), 217–229.

[15] Diamond, P., *Search Equilibrium Approach to the Micro Foundations of Macroeconomics*, The MIT Press, Cambridge, 1984.

[16] Duong, D. V., Symbolic Interactionist Modeling: The Coevolution of Symbols and Institutions, pp. 349–354 in Meystel, A. (ed.), *Intelligent Systems: A Semiotic Perspective*, National Institute for Standards and Technology, 1996.

[17] Durlauf, S., Neighborhood Feedbacks, Endogenous Stratification, and Income Inequality," Chapter 18, pp. 505–534, in Barnett, W. A., Gandolfo, G., and Hillinger, C. (eds.), *Dynamic Disequilibrium Modelling*, Cambridge University Press, Cambridge, 1996.

[18] Epstein, J., and Axtell, R., *Growing Artificial Societies: Social Science from the Bottom Up*, The MIT Press, 1996.

[19] Friedman, D., "Evolutionary Games in Economics," *Econometrica* 59 (1991), 637–666.

[20] Fudenberg, D., and Levine, D. K., *Theory of Learning in Games*, The MIT Press, 1998, to appear.

[21] Gale, D. and Shapley, L., "College Admissions and the Stability of Marriage," *American Mathematical Monthly* 69 (1962), 9–15.

[22] Gode, D. K., and Sunder, S., "Allocative Efficiency in Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy* 101 (1993), 119–137.

[23] Goldberg, D., *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, MA, 1989.

[24] Guriev, S., and Shakhova, M., "Self-Organization of Trade Networks in an Economy with Imperfect Infrastructure," in Schweitzer, F. (ed.), *Self-Organization of Complex Structures: From Individual to Collective Dynamics*, Gordon and Breach Scientific Publishers, London, 1996.

[25] Hofbauer, J., and Sigmund, K., *Evolutionary Games and Replicator Dynamics*, Cambridge University Press, 1998.

[26] Holland, J., *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, Second Edition, MIT Press/Bradford Books, Cambridge, MA, 1992.

[27] Holland, J., and Miller, J., "Artificial Adaptive Agents in Economic Theory," *American Economic Review Papers and Proceedings* 81 (1991), 365–370.

[28] Katzner, D., *The Walrasian Vision of the Microeconomy: An Elementary Exposition of the Structure of Modern General Equilibrium Theory*, University of Michigan Press, Ann Arbor, 1989.

[29] Kirman, A., "Whom or What Does the Representative Individual Represent?," *Journal of Economic Perspectives* 6 (Spring 1992), 117–136.

[30] Kirman, A., "Ants, Rationality, and Recruitment," *Quarterly Journal of Economics* 108 (1993), 137-156.

[31] Kirman, A., "The Economy as an Interactive System," pp. 491–531 in Arthur, W. Brian, Durlauf, S. N., and Lane, D. A. (eds.), *The Economy as an Evolving Complex System, II*, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XXVII, Addison-Wesley, Reading, MA.

[32] Kollman, K., Miller, J. H., and Page, S. E., "Political Institutions and Sorting in a Tiebout Model," *American Economic Review*, 1998, to appear.

[33] Lane, D., "Artificial Worlds and Economics," *Journal of Evolutionary Economics* 3 (1993), 89–107.

[34] Leijonhufvud, A., *On Keynesian Economics and the Economics of Keynes*, Oxford University Press, Oxford, 1968.

[35] Levy, S., *Artificial Life: A Report from the Frontier Where Computers Meet Biology*, Pantheon Books, N.Y., 1992.

[36] Lindgren, K., and Nordahl, M. G., "Cooperation and Community Structure in Artificial Ecosystems," *Artificial Life* 1 (1994), 15–37.

[37] Mailath, G., Samuelson, L., and Shaked, A., "Evolution and Endogenous Interactions," SSRI Working Paper 9426, UW-Madison, June, 1994.

[38] Marimon, R., McGrattan, E., and Sargent, T. J., "Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents," *Journal of Economic Dynamics and Control* 14 (1990), 329–373.

[39] Marks, R. E., "Breeding Hybrid Strategies: Optimal Behavior for Oligopolists," *Journal of Evolutionary Economics* 2 (1992), 17–38.

[40] McFadzean, D., *SimBioSys: A class framework for evolutionary simulations*, Master's Thesis, Department of Computer Science, University of Calgary, Alberta, Canada, 1995. [The source code for SimBioSys is available for downloading as freeware at the current author's Web site.]

[41] McFadzean, D., and L. Tesfatsion, "A C++ Platform for the Evolution of Trade Networks," Economic Report No. 39, Iowa State University, revised September 1997, to appear in *Computational Economics*. [The source code for this trade network game (TNG) framework is available for downloading as freeware, along with instructions, at the current author's Web site.]

[42] Miller, J. H., "The Coevolution of Automata in the Repeated Prisoner's Dilemma," *Journal of Economic Behavior and Organization* 29:1 (January 1996), 87–112.

[43] Mitchell, M., and Forrest, S., "Genetic Algorithms and Artificial Life," *Artificial Life* 1 (1994), 267–289.

[44] Nelson, R., "Recent Evolutionary Theorizing About Economic Change," *Journal of Economic Literature* 33 (1995), 48–90.

[45] Nelson, R. R., and Winter, S. G., *An Evolutionary theory of Economic Change*, Harvard University Press, Cambridge, MA, 1982.

[46] Orbell, J. M., and Dawes, R. M., "Social Welfare, Cooperators' Advantage, and the Option of Not Playing the Game," *American Sociological Review* 58 (1993), 787–800.

[47] Richards, R. J., *Darwin and the Emergence of Evolutionary Theories of Mind and Behavior*, University of Chicago Press, Chicago, 1987.

[48] Roth, A., and Sotomayor, M. A. O., *Two-sided matching: A study in game-theoretic modeling and analysis*, Cambridge University Press, Cambridge, 1990.

[49] Routledge, B., "Artificial Selection: Genetic Algorithms and Learning in a Rational Expectations Model," Working Paper, Faculty of Commerce and Business Administration, UBC, 1994.

[50] Rust, J., Miller, J., and Palmer, R., "Behavior of Trading Automata in a Computerized Double Auction Market," pp. 155–198 in D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, Addison-Wesley, New York, 1992.

[51] Samuelson, L., *Evolutionary Games and Equilibrium Selection*, The MIT Press, Cambridge, MA, 1997.

[52] Sargent, T., *Bounded Rationality in Macroeconomics*, Oxford University Press, Clarendon, 1993.

[53] Schumpeter, J. A., *Capitalism, Socialism, and Democracy*, Harper and Brothers, N.Y., 1942.

[54] Shubik, M., *A Game-Theoretic Approach to Political Economy*, The MIT Press, Cambridge, MA, 1991.

[55] Sigmund, K., *Games of Life: Explorations in Ecology, Evolution, and Behavior*, Oxford University Press, Oxford, 1993.

[56] Stanley, E. A., Ashlock, D., and Tesfatsion, L., "Iterated Prisoner's Dilemma with Choice and Refusal of Partners," pp. 131–175 in Langton, C. (ed.), *Artificial Life III*, Proceedings Volume 17, Santa Fe Institute Studies in the Sciences of Complexity, Addison Wesley, 1994.

[57] Tesfatsion, L., "A Trade Network Game with Endogenous Partner Selection," Economic Report No. 36, Iowa State University, April 1995.

[58] Tesfatsion, L., "How Economists Can Get Alife," pp. 533–564 in Arthur, W. Brian, Durlauf, S., and Lane, D. (eds.), *The Economy as an Evolving Complex System, II*, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XXVII, Addison-Wesley, 1997.

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

[60] Tesfatsion, L., "Preferential Partner Selection in Evolutionary Labor Markets: A Study in Agent-Based Computational Economics," pp. 15–24 in Porto, V. W., Saravanan, N., Waagen, D., and Eiben, A. E. (eds), *Evolutionary Programming VII*, Proceedings of the Seventh Annual Conference on Evolutionary Programming, Springer-Verlag, Berlin, 1998.

[61] Tesfatsion, L., "Ex Ante Capacity Effects in an Evolutionary Labor Market with Endogenous Search," Economic Report No. 48, Iowa State University, November 1998.

[62] Vriend, N., "Self-Organization of Markets: An Example of a Computational Approach," *Computational Economics* 8 (1995), 205–231.

[63] Witt, U., *Evolutionary Economics*, Edward Elgar, London, 1993.

[64] Young, P., "The Evolution of Conventions," *Econometrica* 61 (1993), 57–84.

[65] Young, P., *Individual and Social Structure: An Evolutionary Theory of Institutions*, Princeton University Press, Princeton, 1998.

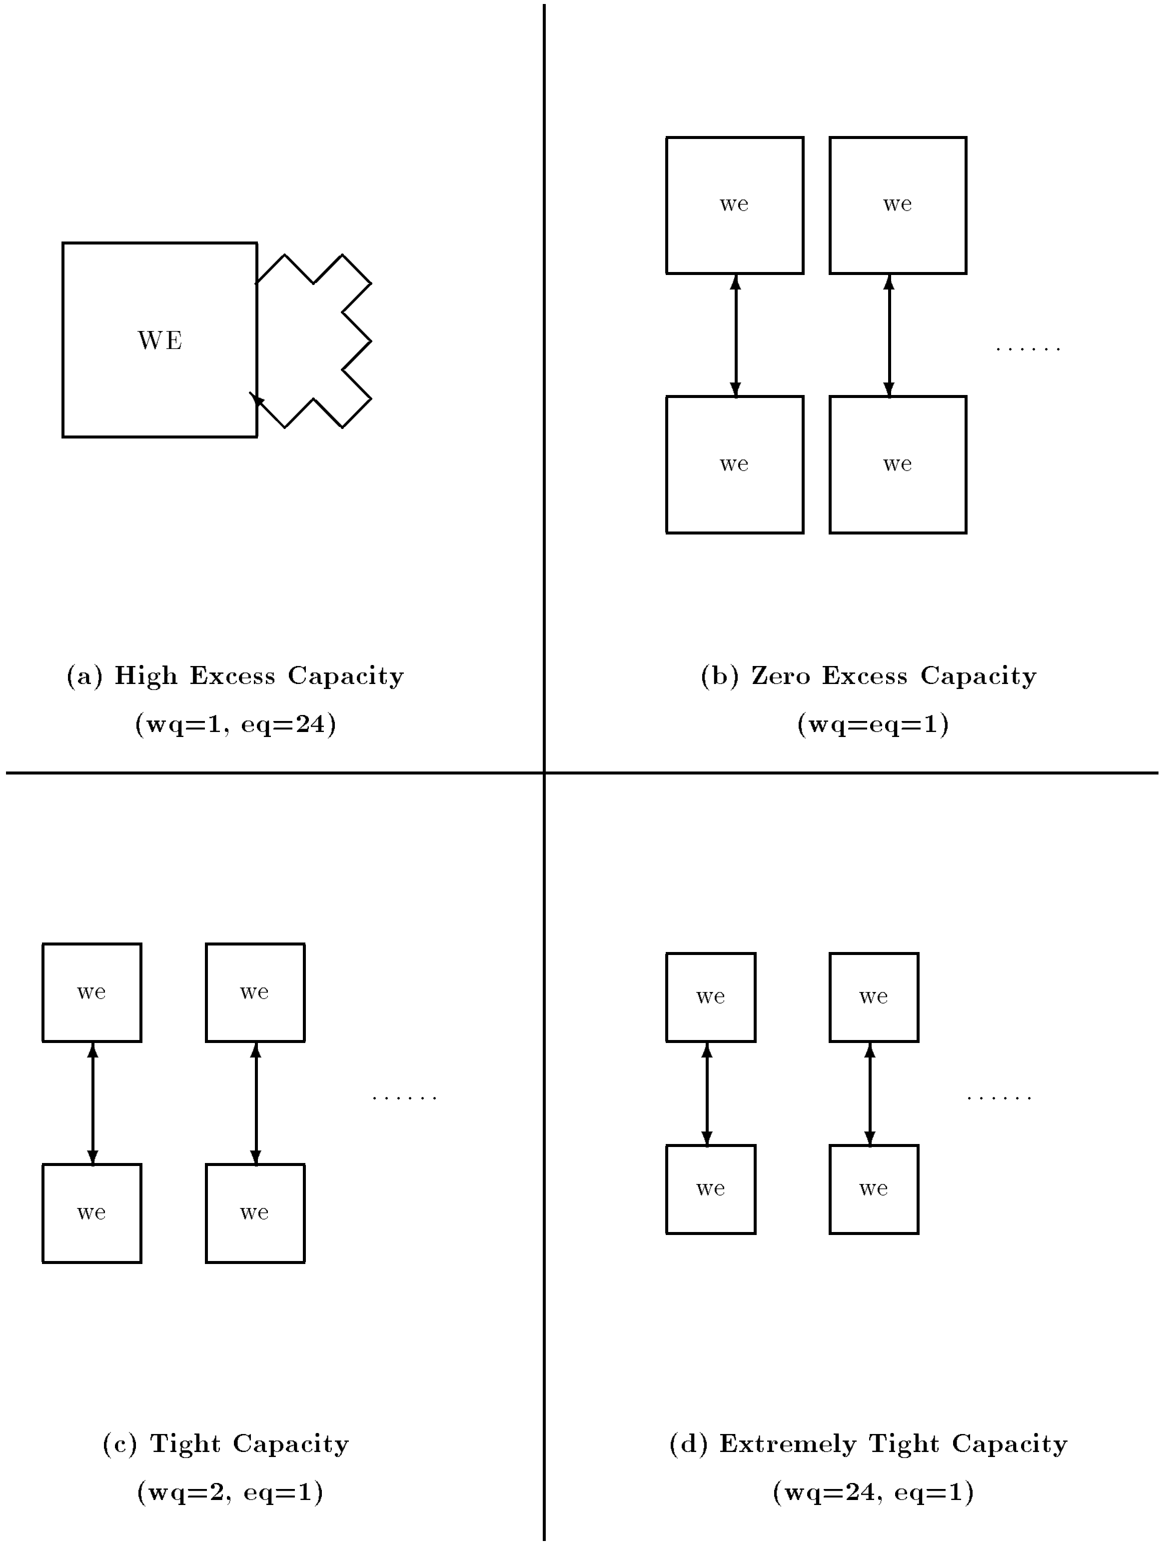

Figure 1: Base Contractual Patterns for Endogenous-Type Labor Markets with Different Ex Ante Capacities. A relatively larger box for the worker-employers (WE) under a particular capacity specification indicates that the worker-employers achieve a relatively higher mean FIT value under this capacity specification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

THIS FIGURE IS AVAILABLE FROM THE AUTHOR UPON REQUEST

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

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

Table 5(a): Endogenous-Type Labor Markets with High Excess Capacity

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

Table 5(b): Endogenous-Type Labor Markets with Zero Excess Capacity

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

Table 5(c): Endogenous-Type Labor Markets with Tight Capacity

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

Table 5(d): Endogenous-Type Labor Markets with Extremely Tight Capacity

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

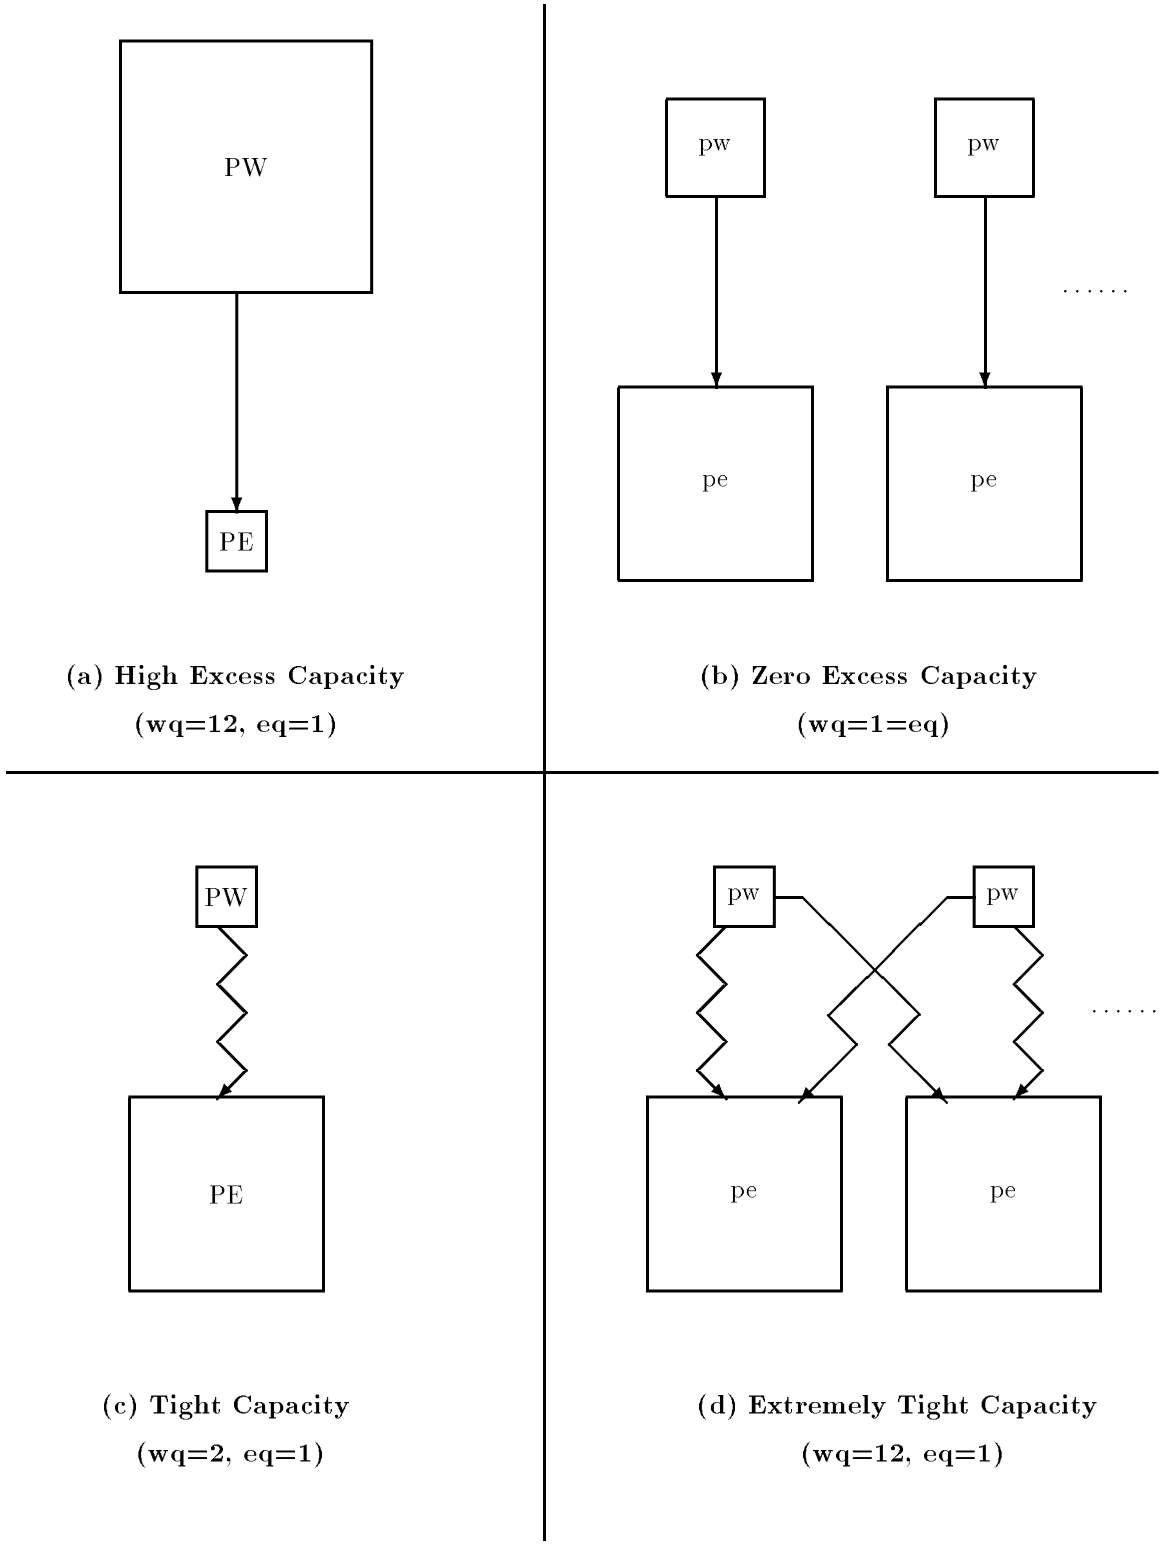

Figure 3: Base Contractual Patterns for Two-Sided Labor Markets with Different Ex Ante Capacities. A relatively larger box for an agent type — pure workers (PW) or pure employers (PE) — under a particular capacity specification indicates that this agent type achieves a relatively higher mean FIT value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

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

Table 6(a): Two-Sided Labor Markets with High Excess Capacity

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

Table 6(b): Two-Sided Labor Markets with Zero Excess Capacity

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

Table 6(c): Two-Sided Labor Markets with Tight Capacity

D^o Cluster	% Runs	Mean UD		Mean PWF		Mean PC		Mean FIT	
		pw	pe	pw	pe	pw	pe	pw	pe
0–6	35%	1%	1%	12%	1%	86%	96%	0.31	1.37
		(3%)	(3%)	(4%)	(3%)	(7%)	(6%)	(.03)	(.06)
15–17	20%	10%	92%	35%	2%	17%	25%	0.35	1.22
		(14%)	(14%)	(7%)	(4%)	(20%)	(34%)	(.17)	(.20)
24	45%	100%	100%	100%	100%	0%	0%	-0.10	-0.01
		(0%)	(0%)	(0%)	(0%)	(0%)	(0%)	(.00)	(.00)

Table 6(d): Two-Sided Labor Markets with Extremely Tight Capacity

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

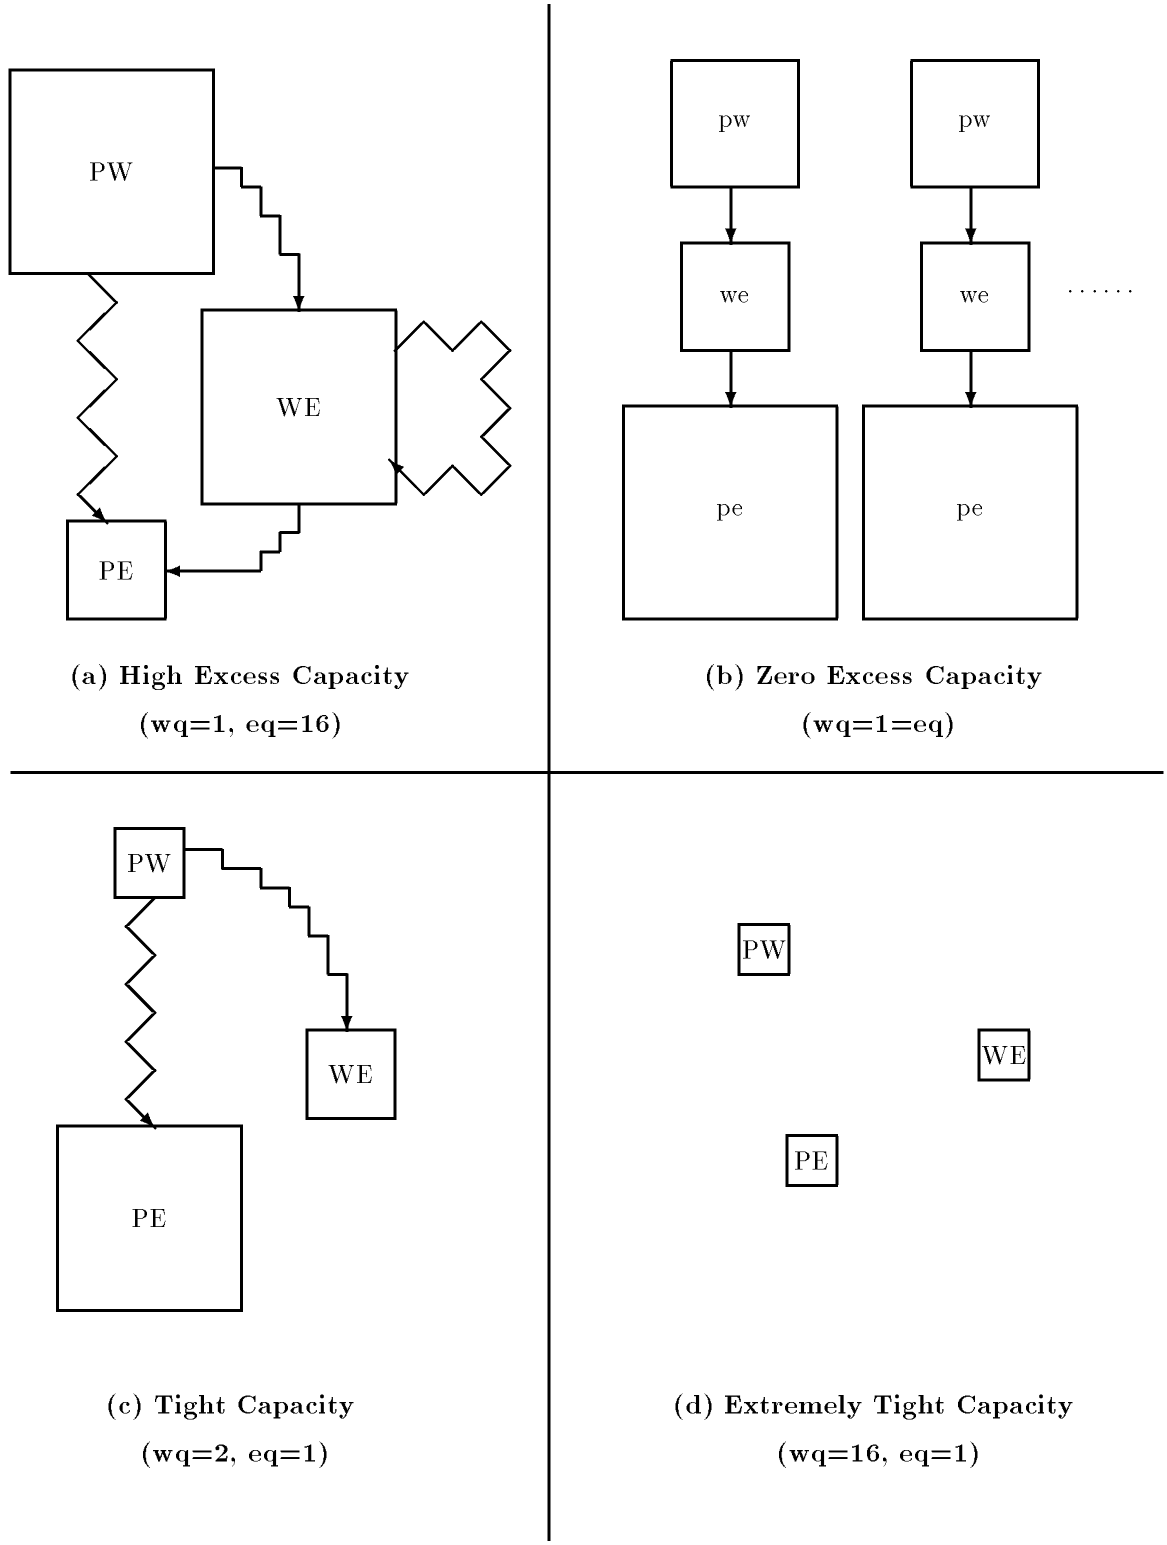

Figure 4: Base Contractual Patterns for Partially-Fluid Labor Markets with Different Ex Ante Capacities. A relatively larger box for an agent type — pure workers (PW), worker-employers (WE), or pure employers (PE) — under a particular capacity specification indicates that this agent type achieves a relatively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

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

Table 7(a): Partially-Fluid Labor Markets with High Excess Capacity

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

Table 7(b): Partially-Fluid Labor Markets with Zero Excess Capacity

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

Table 7(c): Partially-Fluid Labor Markets with Tight Capacity

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

Table 7(d): Partially-Fluid Market Experiments with Extremely Tight Capacity

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