Original Version: September 1995
Revised Version: March 1997
References Updated: October 2000

HOW ECONOMISTS CAN GET ALIFE[1]

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

Abstract: This paper presents a summary overview of the fast-developing field of artificial life, stressing aspects especially relevant for the study of decentralized market economies. In particular, a recently developed trade network game (TNG) is used to illustrate how the basic artificial life paradigm might be specialized to economics. The TNG traders choose and refuse trade partners on the basis of continually updated expected utility, engage in risky trades modelled as two-person games, and evolve their trade behavior over time. Analytical and simulation work is reported to indicate how the TNG is currently being used to study the evolutionary implications of alternative market structures at three different levels: individual trade behavior; trade network formation; and social welfare.

1 Introduction

What is artificial life, or alife for short? And why should economists care?

As detailed in the entertaining monographs by Levy (1992) and Sigmund (1993), the roots of alife go at least as far back as the work of John von Neumann in the nineteen forties on self-replicating automata. The establishment of alife as a distinct field of inquiry, however, must be traced to the first alife conference, organized in 1987 by Chris Langton at the Los Alamos National Laboratory; see Langton (1989).

Alife is the bottom-up study of basic phenomena commonly associated with living organisms, such as self-replication, evolution, adaptation, self-organization, parasitism, competition, cooperation, and social network formation. Alife complements the traditional biological and social sciences concerned with the analytical, laboratory, and field study of living organisms by attempting to simulate or synthesize life-like behavior within computers, robots, and other man-made media. One goal is to enhance the understanding of actual and potential life processes. A second goal is to use nature as an inspiration for the development of solution algorithms for difficult optimization problems characterized by high-dimensional search domains, nonlinearities, and multiple local optima.

The systems studied by alife researchers are complex adaptive systems sharing many of the following characteristics [Holland, 1992]. Most importantly, each such system typically consists of many dispersed

[1]An abbreviated version of this survey appears in Tesfatsion (1995). Links to additional related materials can be found at the author's Web site. Thanks to A. De Vany, J. Duffy, D. Fogel, J. Gray, R. Noll, B. Routledge, and especially N. Vriend for helpful comments.

units acting in parallel with no global controller responsible for the behavior of all units. Rather, the actions of each unit depend upon the states and actions of a limited number of other units, and the overall direction of the system is determined by competition and coordination among the units subject to structural constraints. The complexity of the system thus tends to arise more from the interactions among the units than from any complexity inherent in the individual units per se. Moreover, the local interaction networks connecting individual units are continuously recombined and revised. In particular, niches that can be exploited by particular adaptations are continuously created, and their exploitation in turn leads to new niche creations, so that perpetual novelty exists.

Briefly put, then, alife research tends to focus on continually evolving systems whose global behavior arises from the local interactions of distributed units; this is the sense in which alife research is said to be bottom up. Although the units comprising the systems might be bit strings, molecules, or robotic insects, the abstract description of how the unit interactions result in global behavior is clearly reminiscent of a Schumpeterian economy, only filtered through an unfamiliar terminology.

The study of evolutionary economies has of course been pursued by many researchers in addition to Joseph Schumpeter. For example, one has Armen Alchian's work on uncertainty and evolution in economic systems, 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 Richard Nelson and Sidney Winter's work on an evolutionary theory of economic change. These and numerous other related studies are reviewed by Witt (1993) and Nelson (1995). In addition, as detailed in Friedman (1991), a number of researchers have recently been focusing on the potential economic applicability of evolutionary game theory in which game strategies distributed over a fixed number of strategy types reproduce over time in direct proportion to their relative fitness.

Economists have recently begun to apply the alife paradigm to the computational study of evolutionary economic processes. Exploiting the recent advent of object-oriented programming languages such as C++ and Java, these "agent-based computational economics" (ACE) researchers have been able to extend previous evolutionary economics work in several directions.[2] First, much greater attention is generally focused on the endogenous determination of agent interactions. Second, a broader range of interactions is typically considered, with cooperative and predatory associations increasingly taking center stage along with price and quantity relationships. Third, agent actions and interactions are represented with a greater degree of abstraction, permitting generalizations across specific system applications. Fourth, the evolutionary process is generally expressed algorithmically in terms of genetic (recombination and/or mutation) operations acting directly on agent characteristics. These evolutionary selection pressures result in the continual creation of new modes of behavior and an ever-changing network of agent interactions.

For example, the basic "genetic algorithm" used in many ACE studies evolves a new population of agents from an existing population of agents using the following four steps: (1) *Evaluation*, in which a fitness score is assigned to each agent in the population; (2) *Selection for Reproduction*, in which a subset of the existing population of agents is selected for reproduction, with selection biased in favor of fitness; (3) *Recombination*,

[2] More detailed information about ACE, including surveys, annotated readings, software, and pointers to individual researchers and research groups, can be found at the ACE Web site at http://www.econ.iastate.edu/tesfatsi/ace.htm.

in which offspring (new ideas) are generated by combining the genetic material (structural characteristics) of pairs of parents chosen from among the most fit agents in the population; and (4) *Mutation*, in which additional variations are introduced into the population by mutating the structural characteristics of each offspring with some small probability. See Goldberg (1989) and Mitchell and Forrest (1994).

The central problem for ACE researchers is to understand the apparently spontaneous appearance of regularity 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 from the local interactions of autonomous agents channeled through actual or potential economic institutions rather than through fictitious coordinating mechanisms such as a single representative consumer. In line with this challenge, rationality is generally viewed as a testable hypothesis, or at least as a debatable methodological assumption, rather than as an unquestioned axiom of individual behavior.

Several studies that focus on key ACE-related issues have either appeared or are in the pipeline. See, for example, Anderson et al. (1988), Arifovic (1994), Arthur (1993), Arthur et al. (1997), Bell (1997), Birchenhall (1995), Bosch and Sunder (1996), Bullard and Duffy (1994), De Vany (1996), Durlauf (1996), Epstein and Axtell (1996), Holland and Miller (1991), Kirman (1993;1997), Lane (1993), Mailath et al. (1994), Marimon et al. (1990), Marks (1992), McFadzean and Tesfatsion (1997,1999), Miller (1989), Routledge (1994), Sargent (1993), Tesfatsion (1997), and Vriend (1995).

To illustrate more concretely the potential usefulness of the ACE approach, as well as the hurdles that remain to be cleared, the following two sections briefly outline some ongoing ACE work that appears to be particularly relevant for the modelling of decentralized market economies. Section 2 describes recent attempts to combine evolutionary game theory with preferential partner selection [Stanley et al. (1994); Smucker et al. (1994); and Ashlock et al. (1996)]. Section 3 discusses how a modified version of this framework is being used to study the endogenous formation and evolution of trade networks [Tesfatsion (1997), McFadzean and Tesfatsion (1997,1999)]. Concluding comments are given in Section 4.

2 Evolutionary IPD with Choice and Refusal

Following the seminal work of Axelrod (1984, 1987), the iterated prisoner's dilemma (IPD) game has been extensively used by economists and other researchers to explore the potential emergence of mutually cooperative behavior among non-altruistic agents. As detailed in Kirman (1997) and Lindgren and Nordahl (1994), these studies have typically assumed that individual players have no control over whom they play. Rather, game partners are generally determined by an extraneous matching mechanism such as a roulette wheel, a neighborhood grid, or a round-robin tournament . The general conclusion reached by these studies has been that mutually cooperative behavior tends to emerge if the number of game iterations is either unknown or infinite, the frequency of mutually cooperative play in initial game iterations is sufficiently large, and the perceived probability of future interactions with any given current partner is sufficiently high.

In actuality, however, socio-economic interactions are often characterized by the preferential choice and refusal of partners. The question then arises whether the emergence and long-run viability of cooperative behavior in the IPD game would be enhanced if players were more realistically allowed to choose and refuse their potential game partners.

This question is taken up in Stanley et al. (1994). The traditional IPD game is extended to an IPD/CR game in which players choose and refuse partners on the basis of continually updated expected payoffs.[3]

The introduction of partner choice and refusal fundamentally modifies the ways in which players interact in the IPD game and the characteristics that result in high payoff scores. Choice allows players to increase their chances of encountering other cooperative players, refusal gives players a way to protect themselves from defections without having to defect themselves, and ostracism of defectors occurs endogenously as an increasing number of players individually refuse the defectors' game offers. On the other hand, choice and refusal also permit opportunistic players to home in quickly on exploitable players and form parasitic relationships.

The analytical and simulation findings reported for the IPD/CR game in Stanley et al. (1994), and in the subsequent studies by Smucker et al. (1994), Ashlock et al. (1996), and Hauk (1996), indicate that the overall emergence of cooperation is accelerated in evolutionary IPD games by the introduction of choice and refusal. Nevertheless, the underlying player interaction patterns induced by choice and refusal can be complex and time varying, even when expressed play behavior is largely cooperative. Consequently, it has proven to be extremely difficult to get an analytical handle on the mapping from parameter configurations to evolutionary IPD/CR outcomes.

A reasonable next step, then, is to focus on more concrete problem settings which impose natural constraints on the range of feasible player interactions. In the next section it is shown how a modified version of the IPD/CR game is being used to examine the endogenous formation and evolution of trade networks among resource-constrained traders.

3 A Trade Network Game with Choice and Refusal

The *trade network game* (TNG) developed in Tesfatsion (1997) consists of successive generations of resource-constrained traders who choose and refuse trade partners on the basis of continually updated expected payoffs, engage in risky trades modelled as two person games, and evolve their trade strategies over time. The TNG has been implemented by McFadzean and Tesfatsion (1997,1999) with the support of a general C++ evolutionary simulation framework, *SimBioSys*, developed by McFadzean (1995).

The TNG facilitates the general study of trade from a bottom up perspective in three key ways. First, the TNG traders are instantiated as autonomous endogenously interacting software agents (*tradebots*) with internal behavioral functions and with internally stored information that includes addresses for other tradebots. The tradebots can therefore display anticipatory behavior (expectation formation); and they can communicate with each other at event-triggered times, a feature not present in standard economic models. Second, the modular design of the TNG permits experimentation with alternative specifications for market structure, trade partner matching, trading, expectation formation, and trade behavior evolution. All of these specifications can potentially be grounded in tradebot-initiated actions. Third, the evolutionary

[3] Other game theory studies that have allowed players to avoid unwanted interactions, or more generally to affect the probability of interaction with other players through their own actions, include Fogel (1995), Guriev and Shakhova (1996), Hauk (1996), Hirshleifer and Rasmusen (1989), Kitcher (1993), Mailath et al. (1994), and Orbell and Dawes (1993). A detailed review of this line of work is given in Hauk (1996). There is also a growing body of work on multi-agent systems with endogenous interactions in which the decision (or state) of an agent depends on the decision (or state) of certain neighboring agents, where these neighbors may change over time. See, for example, Brock and Durlauf (1995), De Vany (1996), Ioannides (1997), and Young (1993).

implications of alternative module specifications can be studied at three different levels: individual tradebot characteristics; trade network formation; and social welfare as measured by descriptive statistics such as average tradebot fitness.

Section 3.1 sets out and motivates the general TNG framework. To gain insight into the subtle interplay between game play and the choice and refusal of partners in the TNG, Section 3.2 presents a detailed analytical study of an illustrative TNG with five tradebots. In particular, it is shown that the parameter space for this illustrative TNG partitions into economically interpretable regions corresponding to qualitatively distinct trade network formations. Section 3.3 reports on some illustrative TNG computer experiments for two alternative market structures: buyer-seller markets in which all tradebots can both make and receive trade offers; and two-sided markets in which a subset of buyers makes trade offers to a disjoint subset of sellers.

3.1 The Basic Trade Network Game

The trade network game (TNG) consists of a collection of tradebots that evolves over time. As depicted in Table 1, this evolution is implemented through a hierarchy of cycle loops.

Each tradebot in the initial tradebot generation is assigned a random trade strategy and is configured with a prior expected payoff for each of his potential trade partners. The tradebots then engage in a *trade cycle loop* consisting of a fixed number of trade cycles. In each trade cycle the tradebots undertake three activities: the determination of trade partners, given current expected payoffs; the carrying out of potentially risky trades; and the updating of expected payoffs based on any new payoffs received during trade partner determination and trading. At the end of the trade cycle loop the tradebots enter into an *environmental cycle* during which the *fitness score* of each tradebot is calculated as the total sum of his payoffs divided by the total number of his payoffs and the current tradebot generation is sorted by fitness scores. At the end of the environmental cycle, a *generation cycle* commences during which evolutionary selection pressures are applied to the current tradebot generation to obtain a new tradebot generation with evolved trade strategies. This new tradebot generation is then configured, and another trade cycle loop commences.

The TNG currently uses the particular specifications for market structure, trade partner determination, trade, expectation updating, and trade behavior evolution detailed in Tesfatsion (1997). For completeness, these specifications are reviewed below.

Alternative market structures are currently imposed in the TNG through the prespecification of buyers and sellers and through the prespecification of quotas on offer submissions and acceptances. More precisely, the set of players for the TNG is the union $V = B \cup S$ of a nonempty subset B of *buyer tradebots* who can submit trade offers and a nonempty subset S of *seller tradebots* who can receive trade offers, where B and S may be disjoint, overlapping, or coincident. In each trade cycle, each buyer m can submit up to O_m trade offers to sellers and each seller n can accept up to A_n trade offers from buyers, where the offer quota O_m and the acceptance quota A_n can be any positive integers.

Although highly simplified, these parametric specifications permit the TNG to encompass two-sided markets, markets with intermediaries, and markets in which all traders engage in both buying and selling activities. For example, the buyers and sellers might represent customers and retail store owners, workers and firms, borrowers and lenders, or barter traders. The offer quota O_m indicates that buyer m has a limited amount of resources (credit, labor time, collateral, apples,...) to offer, and the acceptance quota A_n indicates

```
int main () {
  Init() ;                           // Construct initial tradebot generation
                                     //    with random trade strategies.
  For (G = 0,...,GMAX-1) {           // Enter the generation cycle loop.
    If (G > 0) {
      EvolveGen();                   // Generation cycle: Evolve a
                                     //    new tradebot generation.
    }
    InitGen();                       // Configure tradebots with user-supplied
                                     //    parameter values (prior expected
                                     //    payoffs, quotas,...).
    For (I = 0,...,IMAX-1) {         // Enter the trade cycle loop.
      MatchTraders();                // Determine trade partners,
                                     //    given expected payoffs,
                                     //    and record refusal and
                                     //    wallflower payoffs.
      Trade();                       // Implement trades and
                                     //    record trade payoffs.
      UpdateExp();                   // Update expected payoffs
                                     //    using newly recorded payoffs.
    }
    AssessFitness();                 // Environmental Cycle: Record
                                     //    fitness scores.
  }
  Return 0 ;
}
```

Table 1: Pseudo-Code for the TNG

that seller n has a limited amount of resources (goods, job openings, loans, oranges,...) to provide in return.

Three illustrations are sketched below.

Case 1: A Two-Sided Labor Market With Endogenous Layoffs and Quits[4]

The set B consists of M workers and the set S consists of N employers, where B and S are disjoint. Each worker m can make work offers to a maximum of O_m employers, or he can choose to be unemployed. Each employer n can hire up to A_n workers, and employers can refuse work offers. Once matched, workers choose on-the-job effort levels and employers choose monitoring and penalty levels. An employer fires one of its current workers by refusing future work offers from this worker, and a worker quits his current employer by ceasing to direct work offers to this employer. This TNG special case thus extends the standard treatment of labor markets as assignment problems by incorporating subsequent strategic efficiency wage interactions between matched pairs of workers and employers and by having these interactions iterated over time.

Case 2: Intermediation with Choice and Refusal

The buyer subset B and the seller subset S overlap but do not coincide. The pure buyers in $V-S$ are the depositors (lenders), the buyer-sellers in $B \cap S$ are the intermediaries (banks), and the pure sellers in $V-B$ are the capital investors (borrowers). The depositors offer funds to the intermediaries in return for deposit accounts, and the intermediaries offer loan contracts to the capital investors in return for a share of earnings. The degree to which an accepted offer results in satisfactory payoffs for the participants is determined by the degree to which the deposit or loan contract obligations are fulfilled.

Case 3: A Labor Market with Endogenously Determined Workers and Employers

The subsets B and S coincide, implying that each tradebot can both make and receive trade offers. Each tradebot v can make up to O_v work offers to tradebots at other work sites and receive up to A_v work offers

[4]Systematic experiments for two-sided labor markets with endogenous layoffs and quits are reported in Tesfatsion (1999a,b).

at his own work site. As in Case 1, the degree to which any accepted work offer results in satisfactory payoffs for the participant tradebots is determined by subsequent work site interactions. Ex post, four pure types of tradebots can emerge: (1) pure workers, who work at the sites of other tradebots but have no tradebots working for them at their own sites; (2) pure employers, who have tradebots working for them at their own sites but who do not work at the sites of other tradebots; (3) unemployed tradebots, who make at least one work offer to a tradebot at another site but who end up neither working at other sites nor having tradebots working for them at their own sites; and (4) inactive (out of the work force) tradebots, who neither make nor accept any work offers.

The determination of trade partners in the TNG is currently implemented using a modified version of the well-known Gale-Shapley (1962) deferred acceptance mechanism. This modified mechanism, hereafter referred to as the *deferred choice and refusal* (DCR) mechanism, presumes that each buyer and seller currently associates an expected payoff with each potential trade partner. Also, each buyer and seller is presumed to have an exogenously given *minimum tolerance level*, in the sense that he will not trade with anyone whose expected payoff lies below this level.

The DCR mechanism proceeds as follows. Each buyer m first makes trade offers to a maximum of O_m most-preferred sellers he finds tolerable, with at most one offer going to any one seller. Each seller n in turn forms a waiting list consisting of a maximum of A_n of the most preferred trade offers he has received to date from tolerable buyers; all other trade offers are refused. For both buyers and sellers, selection among equally preferred options is settled by a random draw. A buyer that has a trade offer refused receives a nonpositive *refusal payoff*, R; the seller who does the refusing is not penalized. A refused buyer immediately submits a replacement trade offer to any tolerable next-most-preferred seller that has not yet refused him. A seller receiving a new trade offer that dominates a trade offer currently on his waiting list substitutes this new trade offer in place of the dominated trade offer, which is then refused. A buyer ceases making trade offers when either he has no further trade offers refused or all tolerable sellers have refused him. When all trade offers cease, each seller accepts all trade offers currently on his waiting list. A tradebot that neither submits nor accepts trade offers during this matching process receives a *wallflower payoff*, W.[5]

The buyer-seller matching outcomes generated by the DCR mechanism exhibit the usual static optimality properties associated with Gale-Shapley type matching mechanisms. First, any such matching outcome is core stable, in the sense that no subset of tradebots has an incentive to block the matching outcome by engaging in a feasible rearrangement of trade partners among themselves [Tesfatsion, 1997, Proposition 3.2]. Second, define a matching outcome to be B-optimal if it is core stable and if each buyer matched under the matching outcome is at least as well off as he would be under any other core stable matching outcome. Then, in each TNG trade cycle, the DCR mechanism yields the unique B-optimal matching outcome as long as each tradebot has a strict preference order over the potential trade partners he finds tolerable [Tesfatsion, 1997, Proposition 3.3]. As indicated by the computer experiments reported in Section 3.3 below, however, these static optimality properties do not appear to be adequate measures of optimality from an evolutionary perspective.

Trades are currently modelled in the TNG as prisoner's dilemma (PD) games. For example, a trade

[5]The DCR mechanism always stops in finitely many steps [Tesfatsion, 1997, Proposition 3.1]. A particularly interesting aspect of the DCR mechanism is that it requires the tradebots to pass messages back and forth to each other at event-triggered times, a requirement that is easily handled by a C++ implementation; see McFadzean and Tesfatsion (1997,1999).

		Player 2 c	Player 2 d
Player 1	c	(C,C)	(L,H)
	d	(H,L)	(D,D)

Table 2: Payoff Matrix for the Prisoner's Dilemma Game

may involve the exchange of a good or service of a certain promised quality in return for a loan or wage contract entailing various payment obligations. A buyer participating in a trade may either cooperate (fulfill his trade obligations) or defect (renege on his trade obligations), and similarly for a seller. The range of possible payoffs is the same for each trade in each trade cycle: namely, L (the sucker payoff) is the lowest possible payoff, received by a cooperative tradebot whose trade partner defects; D is the payoff received by a defecting tradebot whose trade partner also defects; C is the payoff received by a cooperative tradebot whose trade partner also cooperates; and H (the temptation payoff) is the highest possible payoff, received by a defecting tradebot whose trade partner cooperates. More precisely, the payoffs are assumed to satisfy $L < D < 0 < C < H$, with $(L + H)/2 < C$. The payoff matrix for the PD game is depicted in Table 2.

The TNG tradebots are currently assumed to use a simple form of criterion filter[6] to update their expected payoffs on the basis of new payoff information. Specifically, whenever a tradebot v receives a trade or refusal payoff P from an interaction with a potential trade partner k, tradebot v forms an updated expected payoff for k by taking a convex combination of this new payoff P and his previous expected payoff for k. The inverse of the weight on the new payoff P is 1 plus v's current payoff count with k. As explained in Tesfatsion (1997), this updating procedure guarantees that the expected payoff tradebot v associates with k converges to the true average payoff v attains from interactions with k as the number of interactions between v and k becomes arbitrarily large.

The trade behavior of each tradebot, whether he is a pure buyer in $V - S$, a buyer-seller in $B \cap S$, or a pure seller in $V - B$, is currently characterized by a finite-memory pure strategy for playing a PD game with an arbitrary partner an indefinite number of times, hereafter referred to as a *trade strategy*. Each tradebot thus has a distinct trading personality even if he engages in both buying and selling activities. At the commencement of each trade cycle loop, tradebots have no information about the trade strategies of other tradebots; they can only learn about these strategies by engaging other tradebots in repeated trades and observing the payoff histories that ensue. Moreover, each tradebot's choice of an action in a current trade with a potential trade partner is determined entirely on the basis of the payoffs obtained in past trades with this same partner. Thus, each tradebot keeps separate track of the particular state he is in with regard to each of his potential trade partners.

In the current implementation of the TNG, the only aspect of a tradebot that evolves over time is his trade strategy. The evolution of the tradebots in each generation cycle is thus meant to reflect the formation and transmission of new ideas rather than biological reproduction.

[6] As detailed in Tesfatsion (1979), a criterion filter is an algorithm for the direct updating of an expected return function on the basis of past return outcomes without recourse to the usual interim updating of probabilities via Bayes' rule.

More precisely, each tradebot's trade strategy is represented as a finite state machine (FSM) with a fixed starting state. The FSMs for two illustrative trade strategies are depicted in Figure 1: namely, a nice trade strategy, Tit-for-Two-Tats, that only defects if defected against twice in a row; and an opportunistic trade strategy, Rip-Off, that evolved in an experiment with an initial population of Tit-for-Two-Tats to take perfect advantage of the latter strategy by defecting every other time.[7] As is more carefully explained in McFadzean and Tesfatsion (1997,1999), in each generation cycle the trade strategies (FSMs) associated with the current tradebot generation are evolved by means of a standardly specified genetic algorithm involving mutation, recombination, and elitism operations applied to bit string encodings for the strategies. The effect of these operations is that successful trade strategies are mimicked and unsuccessful trade strategies are replaced by variants of more successful strategies.

—INSERT FIGURE 1 ABOUT HERE—

3.2 An Illustrative 5-Tradebot TNG

Consider a TNG for which the player set contains a total of five buyer-seller tradebots who can both make and receive trade offers.[8] Each tradebot v has the same minimum tolerance level, 0. Also, each tradebot v has the same offer quota, $O_v = 1$, implying that he can have at most one trade offer outstanding at any given time and can receive at most four trade offers from other tradebots at any given time. Each tradebot v is also assumed to have the same acceptance quota, $A_v = 4$, which then implies that each tradebot is effectively unconstrained with regard to the number of trade offers he can have on his waiting list at any given time. The refusal payoff, R, is assumed to be strictly negative, and the wallflower payoff, W, is assumed to be 0.

With regard to trade strategies, three of the tradebots are Tit-for-Two-Tats (TFTTs) and the remaining two tradebots are Rip-Offs (Rips); see Figure 1. A TFTT receives a payoff sequence $(C, C, C, ...)$ in repeated trades with another TFTT and a payoff sequence $(L, C, L, C, ...)$ in repeated trades with a Rip, and a Rip receives a payoff sequence $(H, C, H, C, ...)$ in repeated trades with a TFTT and a payoff sequence $(D, C, C, ...)$ in repeated trades with another Rip. Note that a Rip never triggers defection in a TFTT since a Rip never defects twice in a row. Consequently, in any match-up between a TFTT and a Rip, the Rip would definitely attain a higher fitness score than the TFTT if the TFTT were not permitted to refuse the Rip's trade offers.

One key factor affecting the emergence of trade networks in the TNG is the specification of the tradebots' prior expected payoffs. Low prior expected payoffs encourage tradebots to latch on to the first trade partner from whom they receive even a modestly high payoff. On the other hand, high prior expected payoffs encourage repeated experimentation with new trade partners in the face of continually disappointing payoffs from current trade partners. As will be seen below for the TFTTs, the combination of high prior expected payoffs and wide-spread experimentation can be detrimental to nice tradebots since it increases their chances of encountering opportunistically defecting trade partners whose defections are infrequent enough to avoid triggering refusals.

Two alternative benchmark assumptions will be considered for the tradebots' prior expected payoffs for their potential trade partners. The first assumption, a common prior, is the assumption made in Stanley et

[7]The experimental discovery of Rip-Off was made by Daniel Ashlock of Iowa State University during the preparation of Stanley et al. (1994).

[8]The specifications for this illustrative TNG were chosen to permit comparisons with a 5-player IPD/CR game analyzed in Stanley et al. (1994, pp. 153–156).

al. (1994), Smucker et al. (1994), and Ashlock et al. (1996) for the IPD/CR game. The second assumption, long-run expectations, sets the prior expected payoffs that any two tradebots have for each other equal to the true long-run average payoffs that would result if these two tradebots were to engage in infinitely many trades.

Assumption (CP): Common Prior. Each tradebot associates the same prior expected payoff, U^0, with each other tradebot, where U^0 lies in the open interval from 0 (the minimum tolerance level) to H (the highest possible trade payoff).

Assumption (LR): Long-Run Expectations. Each TFTT associates a prior expected payoff C with each other TFTT and a prior expected payoff $(L + C)/2$ with each Rip; and each Rip associates a prior expected payoff $(H + C)/2$ with each TFTT and a prior expected payoff C with each other Rip.

Another key factor affecting the types of trade networks that can emerge in the TNG is the extent to which the benefits and costs associated with each potential trade partner balance out over time, either triggering eventual refusal or permitting long-term partnership. It is therefore useful to examine the various possible 2-tradebot match-ups for the illustrative 5-tradebot TNG before entering into an analysis of the full-blown model.

Given the form of the criterion filter that the tradebots use to update their expectations, a TFTT never judges another TFTT to be intolerable; for the expected payoff of a TFTT for another TFTT is always nonnegative. Similarly, a Rip never judges a TFTT to be intolerable. However, a Rip can become intolerable for another Rip or for a TFTT. In particular, under assumption (CP) one obtains the following four behavioral regions as the prior expected payoff U^0 ranges from low to high values:

- **(CP.1)** $0 < U^0 < -D$: a TFTT finds a Rip intolerable after only one trade, and a Rip finds another Rip intolerable after only one trade;
- **(CP.2)** $-D \leq U^0 < -L$: a TFTT finds a Rip intolerable after only one trade, and a Rip never finds another Rip intolerable;
- **(CP.3)** $-L \leq U^0$ and $C < -L$: a TFTT finds a Rip intolerable after a finite odd number of trades, and a Rip never finds another Rip intolerable;
- **(CP.4)** $-L \leq U^0$ and $-L \leq C$: a TFTT never finds a Rip intolerable, and a Rip never finds another Rip intolerable.

In contrast, under (LR) one obtains four regions characterized by somewhat different transient behaviors as the mutual cooperation payoff C ranges from low to high values:

- **(LR.1)** $0 < C < -D$: a TFTT finds a Rip intolerable prior to trade, and a Rip finds another Rip intolerable after only one trade;
- **(LR.2)** $-D \leq C < -L$: a TFTT finds a Rip intolerable prior to trade, and a Rip never finds another Rip intolerable;
- **(LR.3)** $-L \leq C < -3L$: a TFTT finds a Rip intolerable after only one trade, and a Rip never finds another Rip intolerable;
- **(LR.4)** $-3L \leq C$: a TFTT never finds a Rip intolerable, and a Rip never finds another Rip intolerable.

What, then, are the possible trade networks that can emerge in this 5-tradebot TNG over the course of a single trade cycle loop, assuming either (CP) or (LR) is in effect?

Possible Trade Networks Under Assumption (CP)

At the beginning of the initial trade cycle, each tradebot judges each other tradebot to be equally tolerable, and he uses random selection to submit a trade offer to one of these tradebots. In turn, each tradebot places all received trade orders on his current waiting list, with no refusals. In accordance with the DCR mechanism, each tradebot then accepts all trade offers on his current waiting list.

Suppose that payoffs and prior expected payoffs are configured as in (CP.1). In this case, even though a Rip receives the highest possible payoff, H, from an initial trade with a TFTT, which encourages him to submit another trade offer to this TFTT in the next trade cycle, the TFTT neither submits trade offers to, nor accept trade offers from, this Rip after their first trade. Moreover, the two Rips cease all trade activity with each other after their first trade. Under the DCR mechanism, a tradebot receiving a refusal payoff from a refused trade offer during the course of a trade cycle immediately submits a replacement trade offer to any next-most-preferred tolerable tradebot who has not yet refused him. Consequently, by the end of the first four trade cycles, a Rip has triggered refusal in every one of his potential trade partners. Thereafter the Rip submits trade offers only to the TFTTS, receiving only negative refusal payoffs in return, until the expected payoff he associates with each TFTT finally drops below zero and he turns into a wallflower.

In summary, by the end of the fourth trade cycle, the only trade networks that are viable for case (CP.1) involve trades among the three TFTTs, with both Rips ostracized and eventually reduced to wallflowers; cf. Figure 2(a). The fitness score of each TFTT thus tends towards the mutually cooperative payoff, C, whereas the fitness score of each Rip tends toward the wallflower payoff, 0. Whether of not the Rips survive and prosper in the generation cycle at the end of the trade cycle loop then depends on the length of this loop. Specifically, in order for a Rip to end up with a higher fitness score than the TFTTs, the loop must be short enough so that the H payoffs received by the Rip from his initial successful defections against the three TFTTs sufficiently outweigh the mutual defection payoff, D, that he receives from his one Rip-Rip trade, any refusal payoffs, R, that he receives from subsequent refused attempts to trade with the TFTTs, and any wallflower payoffs, 0, that he receives after ceasing all trade activity.

—INSERT FIGURE 2 ABOUT HERE—

Cases (CP.2) and (CP.3) are similar to case (CP.1), except that the two Rips end up trading cooperatively with each other rather than as wallflowers; cf. Figure 2(b). Also, in case (CP.3) the TFTTs may take longer to reject the opportunistically defecting Rips. The fitness scores of all tradebots thus tend toward the mutually cooperative payoff, C, but it is now more likely that the Rips will have a higher fitness score than the TFTTs at the end of any given trade cycle loop and hence a reproductive advantage in the subsequent generation cycle.

Suppose, now, that case (CP.4) holds with $-L \leq U^0 < (H+C)/2$. The only long-run trade networks viable in this case consist of the three TFTTs engaged in mutually cooperative trades with other TFTTs, with each Rip latched on to one randomly determined TFTT. Figure 2(c) depicts the case in which each Rip happens to be latched on to a different TFTT.

In this case, then, each TFTT risks becoming a full-time host for a parasitical Rip. Prior expected payoffs are low enough to encourage latching behavior on the part of the Rips, who are delighted with their unexpectedly high payoffs from the TFTTs, but not low enough to induce the TFTTs to refuse the Rips. Although the average payoff, $(L+C)/2$, that a TFTT receives from repeated trades with a Rip is nonnegative and possibly positive, it is not as high as the average payoff accruing to a Rip from such trades, $(H+C)/2$, nor as high as the average payoff, C, that accrues to a TFTT in repeated trades with another TFTT. Hence, the relative fitness of a TFTT is lowered by interactions with a Rip, and this puts him at a reproductive disadvantage in the generation cycle.

It is interesting to note that at least one TFTT always avoids becoming parasitized by a Rip in case (CP.4) with $-L \leq U^0 < (H+C)/2$. The fitness score of any such TFTT tends towards C, whereas the fitness score of each parasitized TFTT is uniformly bounded below C. The structurally identical TFTT tradebots thus end up, by chance, with different fitness scores. Nevertheless, given a sufficiently long trade cycle loop, each Rip exits the loop with a higher fitness score than each TFTT; for the fitness score of each Rip tends towards $(H+C)/2$.

Next, consider case (CP.4) with $U^0 = (H+C)/2$. In this case, as depicted in Figure 2(d), each Rip stochastically switches his trade offers back and forth among the three TFTTs for the duration of the trade cycle loop. Hence, each TFTT always has a positive probability of being parasitized by each Rip. The reason for the formation of this type of trade network is that a Rip is indifferent between any two TFTTs with whom he has traded an even number of times.

Finally, consider case (CP.4) with $(H+C)/2 < U^0 < H$. As depicted in Figure 2(d), each TFTT is now a recurrent host for each of the parasitical Rips by the end of the fourth trade cycle. The intuitive reason for the formation of this type of trade network is that each Rip's prior expected payoff for a TFTT is so high that he essentially always prefers the TFTT with whom he has currently traded the least, and this leads him to repeatedly cycle his trade offers among the three TFTTs.

As in the previous (CP.4) cases, the fitness score of each Rip in this final (CP.4) case tends toward $(H+C)/2$. In contrast to the previous (CP.4) cases, however, no TFTT now has any chance of escaping parasitization by all three Rips. Consequently, the fitness score of all three TFTTs is uniformly bounded below C for all sufficiently long trade cycle loops. Here, then, is an example where optimistic prior expectations, leading to increased experimentation, turn out to be detrimental for the nicer tradebots.

Possible Trade Networks Under Assumption (LR)

Comparing the behavioral regions under (LR) with the behavioral regions under (CP), one sees that the TFTTs tend to behave more cautiously under (LR) because their prior expected payoffs are less optimistic. In particular, a TFTT's prior expected payoff for a Rip is bounded strictly below C under (LR) and may even be negative. Consequently, no TFTT ever submits a trade offer to a Rip. Moreover, a TFTT will not accept an initial trade offer from a Rip under (LR) unless the benefit, C, from a mutual cooperation is at least as great as the cost, $-L$, that is incurred when the Rip successfully defects against him.

Consider, now, the trade networks that can emerge under (LR). In the initial trade cycle, each TFTT uses random selection to submit a trade offer to one particular TFTT, and all such trade offers are accepted. Each Rip likewise uses random selection to submit a trade offer to one particular TFTT; but whether or not these trade offers are accepted depends on the behavioral region.

In case (LR.1), a TFTT refuses all trade offers from a Rip prior to any trades taking place. Under the

DCR mechanism, a tradebot who has a trade offer refused immediately submits a replacement trade offer to any next-most-preferred tolerable tradebot who has not yet refused him. Thus, by the end of the initial trade cycle, each Rip has made one trade offer to each TFTT which was refused, and one trade offer to the other Rip which was accepted. Nevertheless, after this one trade, the Rips find each other intolerable and never submit trade offers to each other again. In subsequent trade cycles each Rip only submits trade offers to the TFTTs, collecting negative refusal payoffs until, finally, his expected payoff for each TFTT drops below zero. Thereafter each Rip subsides into wallflowerdom; cf. Figure 3(a).

—INSERT FIGURE 3 ABOUT HERE—

Case (LR.2) differs from case (LR.1) in only one respect—the Rips never find each other intolerable. Thus, in the first few trade cycles, each Rip uses sequential random selection to submit a trade offer to each TFTT in turn, who refuses the offer, and then to the other Rip, who accepts the offer. Since each refusal results in a negative refusal payoff, R, the expected payoff that each Rip associates with each TFTT eventually drops below the expected payoff that each Rip has for the other Rip, which is always nonnegative. In each subsequent trade cycle the Rips submit trade offers only to each other; cf. Figure 3(b).

The interesting aspect of both (LR.1) and (LR.2) is that each TFTT is able to use refusal to protect himself completely from the Rips, so that he *never* sustains any low L payoffs. The fitness score of each TFTT at the end of the trade cycle loop is thus C, because C is the only payoff he ever experiences. In contrast, each Rip sustains negative refusal payoffs as well as at least one negative defection payoff before finally settling down either to wallflowerdom in case (LR.1) or to mutually cooperative trades with the other Rip in case (LR.2). Consequently, the fitness score of each Rip at the end of any trade cycle loop is definitely below C and may even be negative. It follows that each TFTT has a higher fitness score than each Rip at the end of any trade cycle loop, and hence has a reproductive advantage over each Rip in the subsequent generation cycle.

Case (LR.3) is similar to case (LR.2), except that each Rip is able to obtain one high H payoff from each TFTT (inflicting a low negative L payoff on each TFTT in the process) before collecting refusal payoffs. It is therefore possible for a Rip to end up with a higher fitness score than a TFTT in the subsequent generation cycle. In general, then, as depicted in Figure 3(c), neither the TFTTs nor the Rips have a definite long-run reproductive advantage under case (LR.3).

Finally, suppose case (LR.4) holds. The TFTTs continue to submit trade offers only to each other in each successive trade cycle; but, unlike the previous (LR) cases, they never refuse trade offers received from a Rip. Consequently, a Rip never submits a trade offer to the other Rip; for the trade offers he submits to the persistently more attractive TFTTs are never refused.

Indeed, not surprisingly, the behavior of the Rips in case (LR.4) is very similar to the eventual behavior of the Rips in case (CP.4) with $U^0 = (H + C)/2$. Throughout the trade cycle loop, each Rip randomly selects a TFTT to trade with after every even-numbered trade with a TFTT. The result is that, in every other trade cycle, each TFTT has a positive probability of becoming a host for each parasitic Rip for the next two trade cycles, and the Rips never trade with each other at all; cf. Figure 3(d). It follows that each Rip ends up with a higher fitness score than each TFTT, regardless of the length of the trade cycle loop.

Comparing Figure 3 with Figure 2, the TFTTs have an easier time protecting themselves against the Rips in case (LR), where they have an accurate prior understanding of the payoffs they can expect to obtain from each type of tradebot. While this amount of information may be excessive, it does seem reasonable to suppose that, based on past bad experiences, nice tradebots develop more cautious priors for untested trade partners than do street-wise opportunistic tradebots on the look-out for chumps. Alternatively, nice tradebots might develop high minimum tolerance levels, so that refusal occurs quickly if trades go sour. Having a minimum tolerance level set equal to the wallflower payoff, 0, is locally rational, in the sense that positive payoffs, however small, result in a better fitness score for a player than no payoffs at all. Yet a myopic focus on increasing one's own fitness score in absolute terms might not lead to reproductive success in the generation cycle if other opportunistic tradebots such as the Rips are doing even better; cf. cases (CP.4) and (LR.4). Ideally, then, prior expected payoffs and minimum tolerance levels should be allowed to evolve conjointly with the tradebots' strategies. Some preliminary simulation work along these lines in the context of the IPD/CR game can be found in Ashlock et al. (1996, Section 5.4).

3.3 Illustrative TNG Computer Experiments

Two types of TNG computer experiments are reported in this section: (a) buyer-seller market experiments, in which each tradebot is both a buyer and a seller; and (b) two-sided market experiments, in which a subset of buyer tradebots makes offers to a disjoint subset of seller tradebots. All experimental findings were obtained using the C++ implementation of the TNG (version 104b) developed by McFadzean and Tesfatsion (1997,1999).[9] For each experiment, multiple runs from different initial random seeds are described. The following features are set commonly across all of these experimental runs.

The wallflower payoff W is set at 0, the refusal payoff R is set at -0.6, the PD trade payoffs are set at $L = -1.6$, $D = -0.6$, $C = 1.4$, and $H = 3.4$, and each tradebot's minimum tolerance level is set at 0. Each tradebot assigns the same prior expected payoff, 1.4, to each other tradebot, implying that he is initially indifferent concerning which trade partners he interacts with and is fairly optimistic about the payoffs he will receive; and each tradebot assigns a negative prior expected payoff to himself, thus ensuring that he never trades with himself.

The total number of tradebots is set at 24, the number of generations is set at 50, and the number of trade cycles in each trade cycle loop is set at 150. Trade strategies are represented as 16-state finite state machines (FSMs) with fixed starting states and with memory 1, where the memory is the number of bits used to encode the past actions of an opponent that can be recalled in any given internal FSM state. For implementation of the genetic algorithm, the number of elite trade strategies retained unchanged for the next generation is set at 16, where the elite trade strategies are the trade strategies that have attained the highest relative fitness scores. Also, the probability of crossover is set at 1.0 and the probability of a bit mutation is set at 0.005.

As outlined in Table 1, each experimental run has the same dynamic structure. Each of the 24 tradebots in the initial tradebot generation is assigned a trade strategy encoded as a randomly generated bit string. The tradebots are then configured in accordance with user-supplied parameter values, and a trade cycle loop commences. At the end of the trade cycle loop, the randomly generated bit strings that encode the trade strategies of the current tradebot generation are evolved by means of a genetic algorithm employing two-point crossover, bit mutation, and elitism. The evolved set of 24 bit strings decodes to an evolved

[9]For more detailed computer experiments with alternative market structures, see Tesfation (1997).

set of 24 trade strategies (FSMs) that are reinitialized to their fixed starting states and assigned to a new tradebot generation. This new tradebot generation is then configured with user-supplied parameter values, and another trade cycle loop commences.

Buyer-Seller Markets

Each tradebot in these experiments was both a buyer and a seller, implying that he could both make and receive trade offers.

In the first batch of buyer-seller experiments, the acceptance quota of each tradebot was set equal to the total number of tradebots, 24, and the offer quota of each tradebot was set equal to 1. The tradebots were thus effectively unconstrained with regard to the number of trade offers they could have on their waiting lists at any given time.

As a benchmark, experiments were first run with random partner matching in place of the DCR matching mechanism. Random partner matching was effected by preventing the updating of the prior expected payoff 1.4 that each tradebot initially assigned to each potential trade partner, so that all tradebots remained indifferent concerning their potential trade partners and matching was accomplished by the default mechanism of a random draw. Although occasionally the average fitness score achieved by the tradebots under random matching rose to the mutual cooperation level, 1.4, a more typical outcome was a steady decline[10] to the mutual defection level, -0.6; see Figure 4. The size of the refusal payoff is irrelevant for this finding, since refusals never occur in TNG experiments with random matching and nonbinding acceptance quotas.

—INSERT FIGURE 4 ABOUT HERE—

When the DCR matching mechanism was then restored, the average fitness score achieved by the tradebots typically evolved to the mutual cooperation level 1.4; see Figure 5. These TNG experiments reinforce the previous IPD/CR findings of Stanley et al. (1994) and Ashlock et al. (1996) that a preference-based matching mechanism tends to accelerate the emergence of mutual cooperation in the IPD when each agent is permitted both to make and to refuse game offers, is unconstrained with regard to the number of received offers he can accept, and is permitted to have at most one offer outstanding at any given time.

—INSERT FIGURE 5 ABOUT HERE—

In the second batch of buyer-seller experiments, all parameter settings were retained unchanged except that the acceptance quotas were reduced from 24 to 1. Each tradebot could thus retain at most one trade offer on his waiting list at any one time; all other received trade offers had to be refused. Under random partner matching, the typical outcome was again the emergence of an average fitness score close to the mutual defection payoff level, -0.6. This same outcome obtained even when refusal payoffs were omitted from fitness scores, implying that refusal payoffs resulting from limited waiting lists were not a determining factor.

[10]Recall that each tradebot generation participates in only 150 trade cycles, and each tradebot can direct at most one trade offer to any particular tradebot during a trade cycle. Consequently, the maximum number of trades between any two tradebots in a given generation is equal to 300, which is attained only if each tradebot submits an accepted trade offer to the other tradebot in each of the 150 trade cycles in which they participate. Also, the trade strategies for the initial generation in each TNG experiment are randomly generated. A steady decline to mutual defection over the initial 50 generations is therefore not in conflict with the evolution of mutual cooperation observed for the IPD with random or round-robin matching by previous researchers when much longer player interaction lengths were permitted between players in each generation or when the initial strategy population was sufficiently seeded with cooperatively inclined strategies such as Tit-for-Tat.

When the DCR matching mechanism was restored, however, the average fitness score typically leveled out at about 1.25 instead of evolving to the mutual cooperation payoff level 1.4, the outcome for the first batch of buyer-seller experiments. The explanation for this difference appears to lie in the changed nature of the refusal payoffs.

In the first batch of buyer-seller experiments, the acceptance quota (24) was large relative to the offer quota (1). In these circumstances, tradebots are generally refused by other tradebots only if the latter find them to be intolerable because of past defections. Negative refusal payoffs received in response to defections should rightly count against the fitness of the trade strategies generating the defections, for this induces changes in these strategies in the generation cycle that tend to lead to higher future fitness scores. In the second batch of buyer-seller experiments, however, the acceptance quota (1) was much smaller in relation to the offer quota (1), implying that many more received trade offers had to be refused regardless of their desirability. In these circumstances, tradebots tend to accumulate large numbers of negative refusal payoffs purely as a consequence of the relatively small acceptance quota and the nature of the DCR mechanism, regardless of their trade strategies. Since the quotas and the DCR mechanism are not evolved in the current implementation of the TNG, penalizing the tradebots for these quota and DCR effects by including refusal payoffs in their individual fitness scores tends to lower their current average fitness score without inducing a higher average fitness score in the future.

As expected, the average fitness scores attained by the tradebots in the second batch of buyer-seller experiments markedly improved when refusal payoffs were removed from the calculation of the tradebots' individual fitness scores; see Figure 6. Improvement continued to occur when, in addition, the refusal payoffs were reduced in magnitude from -0.60 to -0.30; but a further reduction in magnitude to -0.06 and then to 0 resulted in increasingly volatile maximum and minimum average fitness scores with no discernible improvement in average fitness scores.

—INSERT FIGURE 6 ABOUT HERE—

The probable cause of this increased volatility is that tradebots receiving refusals during initial trade cycles have little incentive to direct their offers elsewhere in subsequent trade cycles when the magnitude of the negative refusal payoff is small. A negative refusal payoff guarantees that the continually updated expected payoff that a tradebot associates with another tradebot who repeatedly refuses him eventually falls below 0, the minimum tolerance level, at which point he ceases making offers to this other tradebot. Nevertheless, this learning process is slow when the magnitude of the negative refusal payoff is small, and it is non-existent when the refusal payoff is simply set to 0.

Two-Sided Markets

In each two-sided market experiment, the 24 tradebots were evenly divided into 12 pure buyers (makers of trade offers) and 12 pure sellers (receivers of trade offers).

In the first batch of experiments, the acceptance quota of each seller was set to 12 and the offer quota of each buyer was set at 1. Thus, sellers were effectively unconstrained regarding the number of trade offers they could have on their waiting lists at any one time. Experiments were first run with random partner matching in place of the DCR matching mechanism to obtain a benchmark for comparison. Interestingly, in contrast to buyer-seller experiments with nonbinding acceptance quotas and random matching, the average fitness score attained by the tradebots tended to fall to a level between -0.4 and the wallflower payoff 0 rather than dropping all the way down to the mutual defection payoff level -0.6; compare Figure 7 with Figure 4.

—INSERT FIGURE 7 ABOUT HERE—

When the DCR matching mechanism was restored, the average fitness score of the tradebots typically evolved to about 1.2, a payoff level markedly below the mutual cooperation level 1.4 obtained in buyer-seller experiments with nonbinding acceptance quotas and DCR matching. Moreover, the maximum fitness score, the average fitness score, and the minimum fitness score attained by the successive tradebot generations persistently deviated from one another. Compare Figure 8 with Figure 5.

—INSERT FIGURE 8 ABOUT HERE—

As detailed in Tesfatsion (1997, Proposition 3.3), the DCR mechanism is only guaranteed to be Pareto optimal for buyers, that is, for the active makers of trade offers. The effects of this bias on average fitness scores is hidden in buyer-seller markets, where each tradebot is both a buyer and a seller. However, for two-sided markets with buyer offer quotas set equal to 1 and with nonbinding seller acceptance quotas, the DCR mechanism appears to result in a "separating equilibrium" in which the buyers are generally achieving high fitness scores and the sellers are generally achieving low fitness scores. In particular, the extreme pickiness of buyers combined with the acceptance by sellers of all tolerable received trade offers appears to allow buyers to form long-run parasitic relations with sellers, i.e., relations characterized by successful defections within the limits permitted by the sellers' 0 minimum tolerance levels. In the second batch of two-sided market experiments, all parameter specifications were retained unchanged except that the seller acceptance quotas were decreased from 12 to 1. Thus, instead of accepting all tolerable received trade offers, the sellers now accepted at most one received trade offer per trade cycle.

When benchmark experiments were first run with random partner matching in place of the DCR mechanism, the typical outcome was the emergence of an average attained fitness score close to the mutual defection payoff, -0.6. This result obtained whether or not refusal payoffs were counted in the calculation of individual fitness scores. When the DCR matching mechanism was then restored, with refusal payoffs counted in the calculation of individual fitness scores, the accumulation of refusal payoffs tended to result in an average attained fitness score that was markedly below the mutual cooperation payoff level. When refusal payoffs were then omitted from the calculation of individual fitness scores, the average attained fitness score tended to evolve to the mutual cooperation level 1.4 and to be close to the maximum attained fitness scores; see Figure 9.

—INSERT FIGURE 9 ABOUT HERE—

Given the structuring of the two-sided market experiments currently under discussion, with equal numbers of buyers and sellers, common acceptance and offer quotas, and tradebots who are initially indifferent among their potential trade partners, the buyer offer quota roughly predicts the number of trade offers that each seller will receive in the initial few trade cycles. The larger this number is, the more chance there is that mutually cooperative matches between buyers and sellers will be quickly discovered. On the other hand, if the seller acceptance quota is small relative to the buyer offer quota, then many buyers will accumulate negative refusal payoffs unrelated to the nature of their trade strategies; and, if the seller acceptance quota is large relative to the buyer offer quota, opportunistic buyers have a greater chance to seek out and exploit sellers within the limits allowed by their minimum tolerance levels. Either circumstance could slow or even prevent the sustained emergence of mutually cooperative behavior.

It might therefore be conjectured that mutually cooperative behavior will best be induced in the current experimental setting when the seller acceptance quota and the buyer offer quota are equal and sufficiently large. Indeed, this turns out to be the case. In various two-sided market experiments with 12 pure buyers, 12 pure sellers, and equal seller and buyer quotas ranging from 3 to 12, the average fitness score attained by the tradebots tended to evolve to the mutual cooperation payoff level and to be close to the maximum attained fitness score even when refusal payoffs were included in the calculation of individual fitness scores.

4 Concluding Remarks

The hallmark of the ACE approach to social and biological modelling is a bottom up perspective, in the sense that global behavior is grounded in local agent interactions. The agent-based trade network game (TNG) illustrates how the ACE approach can be specialized to an economic context. In particular, the analytical and simulation findings presented in the previous section illustrate how the TNG is currently being used to study the evolutionary implications of alternative market structures at three different levels: individual trade behavior; trade network formation; and social welfare as measured by average agent fitness.

As currently implemented, however, the TNG only partially achieves the goal of a bottom up perspective. The TNG tradebots are surely more autonomous than agents in traditional economic models. For example, in order to determine their trade partners, the tradebots send messages back and forth to each other at event-triggered times. Nevertheless, they are still controlled by a main program that synchronizes the commencement of their trading activities and the evolution of their trade behavior. The advantage of imposing this synchronized dynamic structure is that it permits some analytical results to be obtained concerning the configuration, stability, uniqueness, and social optimality of the trade networks that emerge. The disadvantage is that these networks may not be robust to realistic relaxations of the imposed synchronizations.

As the TNG illustrates, then, the challenges to economists posed by the ACE approach are great and the payoffs are yet to be fully determined. Using the ACE approach, however, economists can at last begin to test seriously the self-organizing capabilities of decentralized market economies.

REFERENCES

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

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

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

4. Arthur, W. Brian, S. N. Durlauf, and D. A. Lane, eds. (1997), *The Economy as an Evolving Complex System, II*, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XXVII, Addison-Wesley, Reading, MA.

5. Ashlock, D., M. D. Smucker, E. A. Stanley, and L. Tesfatsion (1996), "Preferential Partner Selection in an Evolutionary Study of Prisoner's Dilemma," *BioSystems* Nos. 1-2, Vol. 37, 99–125.

6. Axelrod, R. *The Evolution of Cooperation*, Basic Books, N.Y., 1984.

7. Axelrod, R., "The Evolution of Strategies in the Iterated Prisoner's Dilemma," in L. Davis (ed.), *Genetic Algorithms and Simulated Annealing*, Morgan Kaufmann, Los Altos, CA, 1987.

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

9. Birchenhall, C. R. (1995), "Modular Technical Change and Genetic Algorithms," Working Paper No. 9508, University of Manchester.

10. Bosch, A., and S. Sunder (1996), "Tracking the Invisible Hand: Convergence of Double Auctions to Competitive Equilibrium," Economics Working Paper 91, Universitat Pompeu Fabra, Revised.

11. Brock, W., and S. N. Durlauf (1995), "Discrete Choice with Social Interactions," Working Paper No. 95-10-084, Santa Fe Institute, Santa Fe, NM.

12. Bullard, J., and J. Duffy (1994), "A Model of Learning and Emulation with Artificial Adaptive Agents," Economics Working Paper, University of Pittsburgh.

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

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

15. Epstein, J. M., and R. Axtell (1996), *Growing Artificial Societies: Social Science from the Bottom Up*, MIT Press/Brookings, Cambridge, MA. [Reviewed by L. Tesfatsion, *Journal of Economic Literature* XXXVI (March 1998), pp. 233-234 (http://www.econ.iastate.edu/tesfatsi/epaxrev.ps)].

16. Fogel, D. B. (1995), "On the Relationship Between the Duration of an Encounter and the Evolution of Cooperation in the Iterated Prisoner's Dilemma," *Evolutionary Computation* 3, 349–363.

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

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

19. Goldberg, D. (1989), *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, MA.

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

21. Hauk, E. (1996), "Leaving the Prison: A Discussion of the Iterated Prisoner's Dilemma under Preferential Partner Selection," Thesis, European University Institute, Florence.

22. Hirshleifer, D., and E. Rasmusen (1989), "Cooperation in a Repeated Prisoners' Dilemma with Ostracism," *Journal of Economic Behavior and Organization* 12, 87–106.

23. Holland, J. (1992), *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.

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

25. Ioannides, Y. M. (1997), "Evolution of Trading Structure," pp. 129–167 in Arthur et al. (eds.), *The Economy as an Evolving Complex System, II*, *op. cit.*.

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

27. Kirman, A. P. (1997), "The Economy as an Interactive System," pp. 491–531 in Arthur et al. (eds.), *The Economy as an Evolving Complex System, II*, *op. cit.*.

28. Kitcher, P. (1993), "The Evolution of Altruism," *The Journal of Philosophy* 90, 497–516.

29. Lane, D. (1993), "Artificial Worlds and Economics," *Journal of Evolutionary Economics* 3, 89-107.

30. Langton, C., ed., (1989), *Artificial Life*, Addison-Wesley, Redwood City, California.

31. Levy, S. (1992), *Artificial Life*, Pantheon Books, New York.

32. Lindgren, K., and M. G. Nordahl (1994), "Cooperation and Community Structure in Artificial Ecosystems," *Artificial Life* 1, 15-37.

33. Mailath, G., L. Samuelson, and A. Shaked (1994), "Evolution and Endogenous Interactions," SSRI Working Paper 9426, UW-Madison, June.

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

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

36. McFadzean, D. (1995), *SimBioSys: A Class Framework for Evolutionary Simulations*, Master's Thesis, Department of Computer Science, University of Calgary, Alberta, Canada. [The source code for SimBioSys is available for downloading as freeware at the Web site of Leigh Tesfatsion.]

37. McFadzean, D., and L. Tesfatsion (1997), "An Agent-Based Computational Model for the Evolution of Trade Networks," pp. 73–83 in P. Angeline, R. Reynolds, J. McDonnell, and R. Eberhart (eds.), *Evolutionary Programming VI*, Proceedings of the Sixth International Conference on Evolutionary Programming, Springer-Verlag, Berlin. [This proceedings paper is an abbreviated version of the following study.]

38. McFadzean, D., and L. Tesfatsion (1999), "A C++ Platform for the Evolution of Trade Networks," *Computational Economics* 14, 109–134. [The source code for this trade network game (TNG) C++ framework is available for downloading as freeware, along with instructions, at the Web site of Leigh Tesfatsion.]

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

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

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

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

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

44. Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Clarendon.

45. Sigmund, K. (1993) *Games of Life: Explorations in Ecology, Evolution, and Behavior*, Oxford University Press, Oxford.

46. Smucker, M. D., E. A. Stanley, and D. Ashlock (1994), "Analyzing Social Network Structures in the Iterated Prisoner's Dilemma with Choice and Refusal," Department of Computer Sciences Technical Report CS-TR-94-1259, UW-Madison.

47. Stanley, E. A., D. Ashlock, and L. Tesfatsion (1994), "Iterated Prisoner's Dilemma with Choice and Refusal of Partners," 131-175 in C. Langton, ed., *Artificial Life III*, Proceedings Volume 17, SFI Studies in the Sciences of Complexity, Addison-Wesley, Reading, MA.

48. Tesfatsion, L. (1979) "Direct Updating of Intertemporal Criterion Functions for a Class of Adaptive Control Problems." *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-9, 143–151.

49. Tesfatsion, L. (1995), "How to Get Alife," *CSWEP Newsletter*, American Economic Association, Winter Issue, 16–18.

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

51. Tesfatsion, L. (1999a), “Hysteresis in an Evolutionary Labor Market with Adaptive Search,” Economic Report No. 50, Iowa State University, October 1999, to appear in: S. H. Chen, ed., *Evolutionary Computation in Economics and Finance*, Springer-Verlag, Berlin.

52. Tesfatsion, L. (1999b), “Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search,” Economic Report No. 51, Iowa State University, October 1999, to appear in the *Journal of Economic Dynamics and Control*.

53. Vriend, N. J. (1995), “Self-Organization of Markets: An Example of a Computational Approach,” *Computational Economics* 8, 205–231.

54. Witt, U. (1993), *Evolutionary Economics*, Edward Elgar, London.

55. Young, P. (1993), “The Evolution of Conventions,” *Econometrica* 61, 57–84.

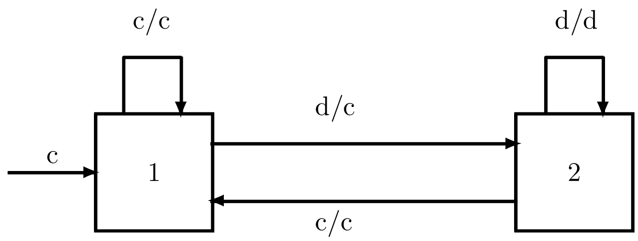

(a) Tit-for-Two-Tats

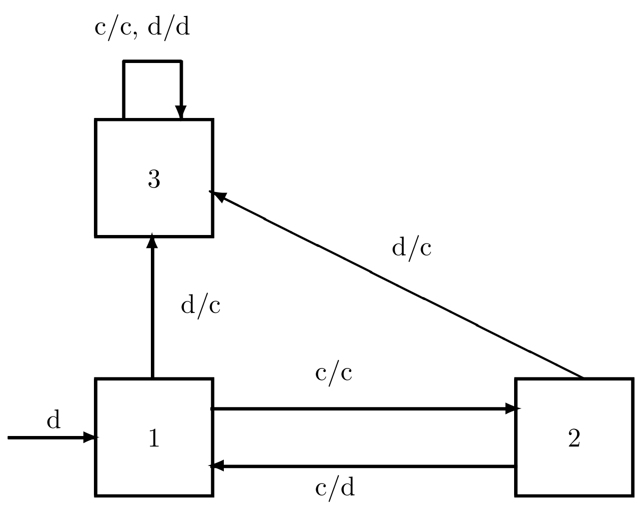

(b) Rip-Off

Figure 1: The FSM Representations for Two Illustrative Trade Strategies. (a) A nice trade strategy that starts by cooperating and only defects if defected against twice in a row; (b) An opportunistic trade strategy that starts by defecting and defects every other time unless defected against. An arrow label x/y means that y is the next move to be taken in response to x, the last move of one's opponent.

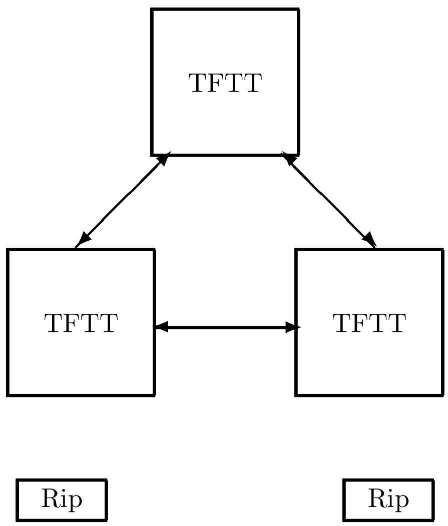

(a) Case (CP.1)

$0 < U^0 <$ ¡ D

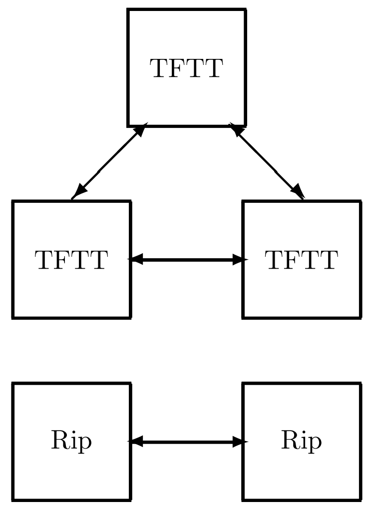

(b) Case (CP.2) or (CP.3)

(¡ D · $U^0 <$ ¡ L) or (C < ¡ L · U^0)

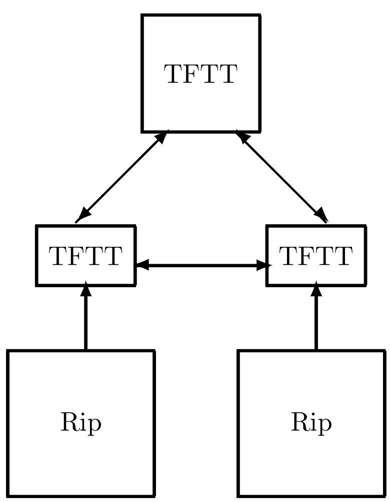

(c) Case (CP.4): ¡ L · C with

¡ L · $U^0 <$ (H + C)=2

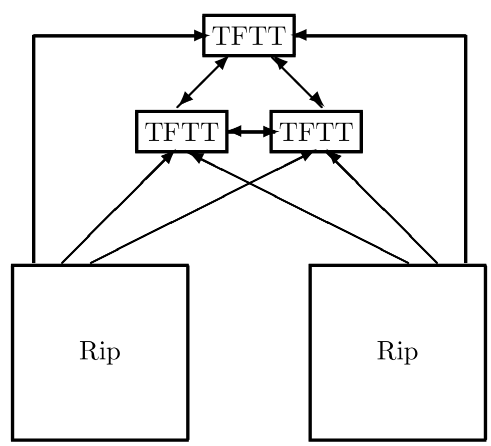

(d) Case (CP.4): ¡ L · C with

(H + C)=2 · U^0

Figure 2: Long-Run Trade Networks Under Assumption (CP) for the Illustrative 5-Tradebot TNG. A relatively larger box indicates a definitely higher fitness score for a sufficiently long trade cycle loop. In case (d), the Rip-TFTT interactions are stochastic if $(H + C)/2 = U^0$ and deterministic if $(H + C)/2 < U^0$.

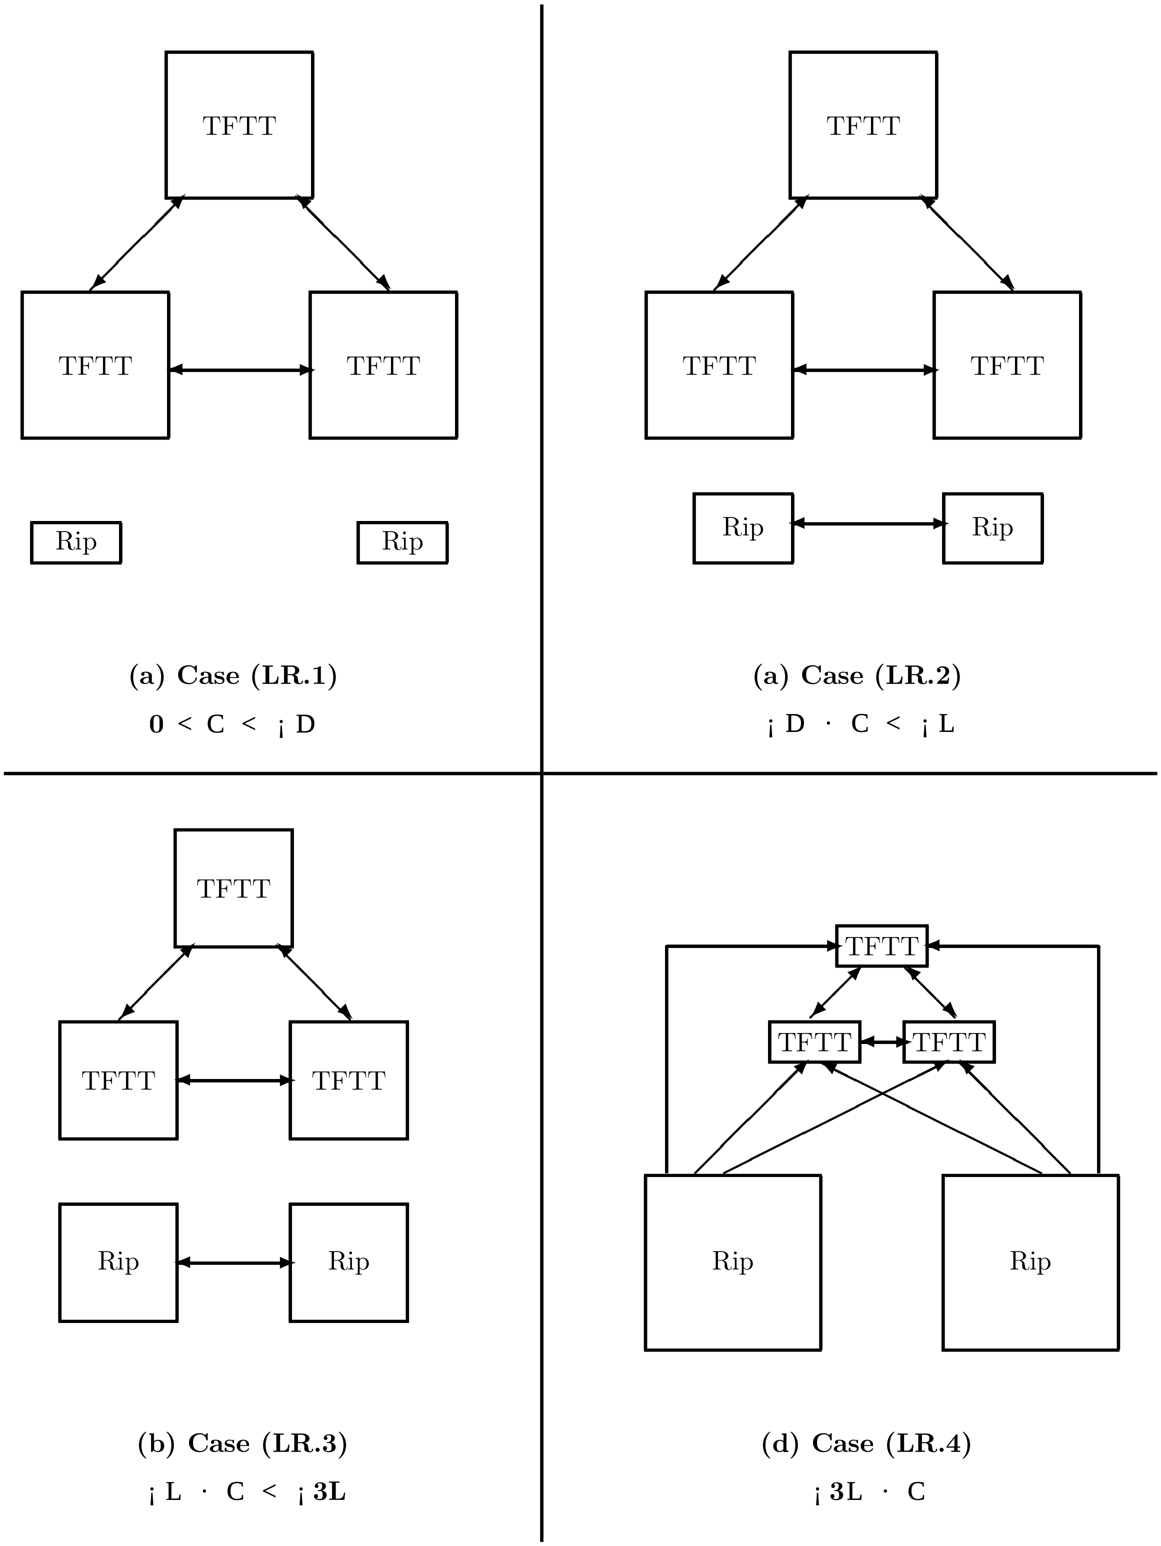

(a) Case (LR.1)

0 < C < ¡ D

(a) Case (LR.2)

¡ D · C < ¡ L

(b) Case (LR.3)

¡ L · C < ¡ **3L**

(d) Case (LR.4)

¡ **3**L · C

Figure 3: Long-Run Trade Networks Under Assumption (LR) for the Illustrative 5-Tradebot TNG. A relatively larger box indicates a definitely higher fitness score for a sufficiently long trade cycle loop. In case (d), the Rip-TFTT interactions are stochastic.

This figure is available from the author upon request.

Figure 4: Buyer-seller average fitness with random matching and seller quotas equal to 24.

This figure is available from the author upon request.

Figure 5: Buyer-seller average fitness with DCR matching and seller quotas equal to 24.

This figure is available from the author upon request.

Figure 6: Buyer-seller average fitness with DCR matching, seller quotas equal to 1, and refusal payoffs omitted from fitness scores.

This figure is available from the author upon request.

Figure 7: Two-sided market average fitness with random matching and seller quotas equal to 12.

This figure is available from the author upon request.

Figure 8: Two-sided market average fitness with DCR matching and seller quotas equal to 12.

This figure is available from the author upon request.

Figure 9: Two-sided market average fitness with DCR matching, seller quotas equal to 1, and refusal payoffs omitted from fitness scores.