An Evolutionary Trade Network Game With Preferential Partner Selection

Leigh Tesfatsion*
Department of Economics & Department of Mathematics
Iowa State University, Ames IA

Economic Report Series No. 38
February 1996

*375 Heady Hall
Ames, IA 50011-1070
Phone: 515-294-0138
E-mail: tesfatsi@iastate.edu

Abstract: An evolutionary *trade network game* (TNG) is proposed for studying the interplay between evolutionary game dynamics and preferential partner selection in various market contexts with distributed adaptive agents. The modular form of the TNG facilitates experimentation with alternative specifications for trade partner matching, trading, expectation updating, and trade strategy evolution. Experimental results obtained using a C++ implementation suggest that the conventional optimality properties used to evaluate agent matching mechanisms in static market contexts may be inadequate measures of optimality from an evolutionary perspective.

February 1996

An Evolutionary Trade Network Game With Preferential Partner Selection[1]

By Leigh Tesfatsion

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

Abstract

An evolutionary *Trade Network Game (TNG)* is proposed for studying the interplay between evolutionary game dynamics and preferential partner selection in various market contexts with distributed adaptive agents. The modular form of the TNG facilitates experimentation with alternative specifications for trade partner matching, trading, expectation updating, and trade strategy evolution. Experimental results obtained using a C++ implementation suggest that the conventional optimality properties used to evaluate agent matching mechanisms in static market contexts may be inadequate measures of optimality from an evolutionary perspective.

Keywords: Evolutionary game; distributed adaptive agents; endogenous interactions; matching; genetic algorithm; artificial life; C++ implementation.

1 Introduction

Evolutionary game studies typically focus on the optimality properties of strategy configurations when agents are matched randomly or deterministically by some extraneous device. The optimality properties of the matching mechanism per se are generally not considered ([10], [13]). In contrast, optimal search studies focus on the optimality properties of preference-based agent matching mechanisms, but generally these studies are set in static contexts [17]. Since actual socio-economic interactions are often characterized both by evolutionary dynamics and by preference-based partner selection, studying both aspects together seems a logical and interesting next step to take.

[1] To appear in L. Fogel et al. (eds.), *Evolutionary Programming V: Proceedings volume for the Fifth Annual Conference on Evolutionary Programming*, MIT Press.

This issue is addressed in Stanley et al. [19]. The standard evolutionary iterated prisoner's dilemma (IPD) is extended to an evolutionary IPD with choice and refusal (IPD/CR) by allowing players to choose and refuse game partners in each iteration on the basis of continually updated expected payoffs.[2] The introduction of choice and refusal fundamentally alters the way in which players interact in the IPD and the characteristics that result in high payoff scores. Choice allows players to increase their chances of encountering other cooperative players, and refusal gives players a way to protect themselves from defections without having to defect themselves. The ostracism of defectors occurs endogenously as an increasing number of players individually refuse the defectors' game offers. Nevertheless, choice and refusal also permit opportunistic players to home in quickly on exploitable players and form parasitic relationships.

The computer experiments reported in [19] and in the subsequent studies by Ashlock et al. [1], Smucker et al. [18], and Hauk [7] indicate that the emergence of mutual cooperation in the standard evolutionary IPD is accelerated by the introduction of preferential choice and refusal of partners. The underlying player interaction patterns induced by choice and refusal can be complex and highly path dependent, however, even when expressed play behavior is largely cooperative. Consequently, it has proved difficult to characterize the mapping from parameter configurations to evolutionary outcomes for the IPD/CR.

A potentially useful way to proceed, then, is to focus on more concrete settings which impose problem-specific constraints on agent interactions. In Tesfatsion [21] an evolutionary *Trade Network Game (TNG)* is proposed for studying the interplay between evolutionary dynamics and preferential partner selection under alternatively specified market structures.

The player set for the TNG consists of buyer and seller tradebots who choose and refuse trade partners on the basis of continuously updated expected payoffs. Buyers make trade offers to preferred sellers which the sellers either accept or refuse. A trade offer is an invitation to engage in a risky trade modelled as a two-player game. Each buyer and seller initially associates a prior expected payoff with each potential trade partner and randomly adopts a strategy for use in subsequent trades. The buyers and sellers then enter into a trade cycle loop consisting of successive rounds of partner matching, resource-constrained trading, and updating of expected payoffs. At the end of the trade cycle loop the buyers and sellers enter into an evolutionary step in which trade strategies successful in past trades are retained while trade strategies unsuccessful in past trades are replaced with variants of more successful strategies. A new trade cycle loop then commences.

The modular form of the TNG facilitates experimentation with alternative specifications for trade partner matching, trading, expectation updating, and trade strategy evolution.

[2] Other game theory studies that have allowed players to avoid unwanted interactions, or to affect the probability of interaction with other players through their own actions, include Guriev and Shakhova [6], Hirshleifer and Rasmusen [8], Kitcher [11], Mailath et al. [12], and Orbell and Dawes [16]. There is also a growing body of work by economists 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 [2], De Vany [3], Durlauf [4], Ionnides [9], and Vriend [22].

This paper presents experimental results obtained for the particular TNG module specifications developed in Tesfatsion [21], using a recently completed C++ implementation (McFadzean and Tesfatsion [15]). As will be clarified in the following section, trade partners are determined in accordance with a "deferred choice and refusal" (DCR) mechanism, a modified Gale-Shapley matching mechanism [5] that retains the static optimality properties of the original Gale-Shapley mechanism. Also, expected payoffs are updated by means of a simple learning algorithm that yields consistent estimates. A trade is modelled as a prisoner's dilemma game, and trade (IPD) strategies are evolved by means of a standardly specified genetic algorithm.

Two types of markets are considered: buyer-seller markets, and two-sided markets. In the buyer-seller market, each tradebot is both a buyer and a seller in the sense that he can both make and receive trade offers. In the two-sided market, the set of buyers (tradebots who can make offers) is disjoint from the set of sellers (tradebots who can receive offers). For each type of market, attention is focused on the average fitness score achieved by the tradebots as the market evolves, and the degree to which they display mutually cooperative behavior. One interesting finding is that the DCR mechanism imposes high transaction costs on the tradebots, which can lower their average fitness score relative to other less sophisticated matching mechanisms. Another interesting finding concerns the evolutionary effects of the bias of the DCR mechanism in favor of buyers (active makers of trade offers). For two-sided markets, this bias appears to permit buyers to form long-term parasitic relations with sellers that reduce seller fitnesses relative to buyer fitnesses and that hinder the emergence of mutually cooperative behavior.

Overall, these experimental findings suggest that the conventional optimality properties used to evaluate agent matching mechanisms in static market contexts may be inadequate measures of optimality from an evolutionary perspective.

2 The Basic Trade Network Game

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. For example, the buyers and sellers might represent customers and retail store owners, workers and employers, borrowers and lenders, or auction traders.

Each generation of tradebots participates in a *trade cycle loop* consisting of a fixed number of trade cycles. 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 O_m and A_n are strictly positive. One interpretation for the buyer offer quota O_m is that buyer m has a limited amount of resources (credit, labor time, collateral,...) to trade in exchange for other items, and one interpretation for the seller acceptance quota A_n is that seller n has a limited amount of items (goods, job openings, loans,...) to provide.

The tradebots determine their submission, acceptance, and refusal of trade offers in each trade cycle using a modifed version of the well-known Gale-Shapley deferred acceptance mechanism [5]. This modified mechanism, referred to below as the *deferred choice and refusal (DCR) mechanism*, presumes that each buyer and seller associates an expected payoff with each potential trade partner. (The way in which these expected payoffs are determined is clarified below.) 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 then 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. If a buyer has a trade offer refused, he 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 already refused him. When all trade offers cease, each seller accepts all buyer trade offers currently on his waiting list.

The buyer-seller matching outcomes generated by the DCR mechanism exhibit the usual 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 [21, Proposition 4.2]. 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 [21, Proposition 4.3].

During the DCR matching process, any tradebot that has a trade offer refused is immediately penalized by receipt of a negative *refusal payoff*, R; the tradebot who does the refusing is not penalized. If a tradebot neither submits nor accepts any trade offers during this matching process, he receives a nonpositive *wallflower payoff*, W.

A trade offer is an offer by a buyer to a seller to participate in a risky trade modelled as a prisoner's dilemma (PD) game. For example, the trade 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 is the lowest possible payoff, received by a cooperative tradebot whose trade partner defects;

Table 1. Payoff Matrix for the Prisoner's Dilemma Game

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

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 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 1.

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 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. No tradebot knows any other tradebot's strategy *a priori*; he can only learn about it by engaging the other tradebot in repeated trades and observing the payoff history that ensues. 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.

At the beginning of the initial trade cycle loop, before any actual trades have taken place, each tradebot v associates an exogenously given *prior expected payoff* $U_v^0(k)$ with each potential trade partner k. Throughout each trade cycle, tradebot v then uses a simple variable-gain criterion filter [20] to update his current expected payoffs on the basis of the new payoffs he obtains from interactions with his potential trade partners. In particular, if tradebot v receives a payoff P from an interaction with a potential trade partner k, then 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, where the inverse of the weight on P is 1 plus tradebot v's current payoff count with k. In this way, tradebot v keeps a running tab on the payoff outcomes of his interactions with k. As explained in Tesfatsion [21, Section 5], 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.

At the end of a trade cycle loop, the fitness score of each tradebot is calculated as the average of all of the trade, refusal, and wallflower payoffs he received during this loop.

Table 2. Pseudo-Code Depiction of the TNG

```
int main () {
    tngInit() ;                        // Specify prior expected payoffs and
                                       //    construct initial tradebot generation.
    For (G = 0,...,GMAX-1) {           // Enter the generation loop.
        For (I = 0,...,IMAX-1) {       // Enter the trade cycle loop.
            MatchTraders() ;           // Determine actual trade partners,
                                       //    given current expected payoffs, and
                                       //    record refusal and wallflower payoffs.
            Trade() ;                  // Implement trades and
                                       //    record trade payoffs.
            UpdateExp() ;              // Update expected payoffs on the
                                       //    basis of newly recorded payoffs.
        }
        AssessFitness() ;              // Assess and record fitness scores.
        If (G < GMAX-1) {
            EvolveGen() ;              // Genetic step: Evolve the trade strategies
                                       //    of the current tradebot generation.
        }
    }
    Return 0 ;
}
```

The tradebots then enter into a *genetic step* in which the trade strategies used by the *elite* (the fittest tradebots) are retained while the strategies used by the remaining tradebots are replaced with variants of the strategies used by the elite. The finite state machine (FSM) and genetic algorithm (GA) used in the genetic step to represent and evolve the tradebots' trade strategies are the same as detailed in Ashlock et al. [1, Section 2.3] for player strategies in the IPD/CR except that two-point rather than one-point cross-over is employed. The particular FSM and GA parameter settings used in all experiments reported in this paper are detailed in the next section.

At the end of the genetic step the memories of the tradebots are reset to zero, their associated FSMs are reset to a fixed initial state, and their expected payoffs are reset to the prior expected payoff levels. A new trade cycle loop then commences. Table 2 depicts the overall logical progression of the TNG in pseudo-code.

Before reporting on some of the TNG computer experiments conducted to date, various special cases of the TNG will be sketched to indicate the range of economic applications it encompasses.

Case 1: A Labor Market Modelled as an Assignment Game with Choice and Refusal

The subset B consists of M workers and the subset S consists of N employers, where B and S are disjoint. Each worker m can submit work offers to a maximum of O_m employers, or he can choose to be unemployed and receive the known payoff W. Each employer n can hire up to A_n workers, and employers can refuse work offers. Once matched, a worker and employer engage in work site interactions modelled as a PD game.

This TNG special case extends the usual assignment problem [17] by incorporating subsequent strategic game play between matched pairs of agents and by having game play iterated over time. Assignment problems are commonly used by economists to model job-matching in labor markets as well as other economic processes, but the payoff outcome for each matched pair of agents is usually specified a priori.

Case 2: A Labor Market with Endogenously Determined Workers and Employers

The subsets B and S coincide. Each tradebot v can submit up to O_v work offers to tradebots at other work sites and receive up to A_v work offers at his own work site. The degree to which any accepted work offer results in satisfactory outcomes for the participant tradebots is determined by subsequent PD game play. 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 submit 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 submit nor accept any work offers.

Case 3: 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 accepted offers are satisfactorily fulfilled is determined by subsequent PD game play.

3 TNG Computer Experiments

Four types of computer experiments are discussed in the present section: (a) buyer-seller market experiments with sellers unconstrained by acceptance quotas; (b) buyer-seller market experiments with seller acceptance quotas set to 1; (c) two-sided market experiments with sellers unconstrained by acceptance quotas; and (d) two-sided market experiments with seller acceptance quotas set to 1. All experimental findings reported below were obtained using *TNG*, a C++ trade platform developed by McFadzean and Tesfatsion [15] that is supported by the C++ abstract base classes developed by McFadzean [14] for a general artificial life platform, *SimBioSys*. These findings are preliminary in the sense that only average fitness scores are considered. A more thorough understanding of these findings will require delving more deeply into the underlying trade patterns and the trade strategies that support these trade patterns.

For each type of experiment, multiple runs from different initial random seeds are reported. The following features are set commonly across all of these experimental runs.[3] The wallflower payoff W is set at 0, the refusal payoff R is set at $-.6$, 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, $U^0 = C$, to each other tradebot; and each tradebot assigns a negative prior expected payoff to himself, thus ensuring that he never trades with himself. Each buyer tradebot has an offer quota of 1, meaning that he can have at most one trade offer outstanding to a seller at any given time. The total number of tradebots is set at 24, and the 16 most fit tradebots in each generation are taken to be the elite. The number of trade cycles in each trade cycle loop is set at 150, and the number of generations is set at 50.

Each trade strategy is represented by a 16-state FSM with a fixed initial state and with memory 1. At the beginning of the first trade cycle loop, a bit string coding for each FSM is randomly generated. At the end of each trade cycle loop, the current population of trade strategies (FSMs coded as bit strings) is evolved by means of a genetic algorithm employing two-point cross-over and bit mutation. The probability of cross-over is set at 1.0 and the probability of a bit mutation is set at 0.005.

3.1 Buyer-Seller Market Experiments

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 at 24, the total number of tradebots. Since offer quotas in these experiments were set at 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.[4] Although occasionally the average fitness score achieved by the tradebots rose to the mutual cooperation level, 1.4, a more typical outcome was a steady decline to the mutual defection level, -0.6; see Figure 1. In contrast, when the DCR matching mechanism was restored, the average fitness score achieved by the tradebots typically evolved to the mutual cooperation level; see Figure 2.

In the second batch of buyer-seller experiments, the acceptance quotas were reduced from 24 to 1. Under random partner matching, the typical outcome was again the emergence of an average fitness score close to the mutual defection level. 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 level of 1.4.

The explanation for the latter finding lies in the changed role of the refusal payoff. With

[3] These settings permit comparisons with the IPD/CR experiments reported in Stanley et al. [19] and Ashlock et al. [1].

[4] Random matching was implemented by commenting out the updating of the tradebots' prior expected payoffs, so that the expected payoff each tradebot associated with each potential trade partner remained at the prior expected payoff level $U^0 = C$. Each tradebot thus remained indifferent among all potential trade partners and partner matching was settled by the default mechanism of a random draw.

high acceptance quotas, a tradebot is generally refused by another tradebot only if the latter finds him to be intolerable because of past defections. 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 genetic step that tend to lead to higher future fitness scores.

With low acceptance quotas, however, each tradebot is forced to be very picky. In each trade cycle he can only have a small number of trade offers on his waiting list at any one time no matter how many desirable trade offers he receives. Consequently, the tradebots tend to amass large numbers of negative refusal payoffs due to excess demand effects. These excess demand effects are consequences of the fixed quota specifications and the DCR mechanism and not the trade strategies per se. Since neither the quotas nor the DCR mechanism are evolved in the current implementation of the TNG, penalizing the tradebots for excess demand effects by including refusal payoffs in their fitness scores simply lowers their current average fitness score without inducing a higher average fitness score in the future.

As expected, the average fitness score achieved by the tradebots improved considerably when either the magnitude of the refusal payoff was reduced or the refusal payoffs were removed from the calculation of the tradebots' fitness scores; see Figure 3. Reducing the magnitude of the refusal payoff all the way to 0 led to disastrous results, however, since tradebots receiving refusals during initial trade cycles were then given no incentive to direct their offers elsewhere in subsequent trade cycles. With a strictly negative refusal payoff, the continually updated expected payoff that a tradebot associates with another tradebot who repeatedly refuses him eventually drops below 0, the minimum tolerance level, at which point he ceases making offers to this other tradebot.

3.2 Two-Sided Market Experiments

In each of these experiments, 12 of the tradebots were pure buyers and the remaining 12 were pure sellers.

In the first batch of experiments, the acceptance quota of each seller was set at 12 so that 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 the buyer-seller experiments, the average fitness score of the tradebots tended to fall to a level slightly below the wallflower payoff of 0 rather than dropping all the way down to the mutual defection payoff; see Figure 4.

Moreover, 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; see Figure 5. Another interesting finding illustrated by Figure 5 is that the maximum fitness score, the average fitness score, and the minimum fitness score attained by the successive tradebot generations persistently deviated from one another; contrast Figure 5 with Figure 2. As discussed in Section 2, the DCR mechanism is only guaranteed to be optimal for buyers. Consequently, it is conjectured that the two-sided market structure with seller acceptance quotas set at 12 results 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 acceptance by sellers of all tolerable received trade offers may 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, the seller acceptance quotas were decreased from 12 to 1. Once again the amassment of refusal payoffs resulting from low seller acceptance quotas tended to lower the average fitness score of the tradebots. When the refusal payoffs were either decreased in magnitude or removed from the calculation of individual fitness scores, the average fitness score of the tradebots tended to evolve to the mutual cooperation level and to be markedly closer to the maximum attained fitness score; see Figure 6. Intuitively, low acceptance quotas better enable sellers to protect themselves against potentially parasitic buyers, thus ameliorating the bias of the DCR mechanism in favor of buyers.

4 Concluding Remarks

Overall, the *TNG* experiments run to date indicate that the choice and refusal mechanism should ideally be allowed to evolve conjointly with the tradebots' trade strategies.[5] Moreover, the exogenously specified offer and acceptance quotas should be replaced by endogenously generated budget constraints and production relations. It will be interesting to see whether Gale-Shapley type matching mechanisms survive in this more stringent evolutionary context.

References

[1] D. Ashlock, M. D. Smucker, E. A. Stanley, and L. Tesfatsion, "Preferential Partner Selection in an Evolutionary Study of Prisoner's Dilemma," *BioSystems* 37, No. 1-2 (1996), 99-125.

[2] W. Brock and S. Durlauf, "Discrete Choice with Social Interactions," Working Paper, Department of Economics, University of Wisconsin-Madison, June 1995.

[3] A. De Vany, "The Emergence and Evolution of Self-Organized Coalitions," to appear in M. Gilli, ed., *Computational Methods in Economics and Finance*, Kluwer Scientific Publishers.

[4] S. N. Durlauf, "Neighborhood Feedbacks, Endogenous Stratification, and Income Inequality," in W. A. Barnett, G. Gandolfo, and C. Hillinger, Eds., *Disequilibrium Dynamics: Theory and Applications*, Cambridge University Press, 1995, to appear.

[5] Some preliminary simulation work along these lines in the context of the IPD/CR can be found in Ashlock et al. [1].

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

[6] S. Guriev and M. Shakhova, "Self-Organization of Trade Networks in an Economy with Imperfect Infrastructure," Proceedings of the International Workshop on Agents, Cooperation, and Communication, San Marco di Castellabate, Italy, 1995.

[7] E. Hauk, "Leaving the Prison: a Discussion of the Iterated Prisoner's Dilemma under Preferential Partner Selection," Working Paper, European University Institute, 1995.

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

[9] Y. M. Ionnides, "Evolution of Trading Structures," Working Paper, Department of Economics, Tufts University, August 1995.

[10] A. Kirman, "Economies with Interacting Agents," Working Paper No. 94-05-030, Santa Fe Institute, NM, April 1994.

[11] P. Kitcher, "Evolution of Altruism in Repeated Optional Games," Working Paper, Department of Philosophy, UC San Diego, 1992.

[12] G. Mailath, L. Samuelson, and A. Shaked, "Evolution and Endogenous Interactions," SSRI Working Paper 9426, University of Wisconsin-Madison, June 1994.

[13] K. Lindgren and M. G. Nordahl, "Cooperation and Community Structure in Artificial EcoSystems," *Artificial Life* 1 (Fall 1993/Winter 1994), 15-37.

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

[15] D. McFadzean and L. Tesfatsion, "A C++ Platform for the Evolution of Trade Networks," *Computational Economics*, Special Issue on Programming Languages, to appear.

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

[17] A. Roth and M. A. Oliviera Sotomayor, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge University Press, Cambridge, 1990.

[18] M. D. Smucker, E. A. Stanley, and D. Ashlock, "Analyzing Social Network Structures in the Iterated Prisoner's Dilemma with Choice and Refusal," University of Wisconsin-Madison, Department of Computer Sciences Technical Report CS-TR-94-1259, December 1994.

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

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

[21] L. Tesfatsion, "A Trade Network Game with Endogenous Partner Selection," pp. 249–269 in H. Amman et al. (eds.), *Computational Approaches to Economic Problems*, Kluwer Academic Publishers, 1997.

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