

Department of Economics

ACE Models of Endogenous Interactions

Nicolaas J. Vriend

Working Paper No. 542

May 2005

ISSN 1473-0278



ACE Models of Endogenous Interactions *

Nicolaas J. Vriend, *Queen Mary, University of London, UK*

May 2005

Abstract:

Various approaches used in Agent-based Computational Economics (ACE) to model endogenously determined interactions between agents are discussed. This concerns models in which agents not only (learn how to) play some (market or other) game, but also (learn to) decide with whom to do that (or not).

J.E.L. classification codes: C6, C7, D1, D2, D3, D4, D5, D6, D8, L1, M3

Keywords: Endogenous interaction, Agent-based Computational Economics (ACE).

N.J. Vriend, Queen Mary, University of London, Department of Economics, Mile End Road, London E1 4NS, UK, <n.vriend@qmul.ac.uk>, www.qmul.ac.uk/~ugte173/.

* I would like to thank Bob Axelrod, Josh Epstein, Giorgio Fagiolo, Nobu Hanaki, Marco Janssen, Jerry Silverberg, Leigh Tesfatsion, participants of the ACE Workshop at Michigan, and two referees for helpful discussions, comments and suggestions. The usual disclaimer applies.

1. Introduction

This paper presents an overview of Agent-based Computational Economics (ACE) models of endogenously determined relationships. This concerns models in which agents not only (learn how to) play some (market or other) game, but also (learn to) decide with whom to do that (or not). Such decisions may depend, for example, on the perceived success of the interactions. These models of endogenous interactions are to be distinguished from models in which the interactions between agents are exogenously determined; for example by the given spatial positions of agents, such as with cellular automata. An alternative way to put this is that in the models with endogenous interactions discussed in this paper, the speed with which connections can be updated is comparable to (or faster than) the speed with which strategies in some underlying game can be updated, whereas in models with exogenously determined interactions the speed of the network updating is so low that the interaction structure can be taken as given.

The main motivation for studying models of endogenous interactions is that endogeneity is a ubiquitous feature of the reality of social interactions. Therefore, a theory of social interactions must take account of it. Consider, for example, the following quote concerning market organization.

“Markets rarely emerge in a vacuum, and potential traders soon discover that they may spend more time, energy, and other resources discovering or making a market than on the trade itself. This predicament is shared equally by currency traders, do-it-yourself realtors, and streetwalkers! Their dilemma, however, seems to have gone largely unnoticed by economists, who simply assume that somehow traders will eventually be apprised of each other’s existence - to their mutual benefit or subsequent regret” (Blin (1980), p. S193).

Therefore, models of market organization going beyond assumptions of perfectly competitive markets (either considering them as black boxes, or with Walrasian auctioneers or invisible hands pulling the strings), explicitly focusing on the "who interacts with whom?" question seem useful. The endogeneity of interactions is equally ubiquitous in other social domains. As Skyrms & Pemantle (2000) observe:

“A child who is being bullied learns either to fight better or to run away. Similarly, a player who obtains unsatisfactory results may choose either to change strategies or to change associates” (p. 9340).

We will focus on ACE models. The basic idea of ACE modeling is that one computes explicitly (either with paper and pencil, using a computer, or just mentally) the actions and outcomes for each and every individual agent at each relevant moment in time. Modeling individual agents computationally does not pose particular

conceptual difficulties to economic theory. After all, as Lucas puts it, doing economics means "*programming robot imitations of people*" (in Klamer (1984), p.49). In fact, ACE modeling follows the same methodology of scientific inference as more traditional mathematical modeling in economics, and should be seen as complementary rather than an alternative to such more standard modeling. Both are modeling approaches using equations and deduction. With standard mathematical modeling in economics one typically specifies a certain micropattern (primitives and rules of possible interaction) and then considers a macropattern as an equilibrium of the thus specified model. In this traditional view, if a certain macropattern is not an equilibrium of such a microspecification, then it is not explained. This is what Varian (1984) calls "*recoverability*" (p. 3). The same applies to ACE modeling, where one also focuses on the question whether it is possible to 'recover' regularities known from reality in relatively simple models (abstracting from many aspects of reality), and analyzes how these regularities depend upon parameter choices or modeled mechanisms. Both ACE models and formal, mathematical models are thus models that are in themselves possible explanations for some real phenomena. Whereas the insights offered by mathematical models are typically presented in the form of theorems or propositions, ACE models seem to produce only computational examples. However, as Judd (forthcoming) explains, even in this respect the two approaches are similar because 'theorem' is just a plural of 'example'. Although examples are produced in a somewhat different way in ACE models, there is no fundamental difference in this respect.

The difference between ACE and more traditional mathematical modeling in economics is a matter of the tools and techniques used: mathematical equations and specifications versus computational instructions. This facilitates different microspecifications and different ways to generate macropatterns, which in turn allows for different types of analysis, addressing somewhat different questions. For example, paying more attention to dynamic and non-equilibrium phenomena such as bounded rationality and learning, while maintaining tractability might be easier with ACE models. And as Tesfatsion (forthcoming) explains, this makes ACE similar to constructive rather than classical mathematics.

Sometimes in the literature one can find people comparing the output of an ACE model run on a computer to the data of laboratory experiments with real (human) subjects, as if such a run were a test of some hypotheses, and as if the computer output were data to be explained. Such a view does not seem very helpful. As explained above, an ACE model is a model as much as a more traditional mathematical economics model is. That is, the computer program in itself is the (possible) explanation of some real phenomena. Running an ACE model on a computer (no matter whether this is called a simulation, a computational test-bed, a wind-tunnel experiment, or an artificial petri dish) is *only* a matter of analyzing the model, checking its internal consistency and examining its properties.

The studies presented in this paper may differ in a number of important aspects from each other. (i) In the way connections are formed in a technical sense (e.g., by sending a communication signal, making a trip to a store, being a neighbor, etc.), and whether there are constraints formed by some underlying topology for the connections (e.g., a lattice). (ii) The way connections are evaluated and established in an economic sense (e.g., as myopic best-replies, or based on some learning process). (iii) The type of game (if any) being played for a given interaction structure or network. (iv) The way agents decide upon their strategies in such a game (e.g., as myopic best-replies, or based on some learning process). (v) Whether the focus is on the emerging interaction structure, or on the emerging strategies used in the underlying game (e.g. the trade-off between risk- and payoff-dominance in coordination games, or the sustainability of cooperation in prisoner's dilemma games). We will focus on the first two of these dimensions, i.e., on the various ways to model the endogenous interactions themselves.

<u>mechanism</u>	<u>paper</u>	<u>section</u>
random		
local		
residential pattern	Schelling (1971)	2.1
resource gradient	Epstein & Axtell (1996)	2.2
predictors	Arthur (1994)	2.3
advertising/patronage	Vriend (1995)	2.4
(threshold) expected payoff	Ashlock et al. (1996)	2.5
arbitrary tags	Riolo (1997)	2.6
trust	Hanaki et al. (2004)	2.7
expected payoff/familiarity	Kirman & Vriend (2001)	2.8
past success rate	Chang & Harrington (forthcoming)	2.9
directed random search	Jackson & Rogers (2004)	2.10

Table 1. Different ways to model interactions in ACE models

Table 1 lists a number of different ways used in the ACE literature to model interactions. The first two ways to model interactions are relatively well-known and straightforward. Considering random interactions has been popular in particular in work originating from evolutionary game theory. Local interactions have often been modeled in the form of interactions with nearest neighbors, e.g., on a grid or lattice. Notice that in these first two approaches the interactions are not endogenous. Instead, they are determined through some exogenous random process or through exogenously determined locations of the agents. Therefore, we will focus on the other approaches listed. In all these approaches, the agents themselves decide whether to establish, maintain, or sever a link with some other agent(s), and these

decisions are usually somehow related to the perceived success of their interactions. In the remainder of this paper we present an ACE paper (see table 1) for each of these ways to model the endogenous determination of interactions.

The overview will focus on the modeling of interactions as such, and will not provide a complete summary of the papers. The prime objective, rather, will be to catalogue the ACE ways to model endogenous interactions. We will also not attempt a comparison to find the best (elements of each) approach, but rather we would argue that the choice of model should depend on 'circumstances' to be modeled in a broad sense, and on the purpose of the model. This includes issues such as the cognitive capabilities of the agents and the opportunities to use them (e.g., for interactions that are immediate, a fast and frugal way to guide interactions may be appropriate), the number of agents involved, whether the interactions are face-to-face (allowing for face recognition and use of simple physical cues as signals) or not, whether the interactions are anonymous or not, whether they are repeated or not, and whether trust is an issue or not.

The objective of this paper is not to attempt to reach a conclusion as to what interaction patterns typically emerge in models with endogenous interactions (e.g., does a fully connected network ever appear?), or whether there are any general differences in this sense between models with endogenous and models with exogenous interactions.¹ On the one hand, it seems much too early for such an attempt. On the other hand, as Wilhite (forthcoming) shows, the relevance of the interaction structure may be different for any different exogenous interaction structure, not to mention the differences among models of endogenous interactions. Similarly, in principle it could be that at some point we will be able to conclude that models with exogenous interactions are satisfactory approximations to the social reality of endogenous interactions. But for the moment any such conclusion would seem premature.

¹ For example, Oechssler (1997), Dieckmann (1999), and Mailath et al. (2000) show that for a certain class of coordination games endogeneity of interactions may directly affect the equilibrium selected.

2. Various approaches

2.1 Schelling (1971): residential pattern

Schelling (1971) presents a spatial proximity model of neighborhood segregation. Although Schelling does not actually use any computers, this must be one of the very first ACE models.

There are two versions of this spatial proximity model: a one-dimensional (1D) and a two-dimensional (2D) model. In the 1D model, individual agents are distributed along a line. An agent's position is defined relative to his neighbors only, and agents can always position themselves next to any agent. A given individual's neighborhood is defined as the k nearest neighbors on either side of him. Agents towards the end of the line will have fewer than $2k$ neighbors. Schelling's 2D model concerns a regular lattice with bounds, such as a checkerboard.² Each agent occupies one cell of the board, and each cell can be occupied by only one agent at a time. Unlike in the 1D model, there are also some free cells. The neighborhood of an individual agent is the so-called Moore neighborhood. For an agent in the interior of the board this consists of the eight cells directly surrounding his own location, with fewer neighbors for agents at the boundary.

In both versions, Schelling considers a finite number of individual agents, distinguishing two types of individuals. Each individual is concerned only with the number of like and unlike neighbors. More specifically, each agent wants, for example, at most 50% unlike neighbors; otherwise agents are indifferent.

--- Figure 1 about here ---

The time structure of the model is given in Figure 1. The starting configuration is created by randomly distributing equal numbers of agents of each type. The dynamics, then, are an iterative process of agents choosing myopic best-responses to the residential locations chosen by the other agents. At each stage all agents that are not satisfied are put in some arbitrary order. When an agent's turn comes, he moves to the nearest satisfactory position. Since in the 1D version all positions are relative only, he simply inserts himself between two agents (or at either end of the line), and his own departure does not lead to an empty position. In the 2D version, each agent who wants to move has to find an empty location. At the next stage a new list of unsatisfied agents is compiled, and so on. This process continues until no agent wants to move anymore.

The interactions are endogenous in the following sense. Individual agents choose their neighbors on the basis of the current residency pattern (neighborhood ratios of like and unlike agents). As individual agents move, this residency pattern evolves. In fact,

² See also Sakoda (1971), which is based on Sakoda (1949), for a very similar model of endogenous interactions.

there are two kinds of externalities with every move. A leaving agent changes the neighborhood ratios for his old neighbors, while a newly arriving agent modifies the ratios in his new neighborhood. In both cases, these externalities may be positive or negative (depending on the perspective of the agents affected). Agents choose their location directly in (myopic) response to the existing residential pattern right from the start, and there is no learning (e.g., to be forward-looking) in this respect. Notice that there is no further underlying game to be played. The only thing that matters to the agents is with whom they interact, i.e., the proportions of each type in their own neighborhood. In some of the other models we will see that the variable guiding the interactions is some intermediate variable, and the agents can learn how this intermediate variable relates to eventual payoffs.

--- Figure 2 about here ---

The random starting state is typically highly integrated. The usual outcome of the dynamic process is a highly segregated state, although nobody actually prefers segregation to integration. Figure 2 gives an example based on Schelling (1971), showing the initial (integrated) and final (segregated) state. Does the endogeneity of the interactions matter? Yes, it is all that matters. Many integrated equilibria exist (see Pans & Vriend (2003)). But externalities of each residential location choice (i.e., the endogenous interaction choice) lead to an unraveling process (i.e., further endogenously determined interactions) resulting eventually in segregation.

2.2 Epstein & Axtell (1996): resource gradient

Epstein & Axtell (1996) study a number of different social behavioral phenomena, ranging from simple gathering and consumption, to mating, cultural transmission, combat, trade, credit, and the spreading of diseases.

These phenomena are studied in a so-called 'sugarscape'. This space consists of a lattice in the shape of a torus (i.e., a 2D grid with each edge folded and connected to the opposite edge). At each site, sugar can grow at a given rate up to some maximum (which may differ from site to site). In the basic model, sugar is the only commodity, and individual agents need it to survive, while in the model with trade there is a second commodity, called spice, as well. Each site will generally have some sugar and some spice, and can be occupied by at most one agent at a time, with each agent occupying one cell. When an agent occupies a site, he increases his wealth by accumulating the sugar (and spice) available at that site. Each individual agent has a given metabolic rate, specifying how much sugar per time step is consumed (decreasing the agent's wealth, i.e., sugar holdings), a given maximum age, and lateral vision up to some given limit. In the variant with trade, agents have also a utility function specified, and they have a metabolism for both sugar and spice, needing both to survive. In other variations of the model, the agents may also have their sex, and an array of arbitrary cultural attributes specified. Each agent has at most four neighbors, comprising the agents occupying sites in his von Neumann neighborhood (i.e., the sites laterally adjacent to his own site).

--- Figure 3 about here ---

As explained above, Epstein & Axtell study a whole range of behaviors. They do this in a modular setup, in which forms of behavior can be added or taken away as one likes. A typical sequence of events, following the initialization of all individual agents and sites, is given in Figure 3. First, at all sites simultaneously, resources grow at the given rate up to the limit of that site. The agents, then, move sequentially (in random order). Each individual agent checks the sites within his field of vision, moves to the best available location therein (if more than one he selects the nearest), makes a record of his new neighbors (calling this his 'neighborhood'), and increases his wealth by collecting the available sugar, while decreasing his wealth through his metabolism. Agents reaching a negative wealth die. Once these basic modules have finished, the optional trade module can be executed, in which agents may trade sequentially (in random order). A trading agent places all agents on his 'neighborhood' list in random order, and processes this list sequentially, making one transaction (if possible) with each of his neighbors by exchanging sugar for spice (or the other way around). The amounts exchanged depend on their marginal rates of substitution (as defined by the agents' utility functions) such that each trade leads to a welfare improvement of all agents involved, and on a pre-defined bargaining rule. After this, the optional combat module may be used, in which all agents sequentially

(in random order) may combat with their neighbors. Finally, the optional agent replacement module may be applied, in which agents who have reached their maximum age, or have died for other reasons (e.g., due to lack of resources or combat), are replaced by new agents with random characteristics.

All interactions are endogenous in the following sense. In all variants of the model, all interactions depend exclusively on the location choices of the agents, and these location choices are guided only by the resource availability, i.e., the distribution of available sugar (and spice) on the landscape. As individual agents move around and harvest sugar (and spice), the pattern of resource availability evolves. As in Schelling (1971), agents choose best-responses to the existing resource pattern right from the start, and there is no learning (e.g., how to react to certain resource patterns). In the basic setup with only consumer-gatherers, there is only indirect interaction between the agents. In the variants studying also other types of behavior (such as sex, trade, and credit), there are direct interactions between the agents as well. These activities take place in 'networks'. But these networks are essentially the (one-step) lagged von Neumann neighborhoods,³ and these neighborhood choices depend only on the resource availability on vacant sites. For example, in the variant with trade, an agent does not take account of the potential gains from trade on a given location, and agents just trade with whoever turns out to be an accidental new neighbor.

--- Figure 4 about here ---

Given the enormous range of behaviors studied in the various modules, we will not try to summarize the results. For each of the modules interesting properties of demographic, economic and other phenomena emerge. What is more, they show that the behaviors of the various modules interact with each other. For example, the outcomes of the economic process are influenced through the demographic dynamics. Given that all interactions are essentially determined through the gradient of the resources in the landscape, the emerging properties are remarkable. Figure 4 gives an example for the model with trade as the only optional module. The figure shows the time series average trading price converging to the "market-clearing" level of 1, which is the emergent property of the model with only bilateral interactions determined through agents myopically following their resource gradient.

³ These networks are formed as follows. When an agent moves in to the nearest best available location within his field of vision, he records his neighbors. Subsequently these neighbors may move on themselves. When, e.g., the trade module is executed, an agent can only initiate trading with those people on his 'neighbor' list.

2.3 Arthur (1994): predictors

Arthur (1994) examines the importance of inductive reasoning, and illustrates this with the so-called 'El Farol' bar problem.

People like spending some time together in this bar, in particular on Thursday nights with Irish music, unless it is too crowded. More specifically, Arthur (1994) assumes that there is a fixed population of 100 agents, that agents enjoy spending time together in the bar if fewer than 60 people are present, but prefer to stay home if more than 60 show up. Hence, the question in the El Farol problem (and in similar coordination problems) is which agents will interact with each other, and how will they decide to do so (assuming that all agents make up their mind each time independently).

Each individual agent is modeled as follows. An agent has an individualized set of predictors in mind. Each predictor determines the expected number of people attending on the basis of a sample of the past weeks' attendance figures. For example, a predictor could be "the average attendance of the last four weeks", or "the trend in the last eight weeks (bounded by 0, 100)". The agent keeps track of the accuracy of each predictor, using the actual attendance figures.

--- Figure 5 about here ---

The time-structure of Arthur (1994) is the following (see Figure 5). The model starts with randomly drawing a set of predictors for each agent individually from an "alphabet soup" of predictors. At the beginning of each period, each agent chooses one of his predictors, the one he currently believes to be the most accurate one. Given the predicted attendance. An agent decides to go to the bar if and only if the predicted number is less than 60. The actual attendance figure determined by all these individual decisions is, then, used to update each agent's belief concerning the accuracy of his attendance predictor.

The interactions are determined endogenously as follows. The individual interaction decisions (whether to go to the bar or not) depend on the past pattern of interactions (attendance figures), as different patterns of past attendances will typically lead to different expected attendance figures and hence different interaction decisions for most given predictors. Through these individual interaction decisions, the pattern of interactions itself evolves, as they will form part of the interaction pattern on which future interactions will be based. What is more, the view of an agent as to how a given pattern of past attendance figures should lead to an interaction decision itself evolves. That is, the agents learn which predictor to use by updating their beliefs as to how accurate these predictors are, where this accuracy depends on the predictors used by the other agents.

--- Figure 6 about here ---

How does the interaction pattern evolve? Figure 6 shows the attendance figures for 100 periods. As we see, it fluctuates around 60% attendance. That is, in each period about 60% of the agents predict attendance below 60%, while another 40% forecast attendance above 60%. Obviously, these predictions cannot be all correct at the same time. Hence, individual choices fluctuate over time as well, not only because the attendance pattern fluctuates, but also because individual agents continue revising the accuracy of their predictors. As Arthur (1994) puts it: *"This is something like a forest whose contours do not change, but whose individual trees do"* (p. 410)"

2.4 Vriend (1995): advertising/patronage

Vriend (1995) presents an example of a computational approach to self-organization of markets, in particular buyer-seller networks. The starting point is the idea that market organization depends in a crucial way on knowledge of the identity of some potential trading partners. Such knowledge requires some kind of communication or interaction between the agents. Markets, then, emerge as the result of interacting individual agents pursuing advantageous contacts. The paper analyzes the emerging trading structure of the self-organized markets, the distribution of firm sizes etc.

--- Figure 7 about here ---

Each day, firms produce a certain commodity in advance, without knowing what the demand on the day will be. They may attract the attention of potential customers by sending information signals randomly into the population, directed at nobody in particular (presenting themselves as sellers to the population), and by offering a reliable service. Both production and signaling are costly. Consumers, then, have the choice to either 'shop around' randomly, stay loyal to their current supplier, or follow one of the information signals they received. Consumers want exactly one unit per day (at a given price), and shopping takes place on a first-come first-served basis. Figure 7 shows the structure of the model.

Each individual firm is specified as a set of alternative rules: binary strings, determining a production and an advertising level. The fitness of each rule depends on the actual payoffs generated using that rule, with fitter rules being more likely to be used. This is a form of reinforcement learning. After each block of 50 days, the sets of decision rules used by the individual firms evolve using a genetic algorithm: some rules are eliminated, while others are reproduced, with selection based on the fitness of the rules, applying crossover, and mutation. See also Brenner (forthcoming) and Duffy (forthcoming) on reinforcement learning and genetic algorithms.

Each individual consumer consists of a set of 15 "if ... then .." rules to decide how to shop: the conditions considered relate to the consumer's shopping experience during the previous day (whether he was satisfied, whether he was late and found only empty shelves, or whether he was simply lost in the mist and could not even find a firm selling the commodity), and to his information state (whether he did or did not receive any advertising signals from firms on this day). The possible actions for a consumer to consider are whether to patronize (return to the last firm visited), to visit one of the firms known to be selling this commodity through the advertisement signals, or to try his luck visiting somebody chosen at random. The fitness of each rule depends again on the actual payoffs generated using that rule, and fitter rules are more likely to be used in the future.

The interactions are endogenous in the following senses. The firms decide with how many people to link up through the number of advertising signals they send, and influence how many of these interactions are successful through their output

decisions. Their views on these decisions evolve as they gain experience about their profitability. This profitability depends also on the shopping behavior of the consumers. These shopping decisions depend on the success of their latest trip and on whether they received advertising signals. These variables may evolve as the result of decisions by the firms and other consumers, and the consumers' view on the importance of these two variables may evolve as well.

--- Figure 8 about here ---

What is the dynamic behavior of this model? First, starting from completely random behavior, all agents are relatively quick to learn reasonable behavior (production, signaling and shopping), leading to high efficiency and a good profit margin for the firms, while heterogeneous behavior emerges among consumers and firms. Second, does patronage occur, and what role does it play? As Figure 8 shows, especially 'strict patronage' (i.e., patronage by a satisfied consumer) emerges. That is, consumers quickly learn that in case they had been disappointed by a firm there is much less reason to return to that firm than in case of previous success. Notice that it is strict patronage that leads to the arbitrage of trading opportunities. For suppose some firms offer higher service rates than other firms. Strict patronage would imply that a firm not able to satisfy its clients is likely to lose some of its customers. Given its level of production, that would mean a higher coefficient of customer satisfaction on the next day. On the other hand, a firm satisfying its customers is likely to enlarge its clientele, thus lowering its service rate. Hence, *ceteris paribus*, strict patronage directly implies arbitrage of trading opportunities, in the sense of the equalization of service rates across firms. Third, does communication matter? Yes, but it is costly and directly related to the endogenous shopping behavior. As can be seen in Figure 2 in Vriend (1995) (showing the costs per unit of sales), the firms, having reached profitable decisions early on, then continue to increase their signaling level steadily, as they are competing with each other to attract the consumers through the advertisement signals, until some constant average level is reached with much lower profits for the firms. Thus, communication matters, although the firms have no (explicit) clue as to why they send such signals. They have no idea what governs shopping behavior. This is illustrated in Figure 11 in Vriend (1995), showing the average signaling level for two versions of the model: the standard version, and a variant in which consumers will always return to a firm after a successful trip (i.e., fixed patronage). Although the firms do not know anything about this, they immediately spot the difference in the value of advertising in the latter setup, avoiding it almost completely, whereas high signaling levels are reached in the standard version.

2.5 Ashlock, Smucker, Stanley & Tesfatsion (1996): (threshold) expected payoff

Ashlock, Smucker, Stanley & Tesfatsion (1996) study the effect of preferential partner selection in an evolutionary study of the prisoner's dilemma game. The Prisoner's Dilemma game studied is a standard two-player simultaneous-move game in which each player can decide to Cooperate or to Defect with the resulting payoffs being as follows: payoffs for mutual cooperation and mutual defection are 3 and 1 respectively, while a unilateral defector gets a payoff of 5, and the sucker payoff equals 0.

Each individual agent is modeled as a finite automaton (Moore machine), represented by a binary string. This string contains two parts. The first part specifies the agent's dynamic game strategy in the iterated prisoner's dilemma. That is, this specifies an agent's action in the first round plus his actions in later rounds, with the latter being dependent on the history of play up to that point. The second part determines the endogenous interactions of the agent (i.e., with whom this agent wants to play the PD game).

--- Figure 9 about here ---

The time-structure of the ACE model of Ashlock et al. (1996) is shown in Figure 9. For a given generation of agents, there are 150 rounds. In each round, each agent proposes to one opponent to play one round of the basic Prisoner's Dilemma game. All proposals are evaluated, and each accepted pair plays the game. After 150 rounds, the set of agents is evolved using a genetic algorithm. That is, depending on their performance, some agents are eliminated, while others are reproduced (applying crossover to recombine successful strings and mutation to induce some experimentation). The performance of an agent is measured by his fitness. This fitness equals the sum of payoffs received by an agent divided by the number of payoffs received. An agent receives a payoff either from playing a round of the Prisoner's Dilemma game, or from the refusal of another agent to interact with him (in the latter case the payoff will be 1.0). There is no payoff for an agent if he rejects himself somebody's offer to play. If an agent neither makes nor receives any offers to play in a given round, he receives a wallflower payoff of 1.6. The model considers 2000 generations.

The interactions are made endogenous as follows. Each individual agent keeps track of the payoffs realized with each other individual agent in the population (either from playing or from refusal by the other). An agent updates his assessment of another agent by taking a convex combination of his existing assessment and his very latest experience with that agent. Hence, this assessment is a weighted average of past payoffs, placing more weight on recent interactions. The initial expected payoff is 3 for each agent. When an agent makes a proposal to play the PD game, he will do so only to the best agent in the population, provided this agent is tolerable (see below). An agent receiving offers, on the other hand, will accept all offers from agents that

are tolerable. An agent is tolerable if and only if the expected payoff with that agent is greater than a certain threshold. This threshold forms part of the individual agent's string, and evolves in the genetic step, such that threshold levels leading to higher fitness are more likely to be reproduced. The initial thresholds of the individual agents are uniform randomly drawn between 0 and 3.

--- Figure 10 about here ---

What does this all imply for the organization of the interactions taking place? Notice that, through their individual threshold levels, the agents care about the payoffs to be expected from other individual agents. First, do agents learn to be picky in this respect? The answer is "yes". The average threshold level increases over time from a level of 1.5 to about 2.1. Second, does being picky matter? Again, the answer is affirmative. The average fitness level increases from a random initial level of 2.25 to a level just above 2.8. In a variant of the model, without allowing for endogenous interactions (which would be the same as having a fixed low threshold level), the average fitness reaches a level of about 2.3. This difference is due to changes in the ways in which agents interact. In particular, the option of refusal gives agents a way to protect themselves from defections without having to defect themselves. As a result, ostracism of defectors occurs endogenously, while parasitic relations are also observed. It is not true in general, however, that higher threshold levels will lead to higher average fitness. There is some risk with caring too much about with whom one will interact. That is, an agent's threshold level might be so high that no agent is acceptable anymore. As a result, only wallflower payoffs are received. Figure 10 illustrates this. The figure shows the frequency distribution over the fitness and threshold levels for all generations over 196 runs. In most cases we observe a high threshold going hand-in-hand with a high average fitness, but there are a good number of generations with a very high threshold and a low fitness. In those generations being too picky led to a breakdown of interactions.

2.6 Riolo (1997): arbitrary tags

Riolo (1997) studies the effects and evolution of tag-mediated selection of partners in populations playing the Iterated Prisoner's Dilemma (IPD) game, analyzing exactly the same basic Prisoner's Dilemma (PD) game as Ashlock et al. (1996).

An individual agent is modeled as a 5-tuple, the first three real-encoded parameters specifying his dynamic game strategy (whether to cooperate or not, conditional on the history of play), and the last two parameters determining the endogenous interactions.

--- Figure 11 about here ---

For a given generation, each agent has to find an opponent ten times. Each successfully matched pair plays a 4-round IPD game. Once this is all done, the set of agents evolves. That is, some agents are eliminated while others are reproduced, with selection based on the agents' fitness (depending on the payoffs realized), and with noise added to the parameter values to induce some experimentation. Figure 11 shows the structure of the model.

The interactions are made endogenous as follows. Each individual agent i uses some arbitrary tag τ_i in $[0, 1]$. This tag is some external label or (behavioral) characteristic that can be easily recognized by other agents. One could, for example, think of the tag τ here as a number written on an agent's forehead. When an agent needs to find an opponent, he first selects a possible opponent randomly. He, then, accepts this opponent on the basis of the similarity of their tags: probability (i agrees to play j) = $1 - |\tau_i - \tau_j|^{b(i)}$, where $|\tau_i - \tau_j|$ measures the absolute distance between the tags of the two agents and $b(i)$ is a parameter in $[0, 100]$ determining the 'pickiness' of agent i . For any given value of $b(i)$, agent i is more likely to interact with others the closer their tags are.⁴ The opponent carries out a similar evaluation simultaneously, and they will play the IPD only if both accept to do so. Otherwise an agent will randomly try another possible opponent. There are search costs (to be subtracted from an agent's eventual payoff) for each failed attempt to find an opponent. After four failed attempts, an agent will have to play against a randomly chosen opponent, who will have to accept. The tags τ_i are an element of the second part of the 5-tuple specifying an individual agent, and evolve in the 'genetic' step, such that tag values leading to higher fitness are more likely to be reproduced. The other element of the second part of the 5-tuple specifying an individual agent consists of the 'pickiness' parameter b . Notice that a high b implies indifference with respect to tags (the distance does not matter), whereas a low b implies that the agent is very picky (the distance

⁴ The similarity in the tag can be seen as a clue that the players can trust each other as they may have a common understanding of the situation. Thus, somebody might be reluctant to play a game with a person with a weird hairdo who does not wear a tie, unless this player happens to go through life without a decent haircut and a tie himself.

must be very small). The pickiness parameter $b(i)$, being part of the 5-tuple specifying an individual agent, evolves as well in the 'genetic' step, such that values leading to higher fitness are more likely to be reproduced. Notice that selection and reproduction take place at the level of the individual agents (each modeled as a 5-tuple). That is, strategies, tags, and pickiness with respect to tags all evolve together such that successful combinations are more likely to prosper.

Riolo's model has some similarity with models in which agents choose a location in space, and then interact with nearby agents. That is, the abstract tag signal can be seen as a location. Notice, however, that in Riolo's model part of the endogeneity concerns the agents' choices whether or not to care about distance. Such endogeneity seems less natural in space, where the economic importance of the distance is typically exogenously given. One could also imagine an evolving matching function as such, allowing, for example, agents to learn to play only against large distance opponents. Again, this seems more natural with arbitrary tags than in real space.

--- Figure 12 about here ---

What are the dynamics of this model to determine endogenous interactions? First, do tags matter? The answer is "yes". As Figure 1 in Riolo (1997) shows, for a given parameter value of $b=0.02$ for all agents, the use of tags leads to quicker and more stable cooperation (resulting in higher average fitness). It is only without the tags that we observe troughs in fitness levels due to systematic defections. The average fitness with tags fluctuates around the expected payoff for random behavior. Hence, what the tags seem to do is allow the agents to 'escape' from systematic defectors (through the evolving tag values). Second, if the parameter b is no longer exogenously fixed, will agents learn to care about tags (through the pickiness parameter b)? Figure 12 shows that this depends on the (indirect) search costs. The figure shows the evolution of the pickiness parameter b over the generations. If the population starts out caring about tags ($b=0.01$ initially) and there are no search costs, then the population continues to care about tags (see the - line in Figure 12). But if there are search costs, then the population slides into indifference with respect to tags (o line in Figure 12). If, however, the population starts being relatively indifferent with respect to tags ($b=2.00$), and there are no search costs, then the population may or may not evolve into one that cares about tags (see \square and Δ lines in Figure 12).

2.7 Hanaki, Peterhansl, Dodds & Watts (2004): trust

Hanaki et al. (2004) study a repeatedly played one-shot version of a standard Prisoner's Dilemma (PD) game. Just as Ashlock et al. (1996), they start from the observation that defection is the dominant strategy, and ask the question whether cooperation could be sustained with endogenously determined local interactions, with the individual agents choosing their (number of) partners.

Each individual agent can choose a strategy for the one-shot PD and he can revise his links with other agents. The agents are restricted to using the same PD strategy for their entire neighborhood (i.e., all partners they are linked to). An individual agent's payoffs are summed over all his interactions, and there are costs attached to interacting with other agents, with the costs increasing in the number of partners.

--- Figure 13 about here ---

The dynamics of the model are given in Figure 13. The model starts with a given number of agents being assigned random actions and beliefs (see below) and without any links between agents. Each period, all agents simultaneously play one round of the PD with all their partners. At the end of each period, with some exogenously given probability, individual agents can update their PD strategy, and with some other exogenously given probability they can update their local network, after which they play another round of the PD. When they update their interaction structure, with some exogenously given probability they either try to sever an existing link or to form a new link, and if this fails they try the opposite action. To choose a PD strategy, each agent copies the most successful strategy in his neighborhood (including his own), where the measure of success is the sum of all current payoffs. If an agent had no partners before, he randomly chooses to cooperate or defect in the next period.

The interactions are endogenous in the following sense. When looking for a link to sever, an agent chooses one of his existing partners randomly, and terminates the relationship if the net benefit of doing so is positive, myopically assuming other agents will not change behavior and the network remains otherwise unchanged too. No consent is needed. To form a link, on the other hand, consent is needed. Hence, both agents, myopically comparing costs and benefits, need to find positive net marginal payoffs. A potential new partner can be selected either among the partners of his current partners (with the probability of any partner being chosen proportional to the number of shared partners) or be a randomly chosen stranger from the entire population. The probability used to decide between these two routes to a new partner is exogenously fixed. That is, the agents do not learn which route to follow. How to estimate the expected payoff of a new partner? If it is a partner of a partner, this partner will inform the agent about the most recent PD strategy of this new partner. If the potential new partner is a stranger, the initial expected payoff depends on trust, which is effectively the subjective probability that such an agent will cooperate. This trust level itself evolves. That is, it is updated every period such that it is a weighted average of the

cooperation levels experienced by the agent (with greater weights for more recent experiences). The experiences that matter in this respect are either (in one version of the model) all others interacted with (including ongoing interactions) or (in another version) only all new partners interacted with. The agents' view on this does not evolve. That is, although the agents update their trust levels, they do not learn on which interactions they should base their trust.

--- Figure 14 about here ---

Hanaki et al. (2004) present an extensive analysis for a wide range of parameter values. Figure 14 shows the relative frequency distribution of average cooperation levels reached in 10,000 runs of 10,000 periods for a population of 1000 agents, with the other parameter values sampled randomly from pre-specified ranges. As we see, substantial amounts of cooperation can occur. They show that the amount (and volatility) of cooperation relies on networks being sparse (both globally and locally), which is facilitated by high connection costs. In a fully-connected network, where cooperators and defectors interact with the same agents, all agents would quickly learn to defect (which is the dominant strategy). They find that cooperation levels are higher when new partners are chosen at random from the whole population rather than friends of friends. On the one hand, this is due to the fact that relying exclusively on friends to find new partners leads to too high connectivity, and hence collapse of cooperation (see above). On the other hand, the advantage of interacting with strangers is that cooperation can expand. Friends of friends are acceptable only if it is known (through the friends) that they cooperate anyway. But strangers are acceptable depending on the trust of the agent looking for new partners. This trust (based on past experiences) is an imperfect substitute for information. As a result, such agents may seek interaction with current defectors, possibly leading to the recruitment of defectors. If these defectors are relatively isolated, they may immediately be converted (through the payoffs) to cooperation if the cooperating agent initiating the link has enough cooperators in his network. Obviously, this implies that assortive matching is essential, i.e., there must be a limit to this willingness to interact with defectors. The eventual amount of cooperation developed is the net result of these two forces. The balance is due to the endogenous trust level. That is, the amount of expansion is determined endogenously. Imperfectly informed agents are open enough for new contacts while all the time updating their beliefs about their environment.

2.8 Kirman & Vriend (2001): expected payoff/familiarity

Kirman & Vriend (2001) study the evolving structure of an actual market: the wholesale fish market of Marseille. They focus in particular on two stylized facts of that real market: price dispersion and the loyalty of buyers to sellers.

Each day the following sequence of events takes place in this model (see Figure 15). In the morning, before the market opens, the sellers purchase their supply for the day, without knowing the demand they will face during the day. The market, then, opens, and the buyers (who want one unit each of the fish) choose the queue of a seller in the market hall. The sellers handle their queues sequentially, giving each individual buyer a 'take-it-or-leave-it' price (thus, prices are not posted). Once the sellers have handled all queues, the morning session is over. In the afternoon, the market re-opens, allowing unsatisfied buyers from the morning sessions to choose again a queue of a seller. With all queues handled by the sellers, the market closes, and all unsold stocks perish. The buyers, then, re-sell their fish outside the market. The model considers 5000 days.

--- Figure 15 about here ---

Each individual seller must decide the quantity to supply, how to handle queues, and which prices to ask during the morning and afternoon sessions. For each decision they use a set of alternative rules. The fitness of each rule depends on the actual payoffs realized when using the rule, and fitter rules are more likely to be used again. An individual buyer chooses a seller in the morning, and possibly (another) one in the afternoon. Whenever a buyer hears a price, he will need to decide whether to accept or reject the price. For each of these decisions an individual buyer has a set of decision rules at his disposal, being more likely to use the fitter rules, with these fitnesses depending on the payoffs generated by these rules.

The interactions are endogenous as follows. The choice of which seller to visit for the buyers depends directly on the average payoffs a buyer realized with each seller, such that more satisfactory sellers (in the sense of offering a better combination of service and prices) are more likely to be visited by a buyer. When the sellers handle their queues, they can do this in any order they like. That is, they may give precedence to some buyers over other buyers. They do this on the basis of the familiarity of the faces of the buyers in their queue. This familiarity is basically a weighted average of past presences of a buyer in a certain seller's queue, and it evolves directly as the result of the buyer's shopping behavior. What is more, the sellers' view concerning the relevance of this familiarity may evolve. That is, an individual seller can move a more loyal buyer either towards the front or the back of a queue. The probability for a buyer to be served next is a function of a buyer's loyalty, and this function depends on a choice parameter, such that different values of this parameter give either more or less advantage or disadvantage to loyal buyers. The sellers learn which parameter value to use through reinforcement, such that values

that led to higher payoffs in the past are more likely to be used again. To decide upon a price to ask from an individual buyer, a seller takes into account the familiarity of the buyer's face too, as well as the remaining stock and remaining queue at that moment. Each seller uses a set of alternative rules linking these two factors to prices, and learns through reinforcement which rule to use.

--- Figure 16 about here ---

What kind of interaction pattern does this imply? First, does loyalty emerge? As Figure 16 shows, loyalty does emerge (on average). The loyalty index used is such that it would be 1 if buyers were perfectly loyal, and 0.10 if buyers were not loyal at all. As buyers do not even know the concept loyalty (they just pick a firm each day), and sellers are indifferent with respect to loyalty to start with, why do buyers become loyal? As it turns out, most buyers get a higher average payoff when returning to the same seller the next day than when switching. This occurs mainly through a better service rate of loyal buyers. Why do sellers offer this advantage to loyal buyers? Sellers realize higher gross revenues when dealing with loyal buyers, which is related mainly to a higher acceptance rate. Second, does this familiarity of faces matter? The answer is "yes", and the role it can play with respect to market organization is illustrated nicely by a setup in which there are three types of buyers. The difference between these three types is in the given prices for which they can re-sell outside the market (imagine, e.g., a cheap corner shop versus a posh restaurant). The model explains how 'high' buyers (those that can re-sell for a higher price) do not only pay higher prices than 'low' buyers, but also find higher prices than the latter. This happens notwithstanding the fact that in this model no trader knows about this difference between types of buyers, and no trader can recognize any type of buyer. But different types of buyers notice their different payoffs at the end of each day. This affects their evaluation of their price acceptance/rejection decisions, and their evaluation of the sellers they visited. Hence, this will influence their shopping behavior. These differences in shopping patterns are indirectly picked up by the sellers through the familiarity of buyer faces. In turn, this leads to different treatments in queues and different prices. What is more, differences among sellers emerge. Some sellers learn to specialize in 'high' buyers, some others in 'low' buyers. The latter ask lower prices, experience nevertheless a higher rejection rate, maintain a lower supply/sales ratio, leading to a lower service rate, and put loyal customers towards the end of the queue.

2.9 Chang & Harrington (forthcoming): past success rate

Chang & Harrington (forthcoming) study the issue of discovery and diffusion of knowledge, and the social networks that may thus arise.

They consider a population of individuals, who all have to solve the same given number of separate tasks. A method to solve a given task is described by a sequence of binary bits. Hence, each individual agent's method to solve his entire set of tasks is simply a binary vector of bits. Each individual agent has an optimal, target vector that describes the optimal way to solve all his tasks. Chang & Harrington (forthcoming) assume that individual agents, although they can not simply pick the optimal method themselves, can rank any two method vectors on their Hamming distance from their target vector (which is effectively the number of bits that is different in the two vectors). The individual target vectors may change over time following some pre-specified dynamic process. This implies that there is a persistent need for the individual agents to discover new methods, and for such knowledge to be diffused.

--- Figure 17 about here ---

The time-structure is the following (see Figure 17). Having drawn all initial method vectors and target methods randomly, and avoiding any bias in favor of imitation or innovation, or any bias favoring the observation of one individual over another, in each period each individual agent goes through the following sequence: An agent decides whether to innovate (all by himself) or to imitate another agent. If he decides to innovate, then he randomly chooses a method for a randomly chosen task. If, however, he decides to imitate, then he chooses an agent to imitate, and copies the method from this other agent for a randomly chosen task. In both cases (innovation and imitation), the method obtained is actually adopted only if its adoption gets the agent closer to his target vector (using the Hamming distance as measure).

The endogenous interactions are modeled as follows. The structure of interactions depends solely on the success of past interactions. This success depends on the distribution of the vectors of methods adopted by the agents as such, and their state relative to the target vector. But the agents do not use any information about this directly. As the agents interact (or not), the distribution of method vectors changes. The agents' view as to whether they should innovate or imitate (and if so, whom) evolves as the agents learn through their own experience. The choice of an agent between imitation and innovation (i.e., whether to interact with others or whether to stay alone) is a probabilistic decision. The decision weight for each depends essentially on the number of successes when choosing that option in the past. In addition, the weights decrease each period through some decay. In case an agent opts for imitation, the choice of the agent to be imitated is made in a similar probabilistic way. The weight for a given agent in the population is increased each

time that agent has been imitated successfully, and it decreases through some decay.

--- Figure 18 about here ---

Chang & Harrington (forthcoming) focus on the properties of the emerging social networks. In much of their analysis they partition the population into a fixed number of groups to get some persistent similarity in goals, as the dynamics of the target vectors of all agents within a group follow some stochastic process (modeling turbulence in the task environment) such that they stay within certain bounds. The individual agents know nothing about this, and one of the questions is whether they will imitate other agents, and if so to what extent they will learn to imitate agents from their own group or from other groups. The analysis is based on 20 runs of 20,000 periods with 20 individual agents. Figure 18 shows the interaction probabilities for all individual agents for a setup with four fixed groups of five agents. The 20 individual agents are ordered identically on both the horizontal and vertical axis according to their group association. Lighter shades indicate higher probabilities of interaction. Notice that individual agents cannot imitate themselves, which shows up as black diagonal cells. As we see, there are four 5x5 blocks that are clearly lighter, indicating that agents learn to interact more with agents within their group than with other agents. Notice that agents within their group are pursuing similar goals, but this is not known to the individual agents. Further analysis shows that this property is stronger the more groups are similar within and different from other groups.

2.10 Jackson & Rogers (2004): directed random search

Jackson & Rogers study some abstract network formation process, and in particular the question which processes may lead to power laws and small worlds, properties that have often been observed empirically in large networks. More precisely, they focus on the following three stylized features of such networks: highly clustered connections (which means that two agents linked to a given agent are themselves likely to be linked to each other as well), small maximal distances between nodes (which means that any two nodes in the network can be linked through a short path), and a power law in the upper tail for the distribution of node degrees (which means that there are more nodes with very few or with very many connections than one would expect if links were formed independently).

--- Figure 19 about here ---

Jackson & Rogers consider an abstract model of network growth, without any further economic interaction. Figure 19 shows the time-structure of the model. At each time step, one individual agent is added to the network. Before joining the network, the individual agent forms two samples of potential links. First, he creates a uniform randomly chosen sample out of all agents in the current network, and second, he forms another sample chosen uniform randomly out of all agents who are currently directly linked to the agents in his first sample. In the basic setup the net benefit of a link is independently and identically distributed across pairs. Given the two samples of potential links, the new agent myopically chooses any links within those samples providing him with positive net utility. Once linked, agents remain linked forever, and no further payoff relevant events occur for these agents.

The links formed are endogenous in the following sense. The choice of agents sampled by a new agent depends on the existing network structure. More precisely, the second sample consisting of partners of the first uniform randomly generated sample of agents is affected directly by the existing network structure. As new agents form links, this network structure evolves. The agents' view on the relevance of certain network structures, however, does not evolve, as the agents do not learn anything. That is, the sizes of the two samples, and the fact to use these two sampling methods are determined exogenously. This is irrelevant in the basic version of the model. Since the net utilities for pairs are independent and identical draws, any sampling method is as good as any other. But as soon as the net utility of connecting to a certain node depends on the existing network structure this is no longer true.

-- Figure 20 about here ---

The analysis of this agent-based model of Jackson & Rogers is in part formal mathematical, in part based on mean-field approximations, and in part computational.

They show that the model explains high clustering, which does not go to zero as the network grows (unlike a number of alternative network formation processes). This seems due to the search method, as any two nodes linked to by a new agent are likely to have been selected in part because they were linked to each other. The diameters of the networks tend to be small, which seems again related to the directed search method. As search is directed towards nodes with relatively large degree, new links are likely to shorten paths for many existing nodes. Finally, they show that the degree distribution of nodes has a scale-free upper tail. Scale-free means that the 'connectedness' (the distribution of links per node) does not vary with the scale of the network. This can be expressed by a power law as the probability of any given node being connected to k other nodes is $(1/k)^n$, where n is some constant parameter. Figure 20 shows a log-log plot of the complementary cumulative distribution function of node degrees. The solid curve is from a mean-field approximation, and the dotted curve from the computational analysis. The latter is based on a run of 10,000 periods, in which both samples were always of size two, and all agents sampled offer positive net payoffs (which means that they are acceptable links). As we see, the upper tail of the distribution is nearly linear, indicating a scale-free distribution, but the lower tail is not scale-free.

3. Concluding remarks

Although the models discussed cover a wide range of possibilities to model endogenous interactions, we can detect some kind of prototype of modeling endogenous interactions. This prototype seems to consist of up to three elements. (i) The interactions are directed (guided) by some variable x , e.g., because the agents are 'picky' with respect to this variable x . (ii) This variable x itself evolves directly as a result of the interactions. (iii) The agents' view of the relevance of variable x evolves, as they may learn, e.g., how and how much they care about it.

Where can the ACE modeling of endogenous interactions go from here? Obviously the approaches discussed could be improved, and alternatives may be created. It seems in particular interesting if various approaches could be used for the same underlying game or economic situation, to analyze the possible differences in dynamics. In the introduction we argued that the choice of interaction mechanism should depend on the 'circumstances' to be modeled in a broad sense, and on the purpose of the model. But it would seem interesting to add one level of endogeneity to the interaction mechanisms discussed in this paper, i.e., to let the type of endogenous interaction itself be determined endogenously. This would allow us to study why certain endogenous interaction mechanisms (depending on past payoffs, proximity, familiarity of faces, simple physical cues or tags, trust, advertisements, ...) seem to be relevant for certain types of interactions but not for others.

4. For further reading

Galouye (1964). For a start, consider the following quotes:

"We can electronically simulate a social environment. We can populate it with subjective analogs - reactional identity units. By manipulating the environment, by prodding the ID units, we can estimate behaviour in hypothetical situations." (p. 7/8).

And:

"... the simulator ... would ... be exploring fully the unpredictable fields of social interaction and human relationships as a means of suggesting a more orderly society, from the bottom up!" (p. 10).

References

- Arthur, W.B. (1994). Inductive Reasoning and Bounded Rationality. *American Economic Review*, 84, *Papers and Proceedings*, 406-411.
- Ashlock, D., Smucker, M.D., Stanley, E.A., & Tesfatsion, L. (1996). Preferential Partner Selection in an Evolutionary Study of Prisoner's Dilemma. *BioSystems*, 37, 99-125.
- Blin, J.M. (1980). Comments on "The Economics of Markets: A Simple Model of the Market-making Process". *Journal of Business*, 53, No. 3, Pt. 2, S193-S197.
- Brenner, T. (forthcoming). Agent Learning Representation: Advice on Modelling Economic Learning. In K.L. Judd & L. Tesfatsion (Eds.), *Handbook of Computational Economics. Volume 2: Agent-Based Computational Economics*. Amsterdam: North-Holland.
- Chang, M.H., & Harrington, Jr., J.E. (forthcoming). Discovery and Diffusion of Knowledge in an Endogenous Social Network. *American Journal of Sociology*.
- Dieckmann, T. (1999). The Evolution of Conventions with Mobile Players. *Journal of Economic Behavior and Organization*, 38, 93-111.
- Duffy, J. (forthcoming). Agent-Based Models and Human Subject Experiments. In K.L. Judd & L. Tesfatsion (Eds.), *Handbook of Computational Economics. Volume 2: Agent-Based Computational Economics*. Amsterdam: North-Holland.
- Epstein, J.M., & Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Washington, DC: Brookings/MIT Press.
- Galouye, D.F. (1964). *Counterfeit World*. London: Victor Gollancz.
- Hanaki, N., Peterhansl, A., Dodds, P.S., & Watts, D.J. (2004). *Cooperation in Evolving Social Networks* (mimeo).
- Jackson, M.O., & Rogers, B.W. (2004). *Search and the Strategic Formation of Large Networks: When and Why Do We See power Laws and Small Worlds?* (mimeo).
- Judd, K. (forthcoming). Computationally Intensive Analyses in Economics. In K.L. Judd & L. Tesfatsion (Eds.), *Handbook of Computational Economics. Volume 2: Agent-Based Computational Economics*. Amsterdam: North-Holland.
- Kirman, A.P., & Vriend, N.J. (2001). Evolving Market Structure: An ACE Model of Price Dispersion and Loyalty. *Journal of Economic Dynamics and Control*, 25, Nos. 3/4, 459-502.
- Klamer, A. (1984). *The New Classical Macroeconomics. Conversations with New Classical Economists and their Opponents*. Brighton: Wheatsheaf.
- Mailath, G.J., Samuelson, L., & Shaked, A. (2000). Evolution and Endogenous Interactions. In Pagano, U. and Nicita, A. (Ed.), *The Evolution of Economic Diversity*. Siena: Routledge.
- Oechssler, J. (1997). Decentralization and the Coordination Problem. *Journal of Economic Behavior and Organization*, 32, 119-135.
- Pancs, R., & Vriend, N.J. (2003). *Schelling's Spatial Proximity Model of Segregation Revisited* (Dept. of Economics Working Paper No. 487) Queen Mary, University of London.

- Riolo, R.L. (1997). The Effects and Evolution of Tag-mediated Selection of Partners in Populations Playing the Iterated Prisoner's Dilemma. In Bäck, Th. (Ed.), *Proceedings of the 7th International Conference on Genetic Algorithms* (pp. 378-385) Morgan Kaufmann.
- Sakoda, J.M. (1949). *Minidoka: An Analysis of Changing Patterns of Social Interaction* (Unpublished doctoral dissertation) University of California, Berkeley.
- Sakoda, J.M. (1971). The Checkerboard Model of Social Interaction. *Journal of Mathematical Sociology*, 1, 119-132.
- Schelling, T.C. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*, 1 (2), 143-186.
- Skyrms, B., & Pemantle, R. (2000). A Dynamic Model of Social Network Formation. *Proceedings of the National Academy of Sciences of the USA*, 97 (16), 9340-9346.
- Tesfatsion, L. (forthcoming). Agent-Based Computational Economics: A Constructive Approach to Economic Theory. In K.L. Judd & L. Tesfatsion (Eds.), *Handbook of Computational Economics. Volume 2: Agent-Based Computational Economics*. Amsterdam: North-Holland.
- Varian, H.R. (1984). *Microeconomic Analysis* (2nd ed.). New York: Norton.
- Vriend, N.J. (1995). Self-Organization of Markets: An Example of a Computational Approach. *Computational Economics*, 8, No. 3, 205-231.
- Wilhite, A. (forthcoming). Economic Activity on Fixed Networks. In K.L. Judd & L. Tesfatsion (Eds.), *Handbook of Computational Economics. Volume 2: Agent-Based Computational Economics*. Amsterdam: North-Holland.

1 run:

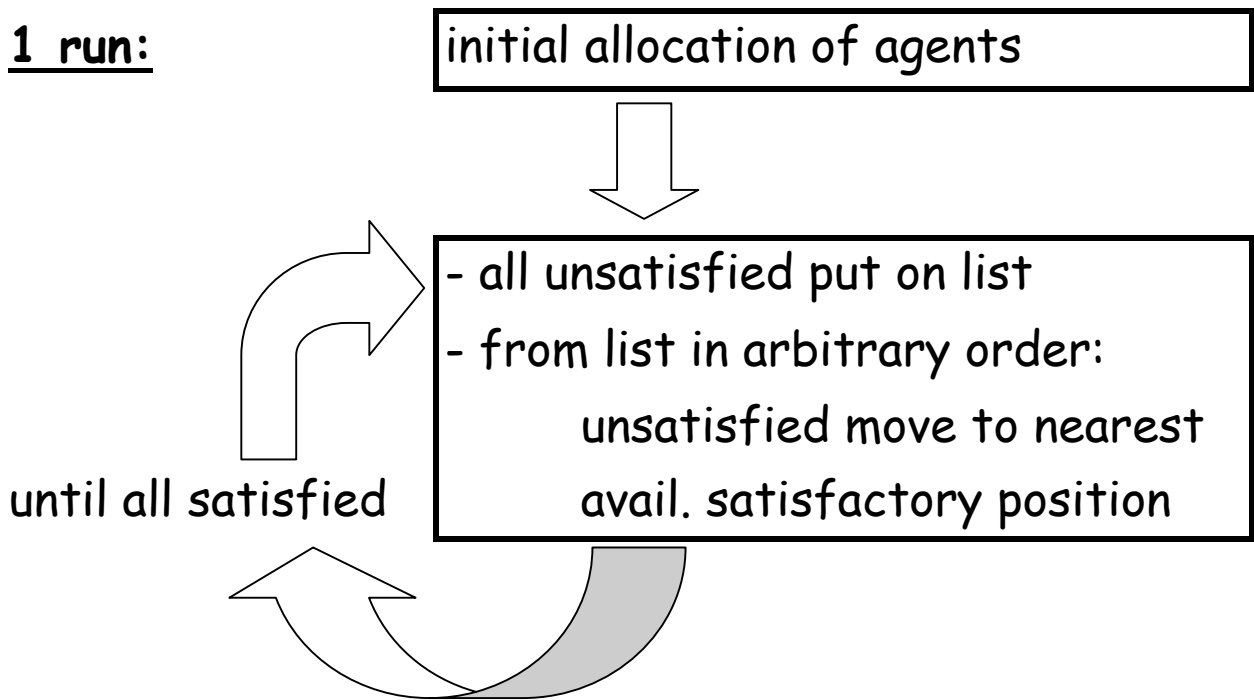




Figure 1. Structure of Schelling (1971)



initial state



steady state

Figure 2. Emergence of segregation
Source: Based on Schelling (1971)

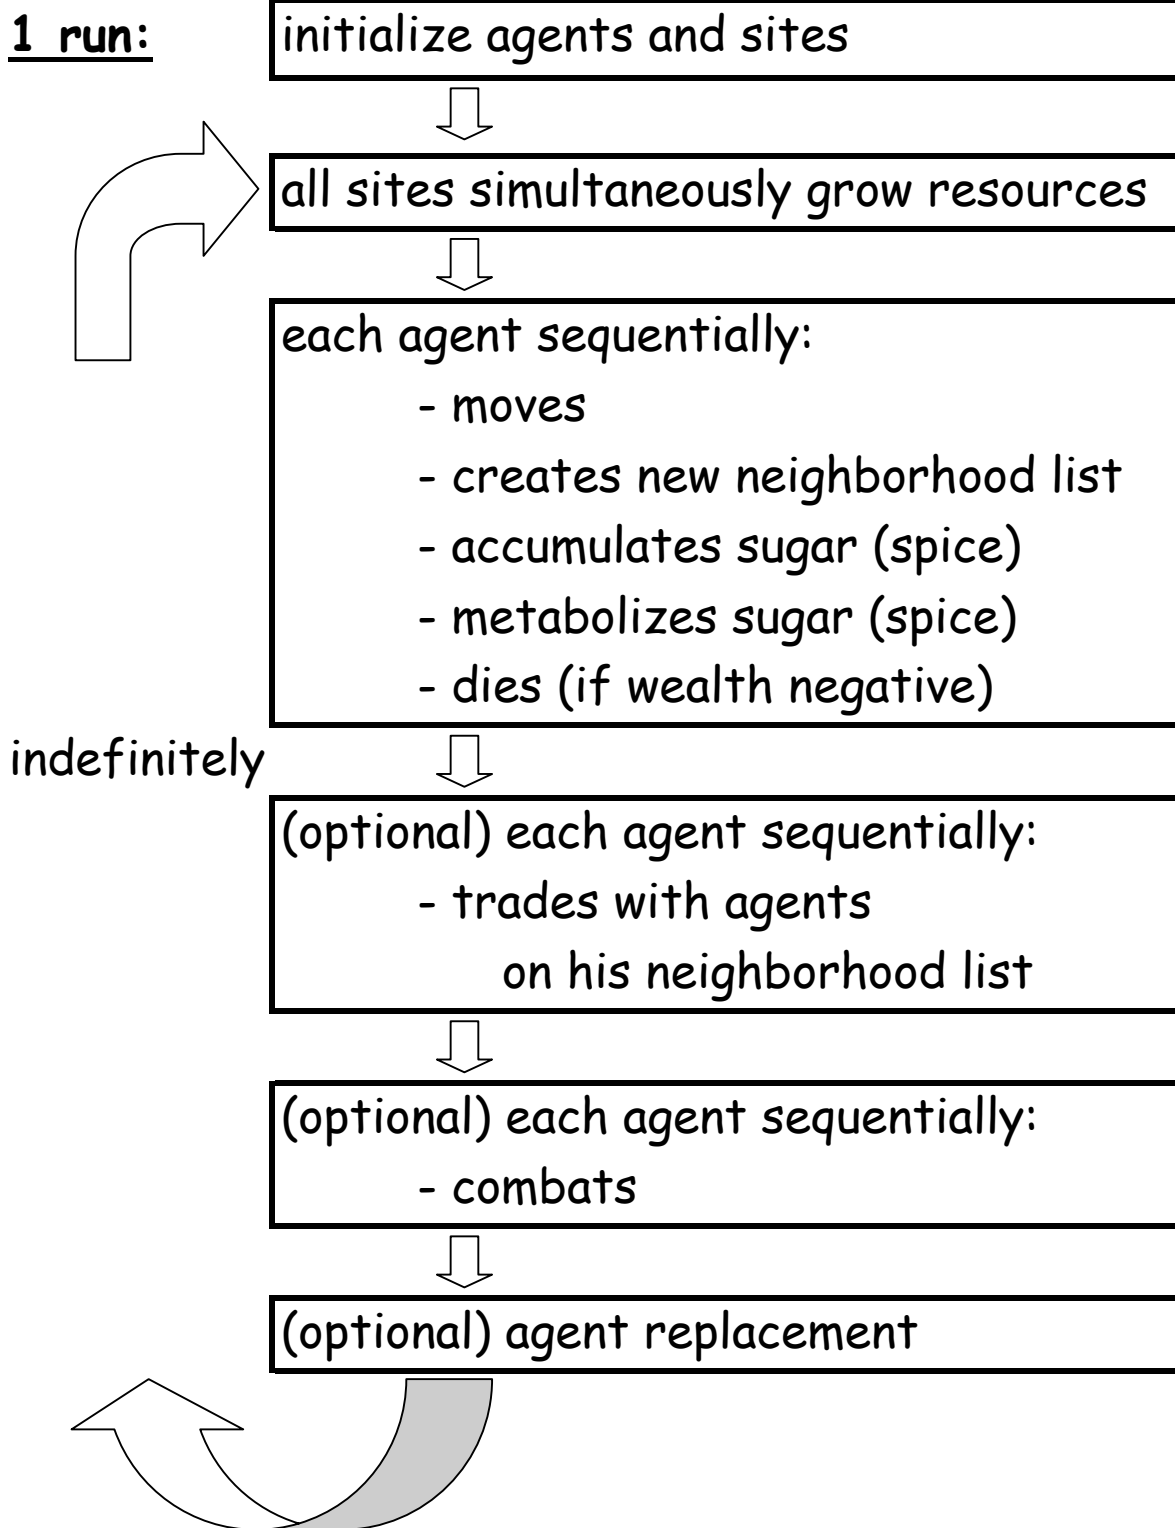


Figure 3. Structure of Epstein & Axtell (1996)




Figure 4. Typical time series for average trade price

Source: Epstein & Axtell (1996)

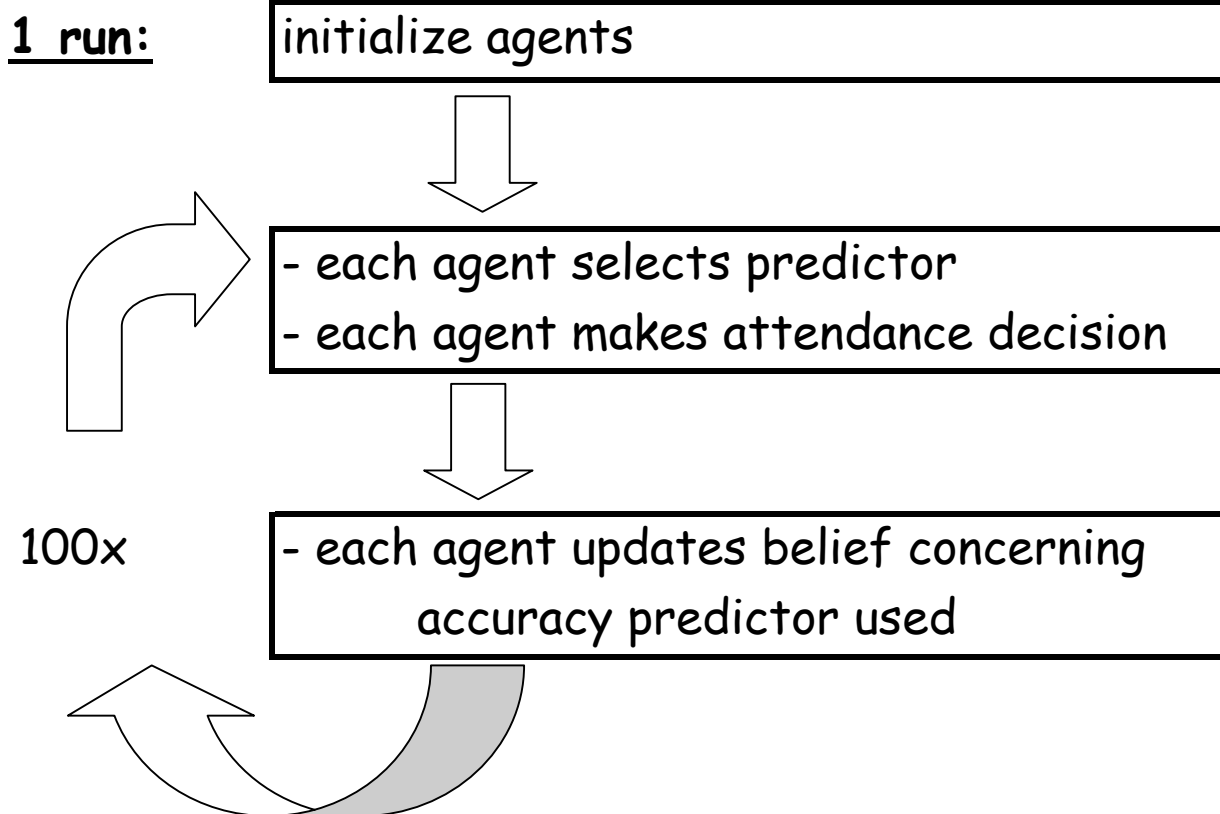


Figure 5. Structure of Arthur (1994)






Figure 6. Numbers attending in first 100 weeks

Source: Arthur (1994)

1 run:


initialize agents



50 days:

- firms produce in advance
- firms present themselves as such
- consumers shop (random order)
- reinforcement decision rules
(firms and consumers)

60x



- evolve rule sets indiv. firms
(selection based on fitness,
crossover, mutation)




Figure 7. Structure of Vriend (1995)






Figure 8. Evolution of patronage rate

Source: Vriend (1995)

1 run:


initialize agents



for 150 rounds:

- each player proposes opponent
- proposals evaluated
- each accepted pair plays one PD
- each player updates expected payoffs from opponents

2000x



- evolve set of players
(selection based on fitness,
crossover, mutation)




Figure 9. Structure of Ashlock et al. (1996)






Figure 10. Distribution of threshold and fitness levels


Source: Ashlock et al. (1996)

1 run:

initialize agents



for each player 10 times:
- player finds opponent
- matched pair plays 4-round IPD



5000x

- evolve set of players
(selection based on fitness,
with noise added)




Figure 11. Structure of Riolo (1997)




Figure 12. Evolution of tag-bias (pickiness) parameter

Source: Riolo (1997)

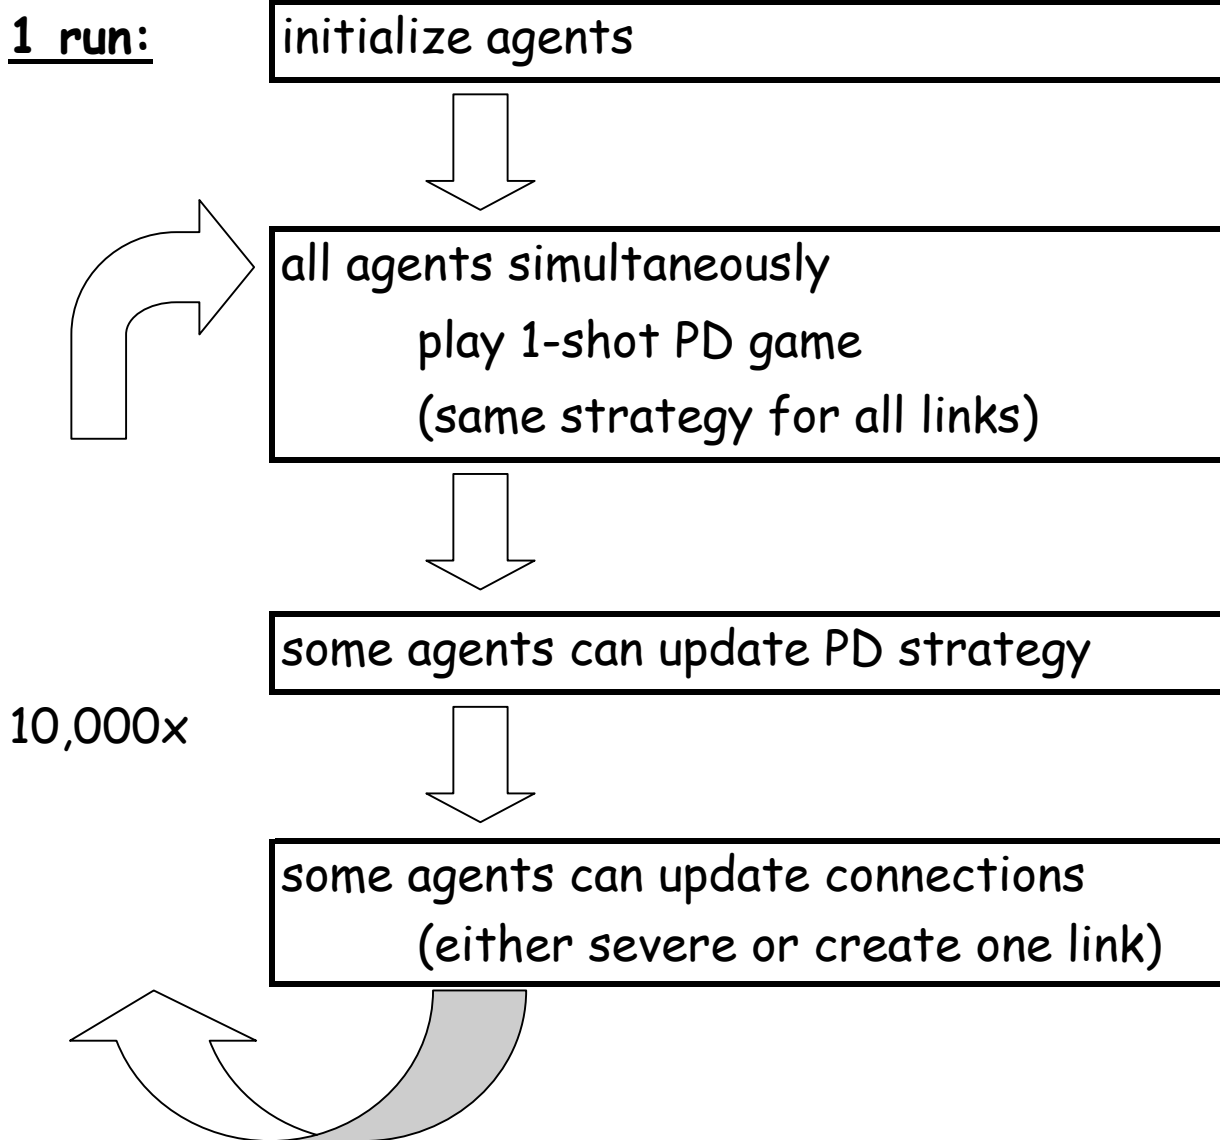


Figure 13. Structure of Hanaki et al. (2004)





Figure 14. Distribution of average cooperation levels reached

Source: Hanaki et al. (2004)

1 run:

initialize agents



morning:

- all sellers purchase supply
- market opens
- all buyers choose queue seller
- all sellers handle queues

afternoon:

- market opens
- all unsat. buyers choose queue
- all sellers handle queues
- unsold stocks perish
- all buyers re-sell outside market

5000x



reinforcement decision rules:

- sellers: quantities to supply,
how to handle queues,
which prices to ask
- buyers: choice seller,
prices to accept/reject




Figure 15. Structure of Kirman & Vriend (2001)






Figure 16. Evolution of loyalty

Source: Kirman & Vriend (2001)

1 run:


initialize agents and target methods



each agent:

- decides to innovate or imitate
- if innovate:
 - choose random method
 - for random task
- if imitate:
 - choose agent to imitate,
 - copy method random task
- adopts new method if better
- updates belief about
 - innovation v. imitation
- updates belief about
 - agent imitated (if any)

20,000x



stochastic change target methods




Figure 17. Structure of Chang & Harrington (2003)






Figure 18. Interaction probabilities between agents i and j

Source: Chang & Harrington (2004)

1 run:

initial state: nothing



one new agent to be added:

- randomly samples potential links
- creates second sample:
 randomly out of those linked
 to agents in first sample
- checks payoff with
 each sampled potential link
- creates all sampled links
 offering positive utility

10,000x




Figure 19. Structure of Jackson & Rogers (2004)




Figure 20. Distribution of node degrees

Source: Jackson & Rogers (2004)

**This working paper has been produced by
the Department of Economics at
Queen Mary, University of London**

**Copyright © 2005 Nicolaas J. Vriend
All rights reserved**

**Department of Economics
Queen Mary, University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 5096
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/papers/wp.htm**