Nicolaas J. VRIEND

ACE MODELS OF MARKET ORGANIZATION

Mots-clés : Organisation des marchés, interactions endogènes, approches computationnelles multi-agents.

Key words : Market organization, endogenous interaction, Agent-based Computational Economics (ACE).

I. — INTRODUCTION

The issue of market organization is presented clearly in the following quote : « *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]).

It seems fair to say that a 'market' is one of most central concepts of economics. The question, then, is : How are they organized ? How do buyers and sellers interact ? Figure 1 presents a straw man version of a market. The dots are individual traders, the box in the center represents the 'market' as such, and the arrows indicate market interactions. Observe that there is no direct interaction between the traders. All interactions go through the market. Individual traders receive price signals and return quantity signals to the market. As to the contents of this box, some economists prefer to leave it a black box, labeled 'perfectly competitive market', whereas others prefer stories about Walrasian auctioneers pulling strings, or invisible hands.

FIGURE 1 : Perfectly competitive market, case 1

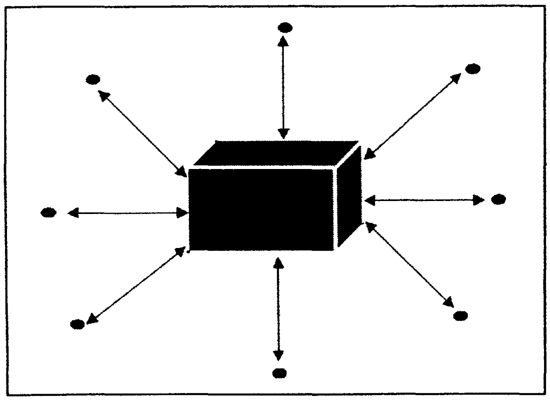

FIGURE 2 : Perfectly competitive market, case 2

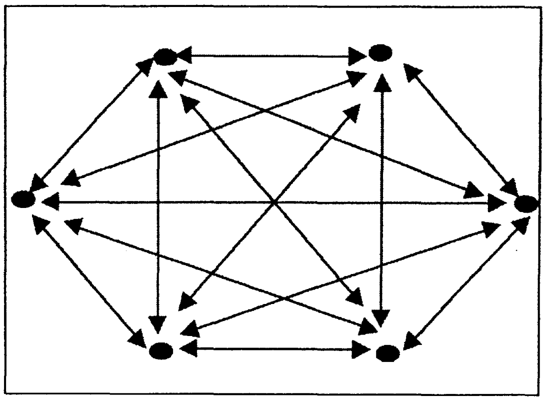

More recently, we have seen an advance of the theory of networks and network formation. It seems fair to say that this literature has paid relatively little attention to markets, or buyer-seller networks. Sometimes in this literature perfectly competitive markets are described as a very special network, in which each individual trader is directly linked to each other trader (see figure 2).

Thus, we have two extreme views of ideal markets. No interaction in one case, and full interaction in the other. Both do not seem particularly realistic, and in both market organization is assumed rather than explained.

If a market is not a star or a fully connected graph, then the crucial question is : « Who interacts with whom ? » In this paper we will focus on Agent-based Computational Economics (ACE) models addressing this question. The ACE literature presents, at least, six possibilities to model market organization (see table 1).

TABLE 1 : Different approaches to model interactions in ACE models

1	random	
2	local	
3	expected payoff individuals	Ashlock *et al.* [1996]
4	arbitrary tags	Riolo [1997]
5	advertising/patronage	Vriend [1995]
6	expected payoff individuals/familiarity	Kirman & Vriend [2001]

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 last four approaches. In all these four approaches, the agents themselves decide whether to establish, maintain, or severe a link with some other agent(s), and these decisions are somehow related to the perceived success of their interactions. In the remainder of this paper we present an early ACE paper (see table 1) for each of these four approaches to modeling 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. We will also not attempt a comparison to find the best (elements of each) approach, but rather we will argue that the choice of model should depend on market 'circumstances' in a broad sense.

II. — ASHLOCK, SMUCKER, STANLEY & TESFATSION [1996]

Ashlock, Smucker, Stanley & Tesfatsion [1996] study the effect of preferential partner selection in an Evolutionary study of the Prisoner's Dilemma. 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.

In later work (*e.g.*, Tesfatsion [1997]), it is argued that one could interpret such a game as a market game, but here we will not pursue that track. Instead, what we want to focus on here is the interaction mechanism as such. As we will see, this could easily be combined with a range of very different games.

Each individual player is modeled as a finite automaton (Moore machine), and represented by a binary string. This string contains two parts. First, a part specifying the player's dynamic game strategy in the iterated prisoner's dilemma. That is, this specifies a player'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. Second, a part determining the endogenous interactions of the player (*i.e.*, with whom this player wants to play the Prisoner's Dilemma game).

The time-structure of the ACE model of Ashlock *et al.* [1996] is shown in figure 3. For a given generation of players, there are 150 rounds. In each round, each player 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 players is evolved using a Genetic Algorithm. That is, depending on their performance, some players are eliminated, while others are reproduced (applying crossover to recombine successful strings and mutation to induce some experimentation). The performance of a player is measured by his fitness. This fitness equals the sum of payoffs received by a player divided by the number of payoffs received. A player

FIGURE 3 : Structure of the Ashlock et al. *[1996] model*

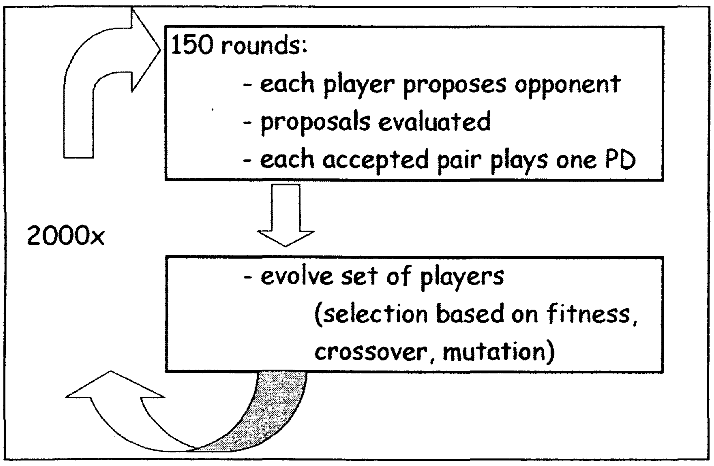

receives a payoff either from playing the Prisoner's Dilemma game, or from the refusal of another player to interact with him (in the latter case the payoff will be 1.6). There is no payoff for a player if he rejects himself somebody's offer to play. If a player neither makes nor receives any offers 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 player keeps track of the payoffs realized with each other individual player in the population (either from playing or from refusal by the other). A player updates his assessment of another player by taking a convex combination of his existing assessment and his very latest experience with that player. Hence, this assessment is a weighted average of past payoffs, placing more weight on recent interactions. The initial expected payoff is 3 for each player. When a player makes a proposal to play the PD game, he will do so only to the best player in the population, provided this player is tolerable (see below). A player receiving offers, on the other hand, will accept all offers from players that are tolerable. A player is tolerable if and only if the expected payoff with that player is greater than a certain threshold. This threshold forms part of the individual player's string, and evolves in the genetic step, such that threshold levels leading to higher fitness are more likely to be reproduced. Initially, the thresholds of the individual players are uniform randomly drawn between 0 and 3.

What does this all imply for the organization of the interactions taking place ? Notice that the players care about the payoffs to be expected from other individual players. First, do agents learn to be picky in this respect ? The answer is « yes ». Figure 10 (a) in Ashlock *et al.* [1996] shows the evolution of the threshold levels in the population. The threshold level increases over time from a level of 1.5 to about 2.1. Second, does being picky matter ? Again, the answer is « yes ». The same figure 10 (a) also shows the average fitness levels, increasing 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, the average fitness reaches a level of about 2.3. This difference is due to changes in the ways in which players interact. In particular, the option of refusal gives players a way to protect themselves from defections without having to defect themselves. As a result, ostracism of defectors occurs endogenously, while we also observe parasitic relations.

There is, however, also some risk with caring too much about with whom one will interact. That is, the threshold levels might be so high that no player is acceptable anymore. As a result, only wallflower payoffs are received. Figure 10 (c) in Ashlock *et al.* [1996] 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.

III. — RIOLO [1997]

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

An individual player 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.

For a given generation, each player has to find 10 times an opponent. Each matched pair plays a 4-round Iterated Prisoner's Dilemma game. Once this is all done, the set of players evolves. That is, some players are eliminated while others are reproduced, with selection based on then players' fitness (depending on the payoffs realized), and with noise added to the parameter values to induce some experimentation. Figure 4 shows the structure of the model.

The interactions are made endogenous as follows. Each individual player uses some arbitrary recognizable tag τ in [0, 1]. When a player needs to find an opponent, he 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 players and b(i) is a parameter in [0, 100] determining the 'pickiness' of a player. For any given value of b, players are more likely to interact with each other the closer their tags are. 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 PD game with a person with long hair who does not wear a tie, unless this player happens to go through life without a decent haircut

FIGURE 4: Structure of the Riolo [1997] model

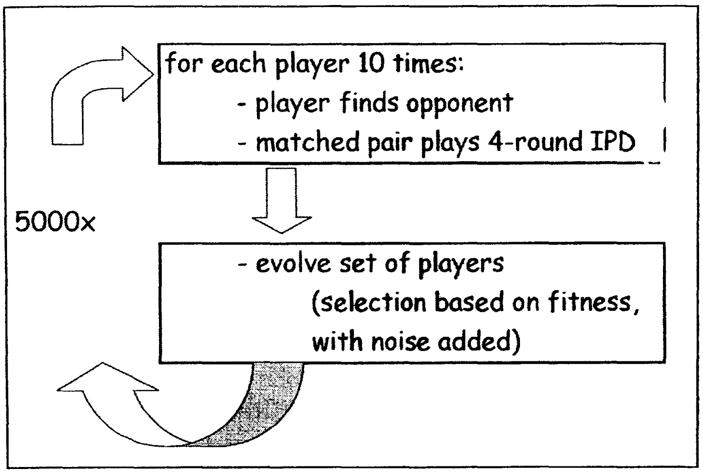

and a tie himself. Notice that a high b implies indifference with respect to tags (the distance does not matter), whereas a low b implies that the player is very picky (the distance must be very small).

The opponent carries out a similar evaluation simultaneously, and they will play the IPD only if both accept to do so. Otherwise a player will randomly try another possible opponent. After four failed attempts, a player will have to play against a randomly chosen opponent, who will have to accept. There are search costs (to be subtracted from a player's eventual payoff) for each failed attempt to find an opponent. The tags τ_i and the pickiness parameter b(i) form together the second part of the 5-tuple specifying an individual player, and both variables evolve in the 'genetic' step, such that values leading to higher fitness are more likely to be reproduced.

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 players, 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 players 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 2 in Riolo (1997) 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 if there are no search costs, then the population continues to care about tags (keeping b low). But if there are search costs, then the population slides into indifference with respect to tags (leading to high b values). 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.

IV. — VRIEND [1995]

Vriend [1995] presents an example of a computational approach to self-organization of markets. 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.

Each day, firms produce a certain commodity in advance, without knowing what the demand on the day will be, and firms present themselves as sellers to the population. They do this by sending a number of signals randomly into the population. Both production and signaling are costly. Consumers, then, shop around (in random order). Consumers want exactly one unit per day (at a given

FIGURE 5 : Structure of the Vriend [1995] model

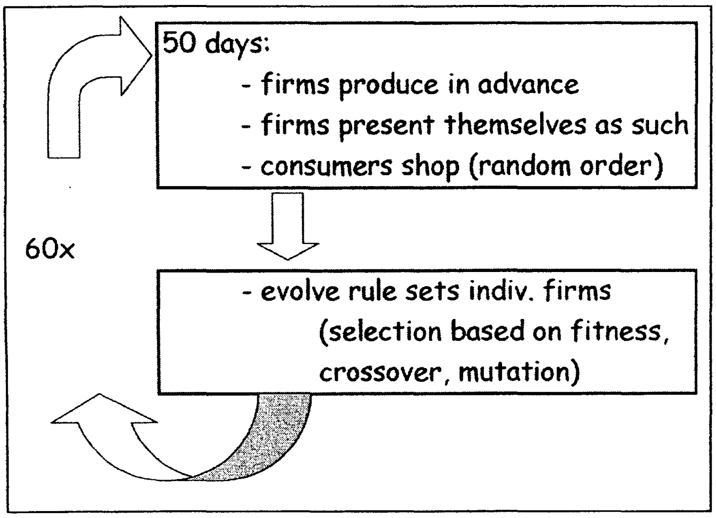

price), and shopping takes place on a first-come first-served basis.

After each block of 50 days, the sets of decision rules used by the individual firms (see below) 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. Figure 5 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. 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 a person 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 this evolves through the updating of the fitness of these rules, and through the genetic steps. The consumers decide how to shop, and this evolves through the updating of the fitness of the rules as well.

What is the dynamic behavior of this model? First, 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. But, as can be seen in figure 2 in Vriend [1995] (showing the costs per unit of sales), the firms, then, continue to increase their signaling level steadily, pushing their profits per unit of sales down, 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 firm 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]. The figure shows 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 very quickly spot the difference in the value of advertising. Whereas in the standard version a high signaling level is reached, in the variant signaling almost disappears completely almost immediately. Second, does patronage occur, and what role does it play? As figure 10 in Vriend [1995] shows, especially 'strict patronage' (*i.e.*, patronage by a satisfied consumer) emerges. Consumers quickly learn that in case they had been disappointed by a firm (arriving late) 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.

V. — KIRMAN & VRIEND [2001]

Kirman & Vriend [2001] study the evolving market structure of the wholesale fish market of Marseille, focusing in particular on price dispersion and the loyalty of buyers to sellers.

Each day the following sequence of events takes place in this model (see figure 6). 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 6: Structure of the Kirman & Vriend [2001] model

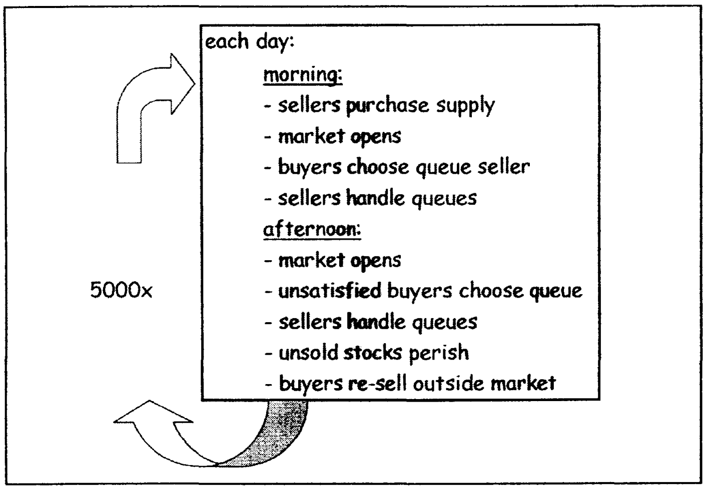

Each individual seller must decide the quantity to supply, how to handle queues, and which prices to ask during the morning and afternoon sessions. For each decision they use a set of alternative rules. The fitness of each rule depends on the actual payoffs realized when using the rule, and fitter rules 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, where this familiarity is basically a weighted average of past presences of a buyer in a certain seller's queue. 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 parameter b, such that different values of b give either more or less advantage or disadvantage to loyal buyers. The sellers learn which parameter value b 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.

What kind of interaction pattern does this imply? First, does loyalty emerge? As figure 6 in Kirman & Vriend [2001] shows, loyalty does emerge (on average). The loyalty index used there is such that it would be 1 if buyers were perfectly loyal, and 0.10 if buyers were not loyal at all. After a few thousand periods the loyalty reaches an average level of almost 0.8. 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.*, the difference between a cheap corner shop and a posh restaurant). The model explains that '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: 9.82 against 9.41 in the mornings, and 11.71 against 11.3 in the afternoons. This happens notwithstanding the fact that 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 pattern are indirectly picked up by the sellers through the familiarity of their faces. In turn, this leads to different treatments in queues and different prices. What is more, differences among sellers emerge. Some sellers specialize in 'high' buyers, some 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 at the end of the queue.

VI. — DISCUSSION

The four papers discussed use different ways to model endogenous interactions. But there are also some important similarities. Most prominently, the pattern of endogeneity of the interaction is similar in all four cases. The players can be 'picky' with respect to some variable (expected payoffs with individual players, tags, advertising signals, familiarity of faces). This variable itself evolves during the process of interaction. And the 'pickiness' of the players (whether and how they care about this variable) evolves too. There are two other similarities that we have neglected somewhat in this presentation. In all four cases, the fact that the interactions are endogenous seems to lead to heterogeneity of the population of players. What is more, the fact that the interactions are endogenous is costly, leading to lower average performance than when the interaction process is less endogenous.

Returning to the differences between the four approaches discussed, one obvious question is which approach seems best suited to model endogenous market interactions. All approaches discussed seem to play an important role in market organization, and their usefulness depends largely on the circumstances in a broad sense (including the nature of the commodity).

First, what assumptions about *cognitive capabilities* of the people can or do we need to make? In all four approaches, agents are relatively ignorant to start with, and can be seen to learn on the basis of experience. Hence, the models discussed are not demanding in this respect. The tags used by Riolo [1997] are a particularly simple and immediate ('fast and frugal') way to guide interactions. The familiarity of faces is very immediate too. Moreover, this familiarity itself develops automatically in the process of interaction, without the need to make any separate decisions.

Second, the *number of agents* involved plays a role. Keeping track of the payoffs experienced with each and every individual is a powerful way to discriminate between players if the number of agents involved is small, but it might be infeasible in large populations. Tags and advertising, on the other hand, work well in large populations too.

Third, recognition of faces requires *face-to-face* interaction. The same applies to tags such as hairstyle, clothing, smiles and handshakes. Simple cues may be sent by email too, and a so-called net-etiquette is evolving rapidly, but as some people observed: « A good glass of Chardonnay and a handshake are still worth a million bytes ».

Fourth, 'old fashioned' (say, late 20th century) shopping implies *anonymity.* As long as a buyer has money in his pocket and the goods are handed over, no further questions are asked about the identity of the seller or buyer. Advertising and tags work well with anonymity. Recognizing the familiarity of faces is possible too, although in itself it erodes the anonymity. Keeping track of the payoffs with individuals will not work with anonymity.

Fifth, *repeat interactions* are necessary to keep track of the payoffs with individuals, for loyalty or patronage to make sense, and for the recognition of faces to be possible. Repeat interactions are not needed to use tags or advertising.

Sixth, *trust* may be an important factor. Keeping track of individual payoffs, patronage, loyalty and the rewarding of familiarity are ways to establish trust in repeated interactions. Tags may be useful to establish trust even in one-off interactions.

Where can the ACE modeling of market organization go from here? Obviously the approaches discussed could be improved, and alternatives may be created. What seems important too, is to apply these approaches to more market games. It seems in particular interesting if various approaches could be used for the same underlying market game, to analyze the possible differences in market dynamics. Another interesting step would be to let the type of endogenous interaction itself be determined endogenously. Finally, it should be noted that in the end the type of endogenous interaction prevailing in itself shapes the nature of the market.

Voir références page suivante

REFERENCES

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.

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.

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.

VRIEND, N.J. [1995], « Self-Organization of Markets: An Example of a Computational Approach ». *Computational Economics*, 8, n° 3, 205-231.