

Agent-Based simulation as a useful tool for the study of markets

Juliette Rouchier

► To cite this version:

Juliette Rouchier. Agent-Based simulation as a useful tool for the study of markets. 2008. <halshs-00334051>

HAL Id: halshs-00334051

<https://halshs.archives-ouvertes.fr/halshs-00334051>

Submitted on 24 Oct 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GREQAM


**Groupeement de Recherche en Economie
Quantitative d'Aix-Marseille - UMR-CNRS 6579
Ecole des Hautes Etudes en Sciences Sociales
Universités d'Aix-Marseille II et III**

**Document de Travail
n°2008-08**

AGENT-BASED SIMULATION AS A USEFUL TOOL FOR THE STUDY OF MARKETS

Juliette Rouchier

April 2008



Agent-Based Simulation as a useful tool for the study of markets

JULIETTE ROUCHIER, GREQAM, MARSEILLE*

1	Introduction.....	1
2	Market and agents' reasoning	5
3	Buyer-seller interactions	11
4	Consumers, producers and chains	19
5	Financial market, auctions	27
6	Concluding remarks.....	37
7	Bibliography.....	38

1 Introduction

In recent years, there has been a growing recognition that an analysis of economic systems, which includes markets, as being complex could lead to a better understanding of the individuals' actions. In particular, one element that could be incorporated in analysis is the heterogeneity of agents and of their rationality. For example, the existence of multiple prices on a market for the same good, sold at the same moment and at the same place, cannot be captured in an equilibrium model, whereas it appears in real life and can be reproduced easily in an agent-based model (Axtell, 2005).

This issue is at the centre of much debate among economists. In classical economy, agents are considered as rational, having a perfect knowledge of their environment, and hence are homogeneous. There were of course discussions about this view of agents. It stayed unchallenged for a moment: for example Friedman (1953) argued that non-rational agents would be driven out of the market by rational agents, who would trade against them and simply earn higher profits. However, in the 60's, the view on rationality evolved. Even Becker (1962) suspected that agents could be irrational and produce same outcomes as rational agents (i.e. negative slope of market demand curve), but that the interest of

* **Acknowledgement:** I wish to thank Bruce Edmonds for his patience, and Scott Moss and Sonia Moulet for their advises.

describing agents as rational is that one needs not choose the way to represent irrationality. The author who definitely changed the view on economic agents is Simon, who stated that any individual could be seen as a “boundedly rational” agent, which means that it has an imperfect knowledge and has limited computing abilities (Simon, 1955). In most markets, agents have no perfect knowledge about the behavior and preferences of other agents, which makes them unable to compute an optimal choice to make. If they do have perfect knowledge, they will require unlimited computational capacity in order to calculate their optimal choices¹.

Indeed, for some contemporary authors, the understanding that one can get of real market dynamics is more accurate if the assumption of a representative agent or of homogeneity of agents is dropped, at the same time as the perfect knowledge assumption (Kirman, 2001). The way bounded rationality is approached can be very formal and tentatively predictive (Kahneman and Tversky, 1979), but some authors go even further by stating that the notion of rationality has been abandoned to be changed to the idea that agents possess rules of behaviors that they select thanks to diverse mechanisms, which are most of the time based on past evaluation of actions (Kirman, 2001). Some authors also compare the difference of results one can get when dealing with a representative agent approach and a bounded rationality approach for agent, and try to integrate both simplicity and complexity of these two points of view (Hommes, 2007). To complete the view that agents are boundedly rational in a market, and that they evolve through time, it is necessary to consider one more aspect of their belonging to a complex system: their interactions. The seminal works on markets with interacting heterogeneous agents date back to the beginning of the 90’s, with the

¹ For example, an issue that anyone representing learning (not only on market) has to face is the exploration-exploitation dilemma. When an action gives a reward that is considered as “good”, the agent performing it must decide to carry on with this action, and hence maybe miss other more rewarding actions, or to search, which implies indeterminate outcomes. Leloup (2002), using the multi-armed bandit (Rothschild, 1974) to represent this dilemma, showed that a non-optimal learning procedure could lead to a better outcome than an optimal – but non computable – procedure.

edition in the US of the Santa Fe institute book where many financial markets were represented (Arthur et al., 1997), which was a book gathering many authors having done this kind of work during the 1990's (Palmer et al., 1993). Since then, a large part of research in agent-based simulation concerns market situation, as can for example be seen in the cycles of WEHIA / ESHIA conferences and ESSA conferences. The former ones gather many physicists who apply dynamic systems techniques to representing heterogeneous interacting agents to deal with economics issues, and they tend to buy the “agent-based computational economics” (ACE) approach of simulated markets promoted by Leigh Tesfatsion² (Phan, 2003). In the later ones, not only economy, but also sociology, renewable resources management, computer science participate and try often to generate subtle representation of cognition and institutions and a strong view of agents as computerized independent entities to deal with broad social issues and being closer to multi-agent social simulation (MAS). Now we will refer to both terms (ACE or MAS) independently.

Not only the techniques can be different pursued when studying markets with distributed agents but also the aims can be. Some try to infer theoretical results about rationality and collective actions (Vriend, 2000) or about market processes (Weisbuch et al., 2001). Others want to create algorithms to represent human rationality on markets, and try to assess the value of algorithms that they use, comparing the simulated behavior with actions of real humans so that to understand the latter a bit more (Hommes and Lux, 2007 ; Duffy, 2001 ; Arthur 1994). Eventually some explorations about the impact of diverse rationalities in a market context enable the identification of possible worlds using a sort of artificial society approach (Rouchier et al., 2001).

Being part of a text book, this chapter should provide tools to be able to build and use agent-based simulation techniques to create artificial markets and analyze results. However, a “know-how” description of the building of artificial market is so dependent on the type of issue that is addressed, than it was decided here to establish a classification of the type of markets and modeling techniques that can be found in agent-based simulation research. We are interested in representations

² <http://www.econ.iastate.edu/tesfatsi/ace.htm>

of markets that focus on a micro-representation of decision processes, with limited information or limited computational abilities for agents, and which takes place in a precise communication framework. This form of representation forces authors to focus on new notions such as interaction or information gathering, and this induces new issues such as the bases for loyalty or the exploration-exploitation dilemma. The use of stylized facts is fundamental in this branch, where modelers try to mimic some elements of the world in their research, focusing either on representations of reasoning that are inferred from observed micro behaviors, or trying to mimic global behavior through learning algorithms, thereby to stay closer to common view orthodox economics.

Contrary to this view, which does not distinguish among individual rationalities and assumes aggregate, centralized knowledge and decision making, researchers involved in the use of multi-agent simulation usually try to understand the local point of view of agents, and its influence on global indicators. Since the agent is then seen as unable to have complete knowledge, it has to accumulate data about its environment and treat those data according to its aims. The study of markets is interesting when it comes to this type of analysis because markets display a much more simple set of possible actions and motivations than many other social settings. Income and reproduction of activity are direct motivations, which imply choices in the short and the long term, prices and quantities are what have to be chosen (as well as sometimes acquaintances, in the case of bilateral bargaining) and information is limited to offers and demands, as well as responses to these two types of proposals.

This chapter presents the main fields of application for multi-agent simulation dealing with markets, to show how researchers have focused on different aspects of this institution and to conclude on the great interest of agent-based simulation when trying to understand the very dynamics of these social environments.

We will describe the main notions that are covered by the term “market” in agent-based literature, and also the main ways to represent rationality and learning that can be found. Then, the main topics of market studies will be described in three parts. In section 3, agents are on a market and actually meet others individually, having private interactions with each other. Choices that have to be modeled are about matching, as well as about all buying or selling decisions. In all other sections, agents are facing an aggregate market and they have to make decisions

based on global data, sometimes associated to networks. In part 4 agents are either consumers or producers in a large market. In part 5 we will deal with auctions, financial markets and market design.

2 Market and agents' reasoning

In recent years the study of market(s) has been predominantly carried on by economists. Ethnographers and sociologists have been active in this domain (Geertz et al., 1979; White, 1988), and the field is now being developed through field studies and analysis. The main difference for those two approaches is that economists generally build models which are rather abstract and formal, whereas ethnologists and sociologists describe actual markets after observing them and generally produce models that are based on classifications of large amount of data. The notion of market itself has a double meaning, even more important with the increasing use of internet: it is at the same time the institution that enables individuals to coordinate their actions through the fixing of a price or, alternatively, a physical place where buyers meet sellers. There is no easy decision to choose the scale to study when dealing with market, neither is the limit of observation that is needed in the supply chain to understand a phenomena easy to set. In agent-based simulation, markets are represented as closed societies with a specified of agents (market being open or closed to new entry) that are possibly interconnected.

Simulations can be based on very specific case studies and in order to describe as accurately as possible the behavior of real actors, but they can also be mainly theoretic and in an attempt to generate expected theoretical results. In all cases, a simulated market cannot be implemented as described in neoclassical economic theory since agents need to be independently specified and interact directly with one another during the simulation. For example, to develop a model where individual agents have to make a decision, a demand curve (that gives for any price of a good, the number of agents that are ready to buy or sell at this price) cannot be imposed on the model but has to be derived from the determinants of agent behavior. One approach is to distribute reservation prices (the maximum price for buying or minimum for selling) among agents which can then be used to aggregate demand and offer curves.

The main elements to define a market model are given in the next section. We will then describe a few approaches to rationality and learning for agents in markets that depend on the type of market and goods that are represented. A similar analysis, more oriented towards consumers' behavior can be found in (Jager, 2007)

Main elements to build an artificial market

Several dimensions are important on a market, and each description of an element is easy to relate to a dimension in the building of an artificial system with agents (simply put: the market institution and the agents' rules of behavior). Axtell (2005) proposes a very abstract view of decentralized exchange on an agent-based market, where he gives no explanation of the bargaining process that organizes the exchange, but shows the existence of a computationally calculable equilibrium to increase all agents' utility. Here, the aim is to find out actual processes that can be used by modelers to represent markets.

The first classification that can be made in the building of a model is to know if one is building the representation of a *speculative market* or a *goods market*. What I call a speculative market, typically a financial one, is such that agents who have a good can keep it, sell it or buy it. They have to anticipate on prices and wait to perform their actions so as to make the highest profit. Seminal work on agent-based simulation were related to speculative markets, which display interesting regularities in their stylized facts. A large body of literature has developed on this topic, which is also due to the fact that data to calibrate models are more easily available than for non-speculative markets. On a goods market, agents have only one role, to sell or buy a certain number of units, and they usually have a reservation price to limit the prices they can accept. The good can be perishable (with an intrinsic value that decreases over time) or durable so that stock management is an issue in case of good markets.

Then, both types of market can be organized either through auctions (with diverse usual institution: double-auction, ascending, descending, with posted prices or continuous announcements) or via pair-wise interactions which imply face-to-face negotiation (with many different institutions, such as "take-it-or-leave-it", one shot negotiation, a series of offers and counter-offers, the possibility for buyers to explore several sellers or not). In the institutional design, it can also be important

to know if money exists in the system or if agents exchange one good for another directly.

Agents cognition must entail some choice algorithm. Agent-based simulation is almost always used to design agents with bounded rationality because agents have limited computational abilities or limited access to information. They have tasks to perform, within a limited framework, and have to make decisions based on the context they can perceive. Most of the time, they are given a function, equivalent to a utility function in economics, that associates a value to each action, enabling the agents to classify the profit it gets and hence to compare actions. First, an agent must have constraints in its actions, so that to be able to make an arbitrage between all possible options:

- Each good is associated to a *reserve price*: if a buyer (resp. seller) goes on a market, there is a maximum (resp. minimum) price it is willing to pay for the good.
- The importance of getting a good can be indicated by using a price of entry on the market. Agents have a greater incentive to buy and get a 0 profit, rather than not buying. The constraint for selling can be the same.
- In some papers, the price is not represented in the system, and the acquisition of a good is limited by a utility function, where the agent acquires a good if and only if it gives enough profit.
- In the case of negotiation, time constraint is usually put on buyers who can visit a limited number of sellers, having hence a limit on their search for good.
- There can be a discount factor: at each period, the risk of seeing the market close is constant and hence agents never know if they will be able to trade at the next period.

The type of decisions that agents have to perform on a market:

- for buyers: which number of units to buy, who to visit, how to decide to stay in a queue depending on its length, which price to propose / accept, which good to accept, and more fundamentally participate or not.
- for sellers: how to deal with the queue of buyers (first-come-first-served or with a preference over the identity of buyers), which offer to make or accept, and in the case of repeated markets how to decide the number of

units to buy for the next time step, in the case of a market where quality is involved: which type of good to propose or build for the next time step. Of course, because of the increasing complexity when adding one decision in a model, it is rare that all of these decisions should be made in one single model. For example, although interactions could potentially be represented in a continuous way, I know of no model where it is the case: all choices and meetings are made and messages sent at discrete time-steps.

Agents' learning

As said before, in most markets that are studied with agent-based models, the central element is that agents are heterogeneous in knowledge as well as in need. This situation can be decided from initialization or can emerge during the course of the simulation, while agents learn. Another element that is rarely given at initialization but is acquired by agents while learning is information about other agents' characteristics. In most cases, this learning takes place as a result of action, at the same time as the acquisition of an object or the acquisition of money.

The way learning is organized is generally linked to a performance of actions, with a selection of actions that “satisfice” or give the best performance. On a market, it is often assumed that agents are interested in getting the highest payoff for their individual actions: the performance is either the profit that agents get from their sells or the utility they get from consuming the product. In most models, learning agents have a set of pre-defined actions they can take and they have to select the one they like the best, following a probabilistic choice. One of the simplest learning models is reinforcement learning (Erev et al. 1999; Bendor et al. 2001; Macy and Flache 2002), where agents attribute a probability of choice for each possible action that follows a logit function. The algorithm includes a forgetting parameter and a relative weight attributed to exploitation (going to actions known as having high value) and exploration (the random choice part in the system). The exploration parameter can be fixed during the simulation, where the level of randomness has to be chosen, or can vary during the simulation (increase) so that there is a lot of exploration at the beginning of the simulation and as time passes agents focus on the “best” actions. This issue of which level of

exploration and exploitation to put in a learning algorithm is of current concern in the literature about markets (Moulet and Rouchier, 2007). Another model of rationality for agents is based on a representation of strategies of other agents, fictitious play: a distribution of past actions is built by each agent and can then infer the most probable set of actions of others, and hence choose their optimal behaviour (Boylan et El-Gamal, 1993). EWA model has been proposed by Camerer (Camerer and Ho, 1999) to gather both characteristics of these models: the agent not only learns what profit it got for each action, but also computes notional gains for each possible action, and attributes the resulting notional profit to each of those possible actions. A slightly more complex representation of knowledge commonly found in the literature is the classifier-system where each decision is made by considering a context, a choice and the past profit made in this precise context by making this special choice (Moulet and Rouchier, 2007). This type of algorithm is very similar to what (Izquierdo et al. 2004) call case-based learning, but it does not seem to be applied to market situations. In general the number of possible actions is fixed from the beginning, but the classifier system can be associated to a genetic algorithm that generates new rules over the time (M.KOPEL H. DAWID Olivier 97). Genetic Algorithm learning is also a quite usual way to represent learning, where the information that is used by agents to estimate the profit of each rule can be based on actual past actions (Vriend, 2000) or also on the imaginary profit of all possible actions considering the actions of others (Hommes and Lux, 2007). The presence of other agents can also be relevant information when agents use imitation or social learning. Brenner³ (2006) undertook an extensive review of usual learning processes. In this paper, an interesting element is developed: the way “satisficing” rationality can be developed, where agents do not look for the best action but for one which enables them to get a “good enough” profit, the notion of “good enough”, called “aspiration level” than can evolve during the simulation (Cyert and March, 1963). An alternative to learning algorithms that are only based on profit is to consider that agents have a social utility, which need not have the same dimensionality as

³ The main objection to Brenner’s exposition is the lack of homogeneity of notation, which makes the algorithms difficult to compare and maybe to implement.

profit. Indeed, it has been demonstrated that translating the social utility in costs or profits that can be added to monetary profit gives a dynamics of learning and behavior that is radically different from a situation where agents reason in two dimensions, social and monetary, with a lexicographic ordering (Rouchier et al., 2001). A way to implement social utility without lexicographic ordering is to include in the utility the appreciation of similarity of members of the network, such as in consumer choice models (Delre et al., 2007).

In most models, action and information gathering are made in one action, and circulation of information as such is not really modelled. The reason is certainly because it would take modellers too far from the neoclassical economic approach to market, where the only information is the observation of transactions, sometimes of intermediate prices as in auctions or, sometimes, bargaining. One model of market by Moss (2002) represents communication among agents before exchange takes place and Rouchier and Hales (2003) (which model evolved into the one of Rouchier (2004)) also allocates one period out of three every time-step for agents to look for information.

Indicators and method

Several points of view can be found in papers about markets, just like for any work in simulation. Some prefer to work at a purely abstract level and others try to fit as well as possible data that they extract from observation and experience. Whatever the approach, indicators that are often observed in markets are prices, efficiency (the total profit that is extracted from agents compared to the maximum profit that could be extracted), and relative power of different agents. The notion of convergence is central to the modeling of markets, since most research refers to economics and has to compare results to economic static equilibrium. In some cases, what is observed is interaction patterns, which can be represented as the random part of the choice of agents when interacting (the disorder), the number of different sellers that a buyer meets in a number of steps. In bargaining models and in general exchange models, the existence of an exchange or not is also something that is observed. Sometimes, the cognition of agents itself is observed: their belief about the others preferences, or even the distribution of propensities to choose sellers.

Among all the data that can be observed in the model, these last ones cannot be observed in real life, and hence cannot be compared to human systems. In a lot of models, agents' behavior is compared to that of humans in order to establish the validity of the cognitive model, expressed as an algorithm. The data are very rarely extracted from real life situation (although it sometimes happens), but are mainly constructed via experiments. Experimental economists control all information circulation and record all actions of agents. It is thus possible to compare in a very precise and quantitative way the global behavior of the group and individual behavior, on one side with artificial agents and on the other side with human agents. Real life situation can also be seen as the mix of human and of artificial agents, such as in financial online markets.

Other researchers do not want to match data too precisely. As Vriend says (Vriend, 2005), agents based models, like any other models, are abstract settings that have to be interpreted as such. The comparison between real and abstract data should go through a first step which is the building of stylized facts that are already a summary of human behaviors, where only the most striking elements are integrated. Vriend is much more interested by the reaction of his abstract model when parameters change, and by its self-consistency. It could be said that by construction, a model can only capture a small part of human cognition, which is full of long-term experiments and memories, and should not be compared to quantitative data without caution.

Eventually, some researchers want their model to influence real life and try to use the results they find to give advices on the way to build markets. Different ways to fit models with real life will be found in each example of a model – be it to fit precisely, to fit stylized facts, or to be an abstract study of the effect of some rules in a social setting.

3 Buyer-seller interactions

In the literature about agent-based markets, a great attention has been given to the analysis of local interactions, hardly ever studied in classical economics, apart from rare exceptions (Rubinstein and Wolinsky, 1985). An aim is to reproduce as well as possible the features of real markets. It is indeed to be noticed in real observation that buyers display regularity in the choice of sellers with whom to interact and that this regularity emerges in time with experience - this attempt to

reproduce patterns of interaction is the only way to understand, rather than just describe, the way individuals coordinate (Weisbuch et al., 2000). Authors study with great care local bargaining processes, as well as repetition of interactions over several days, where choices are not only based on prices but also on the fact that buyers need to buy and sellers to sell. The basic feature of these models is pair-wise interactions on markets with several sellers and buyers where prices are not posted and result from negotiation only. The number of visits buyers can make, the way sellers manage queues or the number of steps of the negotiations are different in all these systems that focus only on small parts of the whole set of stylized facts that are observable on such markets. Some aspects that are often studied are here described and the subsequent choices for modeling the organization of pair-wise interactions are given.

Bargaining processes

Brenner (2002) studies agents' learning in a bilateral bargaining market, focusing on the convergence of prices and the dynamics of bargaining. There is one good in the market, and buyers and sellers meet at every step to exchange this good. Each buyer can choose one seller for each market step, and it chooses according to the fact that the price is acceptable. The seller answers to buyers that wait in the queue in the order of arrival by proposing a price. Buyers have to decide who to visit; sellers have to decide on the first price to propose and then on the number of successive proposals it will make if the buyer rejects the offer, bargaining being costly for both agents. All decisions are made following *reinforcement learning* based on past experience. Hence, all choices are based on the satisfaction that is associated to each past action and on a rigidity variable. A buyer will carry on with choosing a seller as long as he is satisfied. His probability to change depends on his expectations with another agent. A seller also calculates the probability to change behavior depending on the belief he has on what he would gain by performing another choice.

The rigidity parameter, which is the opposite of noise in the system, has a great impact on results. If rigidity is high, buyers keep visiting the same seller. The cost of bargaining also is important: if it is not very low, sellers learn to offer the price they know to be acceptable to buyers, and they do not bargain after a few rounds. In this system, since the relations are so individual, the convergence of price is not

very fast and can be highly variable for a long time, although converging in the end. The model is extremely sensitive to all parameters that define aspiration levels for agents.

Brenner's paper is of the class that compares simulation to theoretical results. Here, sub-game equilibria are used to compare the possible outcomes and their efficiency to the generated prices and behaviors. There is no reference to any real world data. However, it is interesting that both micro behavior (the number of bargaining steps) and macro data (prices) are of importance, justifying an agent-based analysis.

Influenced by this paper, but referring to real data, Moulet and Rouchier (2007) reported a bargaining model based on two sets of data: qualitative, from a field study in the wholesale market of Marseilles by Rouchier (Rouchier and Mazaud, 2004), and quantitative, giving all proposals, intermediate and final prices for a series of transactions in the same market (Kirman et al., 2007). Like the previous model, the market gathers buyers and sellers who meet at every time-step. However, buyers can visit several sellers in one market opening. The choice of buyer has several dimensions: to decide which seller to visit, to decide to accept an offer or to reject it, to propose a counter-offer or leave, and then which value to counter-offer. Sellers must choose the first price to offer, to accept buyer's counter offer or not, and eventually decide on the value of the second offer they can make. In this model, decisions evolve following classifier system learning, where each agent evaluates a list of possible options following his past experience. The results that are produced are compared with indicators derived from real-world data: values of offers and counter-offers of the agents that vary depending of the kind of good that is purchased and *ex post* bargaining power of sellers (which is the difference between the first offer and the price of transaction compared to the difference between the counter-offer and the price of transaction). In the simulations, the values that are obtained fit the data, quite well in that it reproduces the bargaining sequences and agents' behaviors. The two main parameters are the number of sellers that agents can meet (from one to four) and the speed of learning of sellers. The relative importance of learning for the agents can be seen as situating them in a negotiation for in-season goods and a negotiation for out-of-season goods. The model produces results similar to those

of out-of-season goods when agents have to learn repeatedly, when there is no known market price but a lot of heterogeneity in the readiness to pay of buyers. In the case of in-season goods, the market price is more settled, and agents do not explore the possible values of goods as much, relying instead on their experience. Between the different in-season goods, the main difference could be the number of visits buyers make, but this number tends to reduce after a learning period, when buyers have selected their preferred seller. This aspect of the model – the growing loyalty of agents – is not the center of the research and was represented mainly with the aim of matching the actual behaviors of the market actors. Other papers, described in the following section, are more focused on this issue.

Another direction for the study of bargaining processes is related to the creation of robots or artificially adaptive agents (AAA) , to participate in electronic commerce (Oliver, 1997). Such models focus on negotiations that are complicated in terms of business situation in that they integrate several dimensions of trade in the deal: price, quantity and delivery time. The main argument for the value of the algorithm that is proposed in the paper is that the agents learn to negotiate at least “as well as humans”, which means that as many negotiations lead to an agreement as in human bargaining situations so that profit is extracted from both sides of the bargaining. The bargaining is constituted of several steps, where a customer reacts to the first offer by comparing its profit to a threshold, and the offer is accepted if it is higher than the threshold and rejected with a counter-offer that is made in the opposite case. Clearly, such models capture satisficing and bounded rationality rather than profit maximization. The bargaining can then carry on with several successive offers being made by customer and seller. Strategies for accepting and counter-offering evolve through a Genetic Algorithm. Five different games are used to test the learning, in a population of 20 agents with 3 rounds of bargaining at most and each agent is given 20 chromosomes for decision making. It is then proven that AAA perform better than random, that agents are able to learn general strategies that can be used against different bargaining partners, and eventually that AAA perform as well as humans (depending on the game, sometimes better and sometimes worse, maybe depending on affective values for humans) in terms of number of agreements that are reached. This is an interesting result to consider

when one wants to introduce artificial agents into electronic markets, since one wants then to be able to reach as many agreements as possible.

Loyalty

Loyalty is present in quite a few models of markets where agents interact repeatedly. It is popular to deal with this topic with agents, mainly because it is related to two main advances of agent-based modeling: heterogeneity and interactions. There exists two representations of this loyalty in the literature, either fixed loyalties assumed in order to understand its impact (Rouchier, 2004) or emerging loyalties, as the result of endogenous interactions. Vriend refers to “endogenous interactions” when he uses individual learning to generate an evolution of interactions among agents (Vriend, 2005). The idea is that agents learn to select which actions to perform as well as which agent to interact with; it is clear that this can lead to the apparition of loyalty, and that it can take different regular patterns.

One main field where this loyalty issue has been important is the study of perishable good markets (fruits and vegetables and fish). Indeed, in real situation, the participants of these markets are very dependent on the regularity, which implies predictability, of their interactions. The main reason is that buyers need to purchase goods almost every day: they have very little ability to stock and they must provide their customers with all possible goods (a retailer without carrots can turn to be unattractive just because of this lack). In case of shortage, they need to have good relations with a seller to make sure the good will be available to them. Conversely, Rouchier (2004) shows in a model that the presence of loyal agents in a perishable good market is necessary for the sellers to predict the right number of goods to provide every day. In this artificial market, two types of buyers interact with sellers: those that look for the cheapest prices (“opportunistic”) and those that are faithful and try to get the good rather than to get it cheap (“loyal”). To be able to be opportunistic, agents first gather information about prices, and then decide on the seller they want to meet to make the best transaction. The more opportunistic agents are present in the market, the more garbage is produced and shortage take place. Although there is some randomization of needs for the buyers, the presence of loyal agents makes the

sellers estimate their stocks in the best way. This result stands for different learning algorithms (step by step learning, simple reinforcement learning and classifier systems). Considering that the field study on the fruits and vegetables market of Marseilles, France, showed that most of the agents are loyal (according to the definition of the model: first loyal and then try to find all the goods in a minimum of visits to sellers), this result can give a functional explanation to their action choices.

In a slightly different context, Rouchier has represented the shape of emerging patterns of relations that could be created by two types of rationality with agents (Rouchier, 2001). The situation is a market-like situation, where the offer is depending of the preceding step situation, since it was a renewable resource. Agents were herdsmen and farmers, and the later were selling an access to their land. If no one or if too many herdsmen were coming on a land, it would get depleted, and hence the offer would be reduced. Two types of micro behavior were defined: either herdsmen would choose the cheapest farmers, or they would choose the ones that are offered them an access the most often. In the first case, the created situation was a depletion of the resource, with congestion of demand for the cheapest farmers. The links that were created were highly stable (once an agent found the cheapest it would not change), but on the other hand agents could not readapt when there was a shock in the resource quantity (a drought) because everyone would converge to the same farms. With the second rationality, agents had a representation of a “good” farmer, which was only based on individual experience, and hence they would be heterogeneous. They would also have several “good” farmers to visit in case one was not available. This made them much more flexible in their choice, and it would suppress the depletion of the resource, and everyone would be better off. The macro situation, although the process is different, also shows that a loyal micro-behavior is a help to repartition of goods were there can be shortages. In this setting the loyal micro-behavior also enables a more equal repartition of gain among farmers as well as herdsmen.

Vriend and Kirman (2000) explored the emergence of loyalty in an artificial market which representation is based on a field study in the fish market of Marseille. The aim is to see loyalty emerge, and in parallel to see which emergent

behavior display sellers. They use classifier systems to represent learning, where their agents can have a lot of different actions, some of which are, a priori, not good for their profit. Through the exploration of good and bad possible actions, they select those that bring the highest profit in the past. Some buyers learn to be loyal, and those that learn this get higher profit than others in the long term (it is actually a co-evolution where sellers learn to offer lower prices to those that are loyal). The buyers are then differentiated: their reservation price is heterogeneous (for example to represent that they do not sell their fish to the same population, some are in rich neighborhood, some in poor ones). Sellers on the market learn to discriminate, and they offer higher prices to those that have higher reservation prices. Eventually, some of the sellers get themselves specialized since only low prices buyers can visit them. Using a very basic learning where agents are not rational but learn by doing, the results are very satisfying because they reproduce stylized facts of the fish market.

A third model represents the same market but refers much more to quantitative data of this market (Weisbuch et al., 2000). The data represents sales that took place over more than 3 years and concerns 237162 individuals. From them it is possible to observe that buyers that are faithful to a seller are mainly those that buy a lot of quantities every month. The model was built in two parts: one which is simple enough to generate analytical results and a second that displays more realistic hypothesis. In the analytical model, agents use the logit function to select their action (basic reinforcement learning), which means that their choice depends on a β value, between 0 and infinity, which decrease gives a higher propensity to test randomly all sellers and which increase induces a higher propensity to look for the best past interaction. Agents can either imitate others or only base their choice on their own learning. The results can be found using the mean field approach, coming from physics. It is show that there are radically different behaviors – either totally loyal or totally “shop around agents” depending non-linearly of β . The model gets more complex with sellers being able to sell to two different prices, high and low. What can happen in that system is that a buyer gets loyal to a seller when the price that is asked is low, but that he remains loyal even after the price has switched to high. The only important thing is that, as seen before, the good is actually provided. One indicators that is used to synthesize

diverse information of the model is “order” which the quantity of agents that are loyal – the more regular the agents, the more ordered the society. Although the results, once interpreted, are very coherent with real data in terms of possible states of the market, it is a bit difficult to understand precisely some concepts of the paper because it refers mainly to physics indicators that are after translated into social indicators, but this translation is not so straightforward.

Reputation of sellers

A market has been developed as a benchmark for a reputation-based learning algorithm for agents in a social system (Pinyom et al, 2008). The model integrates quality and judgment of a relationship. Reputation is used in the group to enable agents to gather enough information in a context when it is scarce. The market that is used is a rather simple institution, where buyers have to select one seller at each time-step to buy one unit of good. The good is different for each seller, being defined by one parameter called quality. For a buyer, the acquisition of a good of lower quality will give less utility than the acquisition of a good of high quality. Sellers have a limited quantity of units, which they can sell at any period (the good is non-perishable) and they disappear from the system when they sold everything. The relevant information for buyers is first the quality of the good that each seller offers. However, when the number of sellers is large, this information cannot be captured efficiently if the buyer only learns when he meets a seller. This is why information circulates among buyers, who communicate once every time step. A buyer who meets a seller forms an image of this seller; a buyer who gets information about a seller has access to a reputation of this seller. When giving information to another buyer, an agent can decide to give the direct knowledge it has (the image it formed of a seller) or the reputation it has already receives (which is more neutral since it is not its own evaluation). Reputation can also circulate about the buyers, and in that case concerns the validity of the knowledge they give about seller. When a buyer is not satisfied with the information given by another buyer, it can also retaliate and cheat when this very agent asks him a question.

Pinyom et al. (2007) describe in detail the choices that agents make when asking a seller for a good, asking another buyer for information about a seller or a buyer, answering a question and the lying process.

The simulated market contains a large number of sellers (100 for 25 buyers) – simulations are defined by 1/ the type of information that is used by agents (only image or image and reputation) and 2/ the number of bad quality sellers in the system (99, 95, 90, 50). The addition of reputation in the system makes the difference between a normal learning mechanism where buyers select their favorite seller and a learning mechanism where they aggregate several quality of information to (maybe) become more performing. Actually the results show that globally, the agents indeed learn more efficiently when using reputation, in that the average quality that is bought is higher. The quantity of information that circulate is much higher and this enables buyers to increase their utility. This social control mechanism is especially important when quality is really scarce (1% of good sellers). This result is all the more interesting than this is a very rare case of simulated market where communication among sellers is represented, although this behavior is commonly observed in real life situation.

4 Consumers, producers and chains

Another way to look at the notion of good market is to consider large markets, where individual interactions are not important for the agents who do not record the characteristics of the ones they meet, but only the fact that they can or cannot perform an exchange. A large market can indeed include numerous goods, that are distributed among different other agents and not necessarily easy to access.

Another interest in large market is to study endogenous preferences for goods, and imagine their evolution depending on the type of good and some cognitive characteristics of agents. Eventually some authors are interested in the coordination process within the supply chain itself, where the issue is about the amount of information that each agent has to use so that to anticipate needs for far end consumers.

Multi-good economy

Two very abstract models of economy can be found where agents have to produce a good and consume other goods which they can acquire only through exchanges with other agents. The first model was built so that to produce speculative behaviors in agents, which means getting a good that has no value for consumption but a value for exchange (Duffy, 2001); the second model's aim is to

witness the apparition of commonly used equivalence value for the goods, which is analyzed as relative prices (Gintis, 2006). Both models are interesting for their pure description of an abstract economy with minimalist but sufficient assumption to induce economic exchanges. In the works cited here, the use of these two models have different methodological aims: one is purely abstract whereas the other tries to refer to experimental results and mimic human players behaviors. In his paper, John Duffy (2001) designs a model that was originally proposed by Kiyotaki and Wright (1989) to induce some agents to store a good that is not their own consumption good and is more costly to store than their own production good, just because they think this good is easier to exchange for them. There are three different goods in the economy, agents need to consume one unit of good to increase their utility and produce one unit of good each time they have consumed one. Agent 1 needs good 1 and produces good 2; agent 2 consumes good 2 and produces good 3; agent 3 consumes good 3 and produces good 1. To make it easier, Kiyotaki and Wright write that agent i consumes good i and produces good $i+1$. Hence agents have to exchange when they want to consume and not all agents can be satisfied by just one exchange. Indeed, if two agents exchange their own production goods, one can be satisfied but the other would get a useless good, which is neither its own production good nor its consumption good. In this economy, only bilateral trading exists and it takes place after a random pairing of agents. This involves that some agents must keep a good for at least one time step after production before getting their consumption good.

In this economy, speculation is defined as the conservation of the good $i+2$, since the agent does exchange to get this good which it has not produced, only because of the chances to use it as an exchange good at the next time-step. Then, the economy is made non-symmetric by having different costs in the conservation of goods, here $c_1 > c_2 > c_3$. Kiyotaki and Wright paper is all about calculating, given the conservation costs, the discount factor (the probability that the economy stops at the end of a time-step) and the utility of consumption, in which case agents decide to get the most expensive good to store or keep their production good. In an economy with perfect information, the expected profit for each type of agent depends on the proportion of agents 1 holding good 2 (which is 1-number of agents 1 holding good 3), proportion of agents 2 holding good 3 and of agents 3 holding good 3.

This centralized knowledge is not what interest John Duffy who tries to see how agents could learn which choice to make when they are able to acquire a good they do not consume. Especially agents of type 1 are those that should hesitate, since good 3 is the most expensive. A lot of models have been built on this topic (Basci 1999) already, but what Duffy wants to produce is a setting that is close to experiments he has been leading (Duffy and Ochs, 1999) so that to be able to judge if his agents are behaving in a way which is coherent with human actors. So, from a theoretical setting he builds experiments and simulations, and compares all the results this techniques produce. In this paper he therefore proposes an algorithm that is close to his intuition of what individuals should do (he has also asked questions to people involved in his experiments), and he then tries to mimic the results of his experiments, at a global level and a local level. He also proposes some original settings where he mixes human agents with artificial agents to test at the same time his algorithm and how much he can make the human change behaviour depending on the stimuli they get. He is rather satisfied with his results, where his model of learning enables to reproduce human behaviour correctly. It is possible to discuss his algorithm and his way of describing it, since the reproduction of his model is not so straightforward (Rouchier, 2003), but all in all, this description of a very basic economy with few goods and where agents learn in a way which is close to intuition is a very interesting example of market for agent-based modellers.

The paper by Gintis (2006) presents similarities, although the aim and the central question are different. The economy that is presented can be seen in a very general way and is implemented in only one setting which is described here. In the economy there are three goods, and 300 agents. Each agent can produce one good and needs to consume both goods it cannot produce; hence it is forced to exchange with other agents. At the beginning of each period, an agent only holds the good it produces (in a quantity that it can choose and which is costless) and can meet two agents, each producing one of the good he needs to acquire. Each agent has a representation of “prices”, which is here defined as the equivalence quantity between two goods. For this price, there is no common knowledge and each one has its own representation. When an agent meets another agent who can provide him with the needed good, he offers to trade, by sending as a message its

representation of relative “prices” – the exchange takes place at this exchange rate if it is acceptable to both agents and the quantities is the highest quantity that both can exchange. Agents cannot choose who they meet, they just pick randomly in the other producers’ groups. After exchanging, they can consume, which gives them utility, and defines a performance for each individual. Learning in this system is an event that takes place every 20 period, where the 5% of least performing agents (which gets the lowest utility) copy the price representation of the most performing agents.

What is observed in the system is the efficiency of the represented market, meaning the sum of all profits, compared to a setting where prices would be public. When prices are public, all exchanges can take place since all agents agree right away on the equivalence that is proposed and there is no refusal in exchange. In the long-term, the system converges to the highest possible efficiency, so although the agents have private prices, these prices get to be close enough to have regular exchanges. This result in itself is not very surprising in terms of simulation (considering the process at stake), but is interesting in economics since it gives, at least, a process to attain to a common knowledge which is often presupposed to exist.

Adoption by consumers

The study of the behavior of large number of consumer facing the introduction of new product on a market is a topic that is very interesting to approach with agent-based simulation, since it allows, once more, to look for the influence of heterogeneity of individuals and of networks in the evolution of global results. Wander Jager is a prominent character of this area of research, in a position in-between psychology and marketing. In a paper with Marco Janssen (2003), the basic model is presented. The idea behind the study of the acquisition of a new product in a group is that agents have a preference that is based on two main parameters: the individual preference for the consumption of the product and the interest that the agent has to consume the same product as his acquaintances. Hence, a utility function that defines the agent depends on this two parameters, and this will influence his decision to buy a new product. Agents are heterogeneous in such a system, and the representation of “early adopters” (in the real world: people who buy a product when it is just released because they know

well the interest of new technologies) is modeled by a low need to conform to others' belonging. On the opposite, some agents buy a good only because a lot of their acquaintances have already acquired it.

In Janssen and Jager, 2003, the influence of network size and shape is tested, as well as the influence of the utility brought by the product consumption and the way agents choose their action. One aspect that is studied is the type of cognitive process that can be used by the agent (repetition of the same action; deliberation to find a new action; imitation – where other agents' consumption is imitated; social comparison – where other agents' are imitated based on their utility). This indicator is quite rare and shows the psychological grounding of the paper; it is interesting here to observe that the cognitive process changes with the utility gained by the consumption of the considered product. In a small world network, much more products get adopted than in a scale free network, and in the first case, agents with a lot of links are very important for the spreading of product adoption. A discussion is open here about the type of products, which certainly have different influence on the way people copy others – certainly different for milk and for computers or clothes. This question is actually developed in another paper.

In Delre and al., 2007, the question that is at stake is to determine how, depending on the type of good, how to advertise efficiently. "How to advertise" means here: is it better to advertise a lot at the beginning of a campaign, or after a moment; is it better to advertise to a large number of people or to disseminate information among only a few agents?

Two products are differentiated: a brown product, which is a high tech and quite fancy product that can be compared with other agents of the network (CD, DVD player); a white product which is related to basic need and is not really compared (fridge or laundry machine). Agents are gathered in networks of different shapes. In this model, the heterogeneity is similar in the utility formula as in the preceding paper: each one is defined by a propensity to be influenced by others and to be an adopter of a new technology. The first finding is that the timing of promotion is important to the success of the campaign and that who is being touched as well. For a first launch, the best strategy is to « throw gravel », which means that one has to make a bit of advertising to a lot of distant small and cohesive groups of consumers, that will then convince their network neighbors. Another element is

that not too many people must be touched at first, since if they see that others have not adopted the good, they might not want it and become impossible to convince afterwards. This is mainly true for the white good, where it is better to advertise largely when at least 10% of agents have already adopted the good, whereas with brown good the takeoff is much faster and a campaign helps to takeoff.

The issue of the adoption of a practice within a social context has also been studied to understand the adoption of electronic trade for consumers (Darmon and Torre, 2004). The issue at stake is that it should be logical that everyone turns to electronic trade, which reduces transaction and search costs a lot, but we observe that a very small proportion of items are yet traded on internet. This is mainly because consumers have not developed special abilities that are associated to this interaction device, and do not know how to reduce the risk of performing a bad transaction. To study the dynamics of adoption of the electronic institution and learning of agents in a risky setting, a simulation model has been built. The market is composed of agents who can produce a good and exchange it for the good they want to consume (hence all agents are at the same time producers and end consumers). Agents and goods are located on a circle, the location of an agent defines the “quality” of the good it produces. For consumption, each agent is defined by an interval of quality: when consuming a good which quality is within this interval (not including its own production good), it will get a strictly positive profit. The cost of production is heterogeneous and can be stable during the simulation or evolving.

When trading on the traditional market, an agent can identify the quality of a product offered by another agent, but it has to do numerous meetings before finding out who he can exchange with (depending on the quantity of agents and of its interval of choice). The authors also added a notion of friction, which is a probability of succeeding to trade when two agents meet. In the electronic market, an agent sees all other agents at once (no search cost), but cannot identify precisely the quantity that is offered and evaluates it with an error interval. It hence potentially accepts goods that are giving zero utility. Agents are heterogeneous in their ability to perceive the quality and learn about this. If agents learn through individual learning, then it eliminates an agent from its list of

potential sellers whenever the previous trade brought utility. If agents learn through collective learning, then a part of the whole society belongs to a community that shares information about the quality (location on the circle) of the seller they met at this time-step; the agents not belonging to the community learn individually. Then in some simulations, for both learning, agents forget a randomly chosen part of what they learnt at each time-step.

In the case of individual learning, the dynamics that is produced depends on the production cost, which can change at all time-step (and hence all agent have the same average production cost over the simulation) or which can be stable and delineate populations with high or low production costs. When the production cost changes at each time-step, the main result is that all agents switch to electronic market, but in phase. Those that have a good appreciation on quality go to electronic market very fast because it is more profitable for them. Their departure of the classical market reduces the probability to exchange for the remaining agents, who eventually go to electronic market as well. When production cost are heterogeneous, some agents cannot switch from classical to electronic because of their inadequate production cost. Hence they never learn how to identify quality and stay all simulation long in the classical market. When agents forget about their learning, then the size of the electronic market does not get as high as with perfect learning, and a number of agents do not exchange.

When agents participate in a community and exchange their information, a highest number of agents participate in electronic market and overall a lowest number of agents is excluded from exchange. Three groups are created: agents belonging to the community, who get the highest pay-off; agents with low production cost or high expertise who can go on internet market and make high profit; and the remaining agents which sometimes cannot exchange. This result is rather coherent with what could be expected, but it is interesting to have created it with this location-based representation of quality that each individual wants to attain. It is especially clear that there is little risk that traditional market should disappear if the main assumption of the model (that agents need an expertise that is long to acquire before switching to internet market) is true.

Decentralized supply chain

Supply chains are an important aspect of economics, and they are often difficult to consider, mainly because their dynamics spreads in two directions: first the length of the chain, where suppliers have to adapt to demand and buyer to the speed of production so that to be able to put on the end market the right quantity of goods; then another dimension is the fact that suppliers as well as buyers in the chain are substitutable and that each actor is itself in a market position and can choose between several offers or demands. Actually, in agent-based literature, only the first issue is treated. The structure of these models is a series of agents, each representing a firm, that are link to two agents at most : a supplier and a client (except for end supplier and end consumer, of course linked to only one firm). Each agent has to decide on its production level at each time-step, knowing that it needs to use goods from the preceding firm in the production process. It must then anticipate on demand to order enough, before being able to transform the needed quantity. Of course, each firm takes some time (number of steps) to transform the product and be able to sell and there is a cost in keeping its own production when it is not completely sold.

One very important issue of these chains is at the centre of most research: how to avoid bullwhip effect. This effect is a mechanical dynamics that comes from the slow spreading of information and delay in answer because of the length of the production process in each firm. When there is variability in demand coming from end consumers, this variability increases a great deal when it goes in the chain, up to the first producer who exhibits the highest variability. This can be very annoying for organizations to be trapped in such a negative dynamics. Several authors propose algorithms for artificial agents that have to deal with the issue of anticipating demand at each stage of the chain. For example Lin and Lin (2006) describes a system where artificial agents can interact with real agents (and hence be integrated in a real-life company to help deciders) so that to choose the right level of production and order to reduce costs. Several learning algorithms are tested and their efficiency attested, even in situation where the environment is dynamically evolving. The same issue is treated for example by Kawagoe and Wada (2005) who propose another algorithm. They also propose a method to statistically evaluate bullwhip effect. Their method is different from the usual

frequency based statistical measurement (like stochastic dominance) but is based on descriptive statistics.

5 Financial market, auctions

Financial markets have been one of the first example that were developed to prove the interest of Agent-Based modeling. Arthur et al. (1997) indeed reproduced some important stylized facts of asset exchanges on a market and is always cited as the first important use of this modeling techniques for studying markets.

Contrary to models that were presented before, there is no direct interaction among agents in these models, only observation of price patterns. One rare example presented here is an attempt to link a financial market to a consumer market such as one that was seen in previous sections. Another type of market does not integrate any interaction in the economy is the representation of auctions.

Financial market

The literature on financial market is very important in agent based simulation, and dates back to the 1990's (Arthur, ; Arifovic, 1996) and also in the related branch which is called econophysics (the use of physic techniques to deal with economics issues in systems that are composed of a huge number of simple interacting actors (Levy et al., 2000)). A large review of this topic (Lux, *tbp*) will be published in a Handbook of Finance. It describes at the same time the main stylized facts that can be found in financial markets (and hence are meant to be reproduced by simulation) and some models that are candidate for explaining these facts.

Another large review by several authors describes several agent-based simulations models dealing with financial markets but that are not so right in reproducing very general statistical regularities of these markets. (Samanidou et al., 2007). As usual in this section I will describe only some models and ways to represent agents learning in the context of financial market. The basic structure of the market, which defines the type of choice the agent has to make, can vary, as well as the aim and methodology of the researcher building these models and this is why examples seem to be good to give few complete descriptions instead of very generic results.

One reason of using agent-based models is to be able to represent heterogeneous population of agents, and what is very often found in this research is the representation of two type of agents that have different reactions facing information: chartists and fundamentalists. Fundamentalists base their investment decisions upon market fundamentals such as dividends, earnings, interest rates or growth indicators. In contrast, technical traders pay no attention to economic fundamentals but look for regular patterns in past prices and base their investment decision upon simple trend following trading rules. Computer simulations such as those of the Santa Fe artificial stock market (LeBaron et al. (1999), but see also e.g. Kirman (1991), Lux and Marchesi (1999, 2000), Marsili et al. 2000, 2001) have shown that rational, fundamental traders do not necessarily drive out technical analysts, who may earn higher profits in certain periods. An evolutionary competition between these different trader types, where traders tend to follow strategies that have performed well in the recent past, may lead to irregular switching between the different strategies and result in complicated, irregular asset price fluctuations. Brock and Hommes (1997, 1998) have show in simple, tractable evolutionary systems that rational agents and/or fundamental traders do not necessarily drive out all other trader types, but that the market may be characterized by perpetual evolutionary switching between competing trading strategies. Non-rational traders may survive evolutionary competition in the market (see for example, Hommes (2001) for a survey).

In Hommes and Lux (2007), the model of market that is chosen is the one of cobweb experiment, which is a prediction model on a market, not an actual model of selling and buying for agents. The model offers however a rational expectation value, which serves as a benchmark for the experimenters. The methodology is to try to fit agents' behaviour in an artificial world to real behaviours of individuals in experiments. The game is such that participants to the experiments have no clear idea of the structure of the market but must however predict the price of the next period. Neither do they know how many other agents are present, nor do they have in mind the equation that calculates future price based on the realised price and the expectations of all participants. The simulations are made based on rather simple models of agents, one being genetic algorithm, simple learning that copies past prices, and also reinforcement learning. What interests the authors most is the

GA learning, which is the only one to fit stylized facts in different treatments. What the GA learns about is a 40 bit string of 0 and 1 that represents two values, α between 0 and 1 with 20 bits and β with the remaining 20 bits, that predict price at $t+1$ depending on price at t with $p(t+1) = \alpha + \beta (p(t) - \alpha)$.

There are three treatments both for experiments and for simulations, with one parameter defining the stability of the price, being high, medium or low. The genetic algorithm being varied for different mutation rates, it is proven to be largely better than other learning procedures that have been implemented.

“Better” means here that it fits the stylized facts that have been produced by humans in experiments: (1) the mean price is close to rational expectation, and the more stable the market, the closest to this rational expectation value; (2) There is no significant linear autocorrelation in realized market prices. The reason of the good fit of GA is given by authors, which is really interesting because it is not obvious to imagine how GA, which are a random learning process with selection, should be similar to human learning. They assume that two facts create a learning that is similar to human one: the fact that successes are selected positively and that there is heterogeneity of the strategies among the set that agents can use. Once the assessment of the model is thus done, it is used to question the stability of the results of the learning process. One question that arises is to wonder whether humans would adapt the same way when interacting in a very large group as they do in a small group of 6. This opens many questions on the scalability of results concerning market dynamics.

In the paper that is described now, the interaction of agents is direct and not necessarily via price system, as usual in financial markets. Hoffman et al. (2007) indeed consider that many agent-based simulations still take little interest in representing actual behaviours of decision makers in financial markets. They argue that Takahashi and Terano (2003) is the first paper to integrate theories that come from behavioral finance and represent multiple type of agents, such as overconfident traders. In their own paper, Hoffman et al. (2007) present SimStockExchange™, their platform, with agents performing trades and making decision according to news they perceive and prices they anticipate. They argue that their model is based on several theories that are empirically sound and that they check their model thanks to simulations which results are compared to the

Dutch market along a few years. As usual, the platform allows many variations (increase the number of different shares, of agents, change the characteristics of agents) but is tested only with some values of parameters.

Agents receive news that they forget after one time-step, and then can perform two types of action: either sell their stock (if they expect to lose at the next time-step) or buy more shares (in the opposite case). To make sure that they are not making mistakes, agents can use risk-reducing strategies, which can be clarifying strategies (such as collecting more data) or simplifying strategies (i.e. imitating other agents), as well as purely individual (the first one) or social (the later). In the presented simulation, strategies are always social, and hence Agents' confidence, C , determine their use of risk reducing strategies, and these confidence were deduced from empirical studies. Each agent is also defined by a tendency R to perform simplifying strategy or clarifying one. R and C are evaluated on the basis of surveys made with investors. Agents are imbedded in networks that can have two shapes – a Torus network, and a Scale-free one – and from which they acquire the information from others or choose to imitate them. The market itself is designed as an order book, where sells and buys are written, with quantity and price, and are erased as soon as an agent answers positively to this proposal. The market price is the average of all proposed bids and asks of the order book – hence it is not a realized price (average transactions' price) but an aggregation of desired prices for agents.

In the results, some statistical properties of the stock exchange have been reproduced. For example, with weekly data of Dutch stock exchange, linear autocorrelation can be observed, and it is much better reproduced by torus network than scale-free one. When it comes to volatility clustering, it is clear that torus network display differences from Scale-free network and real data, by showing a lot of volatility clustering. This can be due to the speed of circulation of information that disseminates this information so fast that it reduces the shocks that it can cause. The main aspect of the SSE that has to be better modeled is the news arrival, which is a normal distribution around the present price. Maybe this can have a large impact since the use of different networks integrates the importance of information spreading.

Relation between two markets

Sallans et al (2003) report a model where two types of markets are integrated: a financial market and a good market in the same system. Consumers, financial traders and production firms are interacting and the aim is to understand how these two markets influence each other. The good is perishable, and hence needs to be purchased regularly. Consumers make purchase decisions; firms get income on the sales and update products and pricing policies based on performances; traders have shares which they can hold, sell or buy. The product features are described by two binary strings of 10 bits, consumers look for special features in the market, and firms choose the feature of the good they put on the market. In choosing actions, the firm uses an action-value function that integrates expectations about future rewards (firms are not myopic agents), by taking into account the evolution of the price of its share in the financial market and the profit made by selling products on the market. Consumers have preferences on features of a product and its price, and compare any available product to these preferences: they can choose not to buy if the product is too different from their preferences. In the financial market, agents build expectations and built representation of future values by projecting actual and past values in the future. They are parted in two groups: fundamentalists (use past dividend for projection) and chartists (use the history of stock prices); they are also heterogeneous in time horizon. The market clearing mechanism is a sealed bid-auction and the price is chosen to maximize the number of exchanges (and randomly among different prices if they produce the same trade volume).

Since agents from the financial market and those from the firm have different views on the future of the firm, and evaluate future gains in a different way, leading to impact on the firm performance that are not necessarily positive. The simulations' aim is to prove that the model can be used, in certain parameter settings, to reproduce stylized facts of markets.

Although the central issue is very interesting, the paper itself is not as helpful as could be to understanding the dynamics of two markets. In particular, the stylized facts are not very explicit in the paper (appear only once at the end, when obtained results are given). They are classical in financial market analysis, but not clearly shown here: low autocorrelations in stock returns, high kurtosis in marginal return, volatility clustering. Hypothesis on behavior are never explained, hence

there is no understanding of why the stylized facts can be achieved, apart from doing some parameter space random exploration. Although the main issue of the paper is really fascinating, the results are a bit frustrating because the paper is very redundant and self-justified, and hence the reciprocal influence of these market, so important in our real world, stay undisclosed.

Double auctions

In economics double-auction is a very fascinating topic, since it is an extremely stable institution in which predictions can be translated from theory to real life, which is not really the case for most economic systems. When putting real people in a double-auction setting, one can observe that the convergence to equilibrium price occurs. This does not mean that this institution is efficient, since a lot of exchanges take place out of equilibrium price, but at least there is a tendency for the group to converge to a price where the highest number of exchange can be performed, and hence the highest global profit can be extracted. Many authors hence wanted to reproduce double-auction market in artificial society so that to understand what is the source of this high efficiency.

The CDA (continuous double auction) is a two-sided progressive auction. At any moment, buyers can submit bids, (offers to buy). Similarly, sellers can submit asks (offers to sell). Both, buyers and sellers may propose an offer or accept the offer made by others, representing the counteroffers in the market. If a 'bid' or 'ask' is accepted, a transaction occurs at the offer price. An improvement rule is imposed on new offers entering the market, requiring submitted bids (/ asks) at a price to exceed (/ be less than) the standing bid (/ ask). Each time an offer is satisfying for one of the participants, she announces the acceptance of the trade at the given price, and the transaction is completed. Once a transaction is completed the market is cleared (meaning there is no standing bid or ask any more) and the agents who have traded leave the market. At that moment, like at the very opening of the market, the first offer can take any value, and this proposed price imposes a constraint on any following offer. When the market closes, after a time decided beforehand, agents, who have not yet traded, are not allowed to continue. In this market institution all market events are observed by all (bid, ask, acceptance and remaining time before market closing) and hence is said to be common knowledge.

Using this double-action setting, a seminal paper by Gode and Sunder (1993, 2004) shows the strength of institutional constraints on the actions of agents. Indeed in their setting, agents are perfectly stupid from an economics point of view, since they have no understanding of their own interest and that they only follow a rule without having strategic behavior. They are called “zero-intelligence” agents: they are not allowed to sell (/buy) lower (/higher) than their reservation price, and they have to bid within the limits that have been put by others. With these rules, convergence of prices is obtained, very fast. The approach in this paper is quite original in the behavioral economics literature in the sense that it is close to an “artificial life approach”. Authors do not pretend to study human rationality, but instead focus on the abstract reproduction of phenomena. It is interesting to note that it is not so easy to design a double-auction market, and especially in its continuity. Indeed, in a real situation, if two individuals have close reservation price, they will often be able to buy or sell at the same moment. Who will be first is not obvious, since people have different aspiration for profit. Hence, the way to produce a double-auction for Gode and Sunder is to randomly choose an agent between all buyers who can buy or make a bid, and then to pick randomly a seller among those who can sell or make an offer. After trying several methods, they choose this one, explaining that this is a good approximation to continuous actions.

Their work is widely criticized because they are not interested in rationality but in a specific institutional setting and cannot be generalized to diverse settings (Brenner, 2002). However, their result is important and led a lot of researchers to question it. For example Brewer et al. (2002) show that humans have abilities to have markets converge when context changes a lot, which Gode and Sunder’s agents cannot do. In their setting they organize a double-auction market in which agents participate in the public market, but also receive offers from the experimenter in a private way. Only one offer is made at a time, and it is the same for all agents that are proposed the offer, since the equilibrium has to stay the same. The global equilibrium (which value is described in the paper) is thus constant, but individuals can have incentives not to participate to the public market if the offer is interesting. This does change the performance of zero-intelligence a lot, since the prices do not converge anymore in simulations led with this new institution. On the opposite, humans performing experiments attain convergence, which could

mean that only very specific institutions constrain participants enough so that they have no choice but to converge, even while not understanding more than the rules (zero-intelligence).

Another paper has been inspired by Gode and Sunder, but also by the theoretical model of Easley and Ledyard (Rouchier and Robin, 2006), and tried to establish the main elements that an agent should have in its rationality to be able to choose the right action in a context of double-auction. To differentiate among different possible learning procedure, a comparison to some experimental results were made. Learning procedure was a simple algorithm that made the agent revise its reservation price towards past average perceptible prices, but it depended on two elements. First, the moment when the agent would change its reservation price (i.e. buyer (/seller) accepting higher (/lower) prices), called here the “stress time”, could change by increasing after a successful transaction and decrease after a day with no transaction. Second the agent could perceive only its own transactions or any successful transactions in the market. What was demonstrated in the paper is that agent would learn quicker to converge to the equilibrium price (making the highest global pay-off) if they did not revise their stress-time and had a global perception of prices. This quick learning would at the same time correspond the best to the speed of convergence that could be found in experiments. What is a bit surprising in this result is that more “clever” agents (reacting to risk and failure from one day to another) would neither copy human behavior well nor get to efficient situation very fast.

Market design / agent design

In a chapter of the handbook for computational economics published by Tesfatsion and Judd (2007), Robert Marks (2007) reviews recent work in market design using agent-based simulation. Market design is the branch of economic research which aim is to provide insights about which interaction structure (and hence information circulation rules) is the best to obtain certain characteristics out of a market. As said repeatedly in this chapter, this choice is crucial to have certain part of population have more power than other, or have efficiency attained in a short time. Hence, many scientists have been thinking about this issue, using game theory (Roth, 1991), as well experimental economics, and more recently

computational exploration. As seen before, sophisticated agents are not the one who do the best in market situations, and are not the ones that copy humans the best.

When designing an institution it is important to see two problems. First the “aim” of the institution should be made clear. For example using Dutch auction has one advantage over English auction: it is fast ; double-auction is good because it extract the highest global profit for the participants (on both sides), but one could wish to extract the highest profit for buyers only, for example. Since not all good aspects can be achieved in one institution (see for example Myerson and Satterthwaite, 1983). Lebaron (2005) explains that the fitness of a model is as important as all other elements (what is traded, what are the motivation of agents, what is the interaction and information circulation organised, ...). This trade-off between different characteristics to achieve is already a huge choice before starting the design.

Then one has to think on how to achieve this aim. It is indeed not easy to know how individuals will react to an interaction and information constraint. The basic use of agent-based simulation can then be to either test a certain rationality of agent and compare institution to see what difference it makes in prices (or other indicators) (Audet at al. 2002; Bottazzi et al. 2003; Moulet and Kirman, 2008) or be used to test different learning algorithms in the same setting (Chan and Shelton, 2001). Both approaches are really uniquely developed using agent-based, and can indeed help understand the relation between behavior and institution.

Many works, be it for computer scientists or economists, were designed to fit the context of electricity market, which is very central since the problems can be extremely heavy for society (when there are huge unpredicted shortages) and the variations in price can be very fast. In the models that are created agents are not designed to represent human’s rationality, but to try to be as optimal as possible in the adaptation to the electricity market. Many institutions can be used, but auctions (which are theoretically supposed to be the most efficient of all market institutions) are widely used. Bidding behaviors, but also concentration of sellers and buyers and capacity to produce and sell (Nicolaisen at al., 2000) have an impact on the efficiency and this can be explored. As said before what is explored is the impact of the institution on efficiency and on market power. Two ways of

learning are used for the agents, and authors sometimes disagree on the way to choose: either social with a Genetic Algorithm, or individual with reinforcement learning, and it is already well known that this has a huge impact on global results (Vriend, 2000). We cannot here decide on the best choice to make. However, to our view, most results cannot really be extended to real life design since the representation of learning for agents can sometimes be badly adapted to the application context (necessity to have long learning in case of GA or even reinforcement learning).

One original approach that is cited by March (2007) is the “evolutionary mechanism design” (Phelps et al. 2002), where strategy of three actors, sellers, buyers and auctioneers are all submitted to evolution and selection (the fitness of the auctioneer’s strategy being linked to the total profit of the participants). This approach is logically different since the institution itself (through the auctioneer) is what evolves to get to a better result, with the characteristic of the participants being fixed (relative number of each and relative production and demand).

It is interesting to note that another branch deals with the representation of individual agents on large markets, and is also quite close to an idea of design of markets, but from the opposite perspective: by introducing agents in real markets. Computer scientists interesting in the analysis of cognition indeed have the goal of making artificial agents as efficient as possible in a context of bidding in auctions, at the same time from the point of view of the seller and the buyer (Greenwald and Kephart, 2002). They usually are not interested in the understanding of real humans’ behaviour and decisions, but rather in explaining the properties that can emerge in markets in which many artificial learning agents interact (with each other or humans), differentiating their strategies, getting heterogeneous payoffs and creating interesting price dynamics. The focus is very much put on information treatment. This very applied approach is interesting in so that many algorithms that are used for building programs can also be used for economic analysis in the framework of models of the type that have been explored here. However the aim is slightly different, since the indicator is in this latter case the individual success of a strategy, whereas the previous works on markets is based on global properties of the system.

6 Concluding remarks

This chapter is not an extremely general view of market simulation in recent research, and instead of giving many examples, we focused on a few to show the diversity of questions, models, rationality and eventually results that can be found in the literature, coming from sometimes different backgrounds (classical economy, experimental economy, computer science). The representation of a market is always linked to the aim of the research, and there is never one way forward. The quantity and substitutability of goods, the possibility to interact with one or several sellers, with other buyers, the memory of the agents themselves, all depends on the type of issue, and this is why we have build the chapter in this manner: to give some ideas of the issues that have been addressed up until now with agent-based. What is noticeable is the real difference between this approach and the classical approach in economics, where the dynamics is not a central question. The achievements with the method are now numerous enough to prove that agent-based simulation can really participate in a better understanding of market institution, of behavior of individuals on market, and enhance the institutional choices of politics. What can be noted in conclusion is that several issues are still at stake when it comes to the representation of markets.

First, like in most models, the temporal issue is huge. Most models use discrete time to organize the simulation. This is not so easy to understand the meaning of discrete time when it comes to an auction, where different agents can act precisely at the same time and have a different impact on prices than when they act sequentially. Some people are specifically working on this issue and build platforms that support a simulated continuous time⁴ (Daniel, 2006).

Another technical issue is the one of learning sequences of actions. In a situation where agents evaluate their actions with profit, if they have to perform several actions in a row (i.e. choosing a seller and then accepting a price or not), it is impossible to decide which of these actions is the reason for a success or a failure. Facing this issue, economic papers describe agents that associate the profit to all

⁴ Natlab, which can be found at: <http://www.complexity-research.org/natlab>

actions, as if they were separated. This is clearly not very satisfying in terms of logic, but no alternative modelling has been proposed yet.

Eventually there is a conceptual gap in all the cited models. For the moment, another element has never been taken into account in the representation of agents' reasoning on markets, which would fit in models where agents try to maximize their profit by choosing the best strategies. In this case, they can scan past actions and the following profits, or their past possible profit with all actions they could have undertaken and then select the best action in all contexts. This last vision is a bit more general than the first one, but in no case do the agents imagine that a change in their action will modify other agents' behavior as well. This is strange enough, since a lot of people interested in game theory have been working on agents in market, but they have not produced models of anticipation about others' choices. In markets where bargaining is central, it could however be a central feature in the understanding of real human behavior.

7 Bibliography.

Arthur, B. (1991) On designing economic agents that behave like human agents: A behavioral approach to bounded rationality *American Economic Review*, n. 81, pp. 353-359.

Arthur, B. , 1994, Inductive reasoning and bounded rationality *American Economic Review*, n. 84, p. 406.

Arthur B., Holland J., LeBaron B., Palmer R. and Tayler P., 1997, *Asset Pricing under Endogenous Expectations in an Artificial Stock Market*, In: Arthur B., Durlauf S. and Lane D., *The Economy as an Evolving Complex System*, Studies in the Sciences of Complexity XXVII, Addison-Wesley, Santa Fe, pp 15-44.

Arifovic, J., 1996, The behavior of the exchange rate in the genetic algorithm and experimental economies, *Journal of Political Economy*, vol. 104 n. 3, pp. 510-541.

Axtell R., 2005, The complexity of exchange, *The Economic Journal*, Vol. 115 (June) pp 193-210.

Basci E., 1999, Learning by imitation, *Journal of Economic Dynamics and Control*, 23, pp 1569-1585.

Becker G., 1962, Irrational behavior and economic theory, *Journal of Political Economy*, Vol 70, pp 1-13.

- Boylan, R., El-Gamal, M. (1993). "Fictitious Play: a Statistical Study of Multiple Economic Experiments," *Games and Economic Behavior*, **5**, 205-222.
- Camerer, C., Ho, T.H. (1999). "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica*, **67**, 827-874.
- Brenner, T., 2002, A Behavioural Learning Approach to the Dynamics of Prices, *Computational Economics*, Vol. 19, 67-94.
- Brenner, T., 2006, Agent learning representation: advice on modelling economic learning, In : Tesfatsion and Judd (eds) *Handbook of computational economics vol. 2: Agent-based computational economics*, Elsevier / North Holland, chapter 18.
- Brewer P.J., Huang M., Nelson B. and Plott C., 2002, On the Behavioral Foundations of the Law of Supply and Demand: Human Convergence and Robot Randomness, *Experimental economics*, **5**, pp 179-208.
- Cyert, R. M. and March, J. G., 1963, *A Behavioral Theory of the Firm*, Prentice-Hall.
- Daniel G., 2006, Asynchronous simulations of a limit order book, PhD thesis, University of Manchester, <http://gillesdaniel.com/PhD.html>.
- Darmon E. and Torre D., 2004, Adoption and Use of Electronic Markets: Individual and Collective Learning" (with D.Torre): *Journal of Artificial Societies and Social Simulations (JASSS)*, **7**, 2.
- Delre S.A., Jager W., Bijmolt and Janssen, 2007, Targeting and timing promotional activities: An agent-based model for the takeoff of new products, *Journal of Business Research*, **60**, pp 826-835.
- Duffy J., 2001, Learning to Speculate: Experiments with Artificial and Real Agents, *JEDC*, **25**, pp 295-319.
- Duffy J and Ochs J, 1999, Emergence of money as a medium of exchange: An experimental study, *American Economic Review*, **89**, pp 847-877.
- Friedman, Milton (1953), «The Methodology of Positive Economics», in *Essays in Positive Economics*, Chicago, The University of Chicago Press, p. 3-43. Republ. in *Readings in Microeconomics*, William Breit and H. M. Hochman (eds.), 2nd ed., New York, Holt, Rinehart and Winston, 1971, p. 23-47.
- Geertz C., Geertz H. and Rosen L., 1979, *Meaning and order in Moroccan society. Three essays in cultural analysis*, Cambridge University Press.

Gintis H. (2006) "The Emergence of a Price System from Decentralized Bilateral Exchange," *Contributions to Theoretical Economics*: Vol. 6 : Iss. 1, Article 13.
Available at: <http://www.bepress.com/bejte/contributions/vol6/iss1/art13>

Greenwald A and Kephart J., 2002, Shopbot Economics, Autonomous Agents and Multi-Agent Systems: Special Issue on Game-theoretic and Decision-theoretic Agents. Kluwer Academic Publishers, 5(3):255-287.

Hoffmann A.O.I., Jager W. and Von Eiji J.H., 2007, Social simulation of stock markets: taking it to the next level, JASSS, vol 10, n° 2, 7.

Hommes C., 2007, Bounded Rationality and learning in complex markets, CeNDEF Working paper 07-01 University of Amsterdam,
<http://www1.fee.uva.nl/cendef/>

Hommes C. and Lux T., 2007, Bounded Rationality in Complex Markets: Learning Algorithms to explain Strategic Behavior in Cobweb Experiments,

Jager W., 2007, The four P's in social simulation, a perspective on how marketing could benefit from the use of social simulation, *Journal of Business Research* 60 (2007) 868–875

Kahneman D, Tversky A., 1979, Prospect theory: an analysis of decision under risk, *Econometrica*, 2, pp 263-292.

Kawagoe T., Shihomi Wada, 2005, A Counterexample for the Bullwhip Effect in a Supply Chain, In: Philippe Mathieu, Bruno Beaufils, Olivier Brandouy (eds), *Artificial Economics, Agent-Based Methods in Finance, Game Theory and Their Applications*, Lecture Notes in Economics and Mathematical Systems, vol 564, Springer, pp 103-111.

Kirman A, 2001, Some Problems and Perspectives in Economic Theory, In: G. Debreu, W. Neufeld and W. Trockel (eds) *Economic Essays, A Festschrift for Werner Hildenbrand*, pp 231-252, Springer, Berlin.

Kirman A.P., Vriend N.J., 2001, Evolving market structure: an ACE model of price dispersion and loyalty, 25, 459-502.

Kopel M., Dawid H., 1998, On economic applications of the genetic algorithm: a model of the cobweb type, Journal of Evolutionary Economics, Vol 8 (3) pp 297-315.

Leloup B. (2002) Dynamic Pricing with Local Interactions: Logistic Priors and Agent Technology Proceedings of the 2002 International Conference on Artificial Intelligence, CSREA Press, June 24-27, Las Vegas.

Levy, M., Levy H. and Solomon S., 2000, Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena, Academic Press.

Lin F-R and Lin S-M, 2006, Enhancing the Supply Chain Performance by Integrating Simulated and Physical Agents into Organizational Information Systems, Journal of Artificial Societies and Social Simulation vol. 9, no. 4<<http://jasss.soc.surrey.ac.uk/9/4/1.html>>

Lux, T., 1998, The Socio-economic Dynamics of Speculative Markets: Interacting Agents, Chaos and the Fat Tails of Return Distribution, Journal of Economic Behaviour and Organization, 33 (2) pp 143-165.

Lux, T. and Marchesi, M., 1999, Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market, Nature, vol. 397, pp 498-500.

Marks R., 2007, Market desing using agent-based models, In L. Tesfatsion & K.J. Judd (Eds.), Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics, chapter 27,. (pp. 1339-1380.). Amsterdam: Elsevier.

Moss S., 2002, Policy analysis from first principles, PNAS, 99 (3), pp 7267–7274.

Moulet S. and Rouchier J., 2007, The influence of sellers' beliefs and time constraint on a sequential bargaining in an artificial perishable goods market.
Palmer, R., Arthur, W. B., Holland J.H., LeBaron, B. and Taylor, P., 1993,

Artificial Economic Life: A Simple Model for a Stock Market, *Physica D*, vol 70, pp 264-274.

Oliver J., 1997, Artificial Agents Learn Policies for Multi-Issue Negotiation , *International Journal of Electronic commerce*, 1 (4) pp 49-88.

Phan D., 2003, From Agent-Based Computational Economics towards Cognitive Economics, In : P. Bourguine and J.-P. Nadal Editors, *Cognitive Economics*, Springer-Verlag.

Pinyol I., Paolucci M., Sabater-Mir J., Conte R., 2007, Beyond Accuracy. Reputation for Partner selection with Lies and Retaliation, MABS 07, Eighth International Workshop on Multi-Agent-Based Simulation. Honolulu, Hawaii.

Rothschild M. (1974) A Two-Armed Bandit Theory of Market Pricing, *Journal of Economic Theory*, 9 p.185-202.

Rouchier J., Bousquet F., Requier-Desjardins M. and Antona M., 2001, A Multi-Agent Model for Describing Transhumance in North Cameroon: Comparison of Different Rationality to Develop a Routine, *Journal of Economic Dynamics and Control*, 25, 527—559.

Rouchier J., Hales D., 2003, How To Be Loyal, Rich And Have Fun Too: The Fun Is Yet To Come, ESSA Conference. GREQAM Working Paper 03B13

Rouchier J., Robin S., 2006, Information perception and price dynamics in a continuous double auction, *Simulation and gaming*, 37, pp 195-208.

Sallans B., Pfister A., Karatzoglou A., Dorffner G., 2003, Simulation and validation of an integrated markets model, *JASSS*, Vol 6, n°4, 2.

Simon H., 1955, A behavioural model of rational choice, *Quarterly Journal of Economics*, Vol 69, pp 99-118.

Strader T., Lin F.-R. and Shaw M.J. (1998) Simulation of Order Fulfillment in Divergent Assembly Supply Chains *Journal of Artificial Societies and Social Simulation* vol. 1, no. 2, <<http://www.soc.surrey.ac.uk/JASSS/1/2/5.html>>

Tesfatsion L. and Judd K.L., 2007, *Handbook of computational economics vol. 2: Agent-based computational economics*, Elsevier / North Holland.

Vriend N.J., 2000, An illustration of the essential difference between individual and social learning and its consequences for computational analyses, *Journal of Economic Dynamics and Control*, 24, pp 1-19.

Vriend N.J., 2004, ACE models of market organization, *Revue d'économie industrielle*, 107, 3rd trimester.

Vriend, N.J., 2006, ACE Models of Endogenous Interactions. In L. Tesfatsion & K.J. Judd (Eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. (pp. 1047-1079). Amsterdam: Elsevier.

Weisbuch G., Kirman A., Herreiner D., 2000, *The economic Journal*, Vol. 110, N° 463, pp 411-436.

White H., 1995, "Varieties of Markets." in B. Wellman and S.D. Berkowitz (eds.), *Social Structures: A Network Approach*, New York: Cambridge University Press.