

Agent-Based Computational Economics and Finance: early research and design issues

Yann Semet, Sylvain Gelly, Marc Schoenauer, Michèle Sebag
DREAM Project - CMAPX, LRI, INRIA

May, 2003 - DRAFT

Abstract

Modelling economic or social systems in general and financial markets in particular with distributed networks of evolutionary agents is a very active and growing field known as Agent-based Computational Economics (ACE). It aims at explaining global behaviours and structures of social systems in terms of multiple iterative interactions of simple but adaptive localized agents. A concise survey of literature is conducted here that outlines key seminal works. It essentially builds over a broad survey by Tesfatsion [28] and two by LeBaron [17, 18]. It proceeds as follows. After positioning the field in terms of economic theory and early intuitions, it outlines, after Tesfatsion [28], the eight main directions of ACE research. Focusing then on the particular case of financial markets, it follows LeBaron [17] for a review of early and influential experiments. Still with LeBaron [18], one finally enumerates some of the key practical design questions any artificial market designer is going to be confronted to.

1 Introduction

“I will buy with you, sell with you, talk with you, walk with you, and so following; but I will not eat with you, drink with you, nor pray with you.”

William Shakespeare, The Merchant of Venice (I, iii, 35-39)

1.1 Smith and Hayek’s insights: invisible hand, decentralized thinking and individualism

Thinking the economy in terms of competitive interaction of localized individualities is not new, Adam Smith already [13, 27], a man considered the father of modern economics, explains how, using the famous metaphor of an “invisible hand”, economic society is structured by the competitive interaction and exchanges of localized and individualist individuals without any explicit agreement on the promotion of collective well-being in general:

Every individual necessarily labours to render the annual revenue of the society as great as he can. he generally neither intends to promote the public interest, nor know how much he is promoting it... He intends only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention.

This insightful concept has persisted over centuries and is central to much research conducted in modern economics and even in Artificial Intelligence where Swarm Intelligence techniques are based on barely more than this idea that coherent collective behaviour can arise from the interaction of a colony of short-sighted individuals.

Much later on, in 1945 another brilliant and insightful work [12] will come out by Friedrich August von Hayek, 1974 nobel prize in economics, a key figures of the 20th century’s economics, well known for his positions on individual liberty. In his 1945 paper entitled “The use of knowledge in society”, Hayek clearly advocates for an economy thought in decentralized terms and strongly denigrates the use global control and analysis based on statistical aggregates that, according to him, do not take into account a key feature of economic systems in their quest for efficiency and happiness: the ability to conduct *rapid adaptations to changes* thanks to distributed information and knowledge. In Hayek’s own terms:

The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form, but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess.

Hayek’s thought clearly supports most arguments used by ACE researchers and a recent paper [32] investigates in depth the question of Hayek being a precursor of modern ACE research in terms of methodology.

1.2 Financial Market Efficiency and Rational Competitive Equilibria

Following Hayek in seeing market prices as a fantastic information conveyors, Roy Radner introduces in 1979 [23] the notion of Rational Expectations Equilibrium. In a still influential paper entitled “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices”, Radner shows how asymmetry in information born by different traders is flattened out by the information one can take out of the evolution of prices, leading the economic society to a state of competitive equilibrium based on expectations each individual has on the link between information and prices. The general theory that comes out of this work is known as the *efficient market theory* and stipulates that prices are *fully revealing*, which means that they contain all of the information available to traders and that consequently prices follow an unpredictable random walk [21]. This theory, prevailing among many academic economists has huge consequences as it forbids any kind of arbitrage profit: it is impossible to beat the market by acquiring private information as this information would be immediately reflected by prices and therefore available to all traders. Radner advocates that such equilibria exist in general and form the basis of what happens in real markets.

Two strong controversies exist however concerning this theory. First of all, it seems to be denied by empirical evidence. Second of all, seeing all the efforts spent by financial traders to acquire private information as vain because it can all be inferred from price movements is rather counterintuitive. In an antagonist but just as influential work, Grossman and Stiglitz [11] show that the existence of competitive equilibria is inconsistent with costly information acquisition and therefore conclude on the “Impossibility of Informationally Efficient Markets.”

Finance oriented ACE research is at the very heart of this controversy between academic practice and real world practice that is still very active today as illustrated by this recent paper by Krebs [15] where it is claimed that competitive equilibria are actually consistent with costly information acquisition. ACE experiments could greatly help determine what the boundaries are between a competent practice of asset trading based on information competition and a theoretically limited potential profit.

2 Survey 1 [28]: eight areas of research

In a broad and extensive survey of existing ACE literature, Leigh Tesfatsion gives an idea of the variety of fields and interests this research covers (the interested reader can also refer to [31, 29, 30, 1, 25] for an extensive wealth of recent material on ACE research in general). She enumerates eight non exclusive directions of research, or reasons why one might get interested in ACE:

Learning and the embodied mind . A great deal of attention has been drawn on the learning aspect of ACE systems. Three key questions call for answers:

- *Who is learning what ?* Designers have to decide what entities (traders, brokers, markets themselves, etc.) learn what (strategies, representations) and based on what information (time series, aggregate indicators, etc.). To do so, they have to know what their ultimate objectives or intellectual interests are : coming up with original and efficient trading strategies, designing optimal and fair market rules or deriving relevant market indicators ?

- *What motive for a learning algorithm ?* The learning algorithm should not be seen only as an optimization tool. It can certainly be used as such in certain context where one aims at maximizing precise indicators such as market efficiency or average wealth by tweaking well defined parameters in a clever way but the learning process itself is a topic of considerable interest when studied as a model of what actually occurs in real human driven markets. In the latter case, reaching optimality becomes less than important and learning dynamics draws most of the researchers' attention. Typical human and social characteristics then come into play such as fuzziness, irrationality, imperfect learning or clustering of information networks.
- *How to implement it ?* Evolutionary algorithms are widely used being a seducing metaphor of natural social competition. Numerous other possibilities were explored among which: reinforcement learning, Q-learning and classifier systems. See [28] for detailed pointers to examples of these. The natural trend to use off the shelf optimization methods as black box tools tends to fade out as studies such as that of Dawid [8] show that parameters settings and algorithmic choices can have a great qualitative impact on the resulting behaviour and auto-organisation of the system.

Evolution of Behavioral Norms This particular direction of research is concerned with the dynamics of creation, life and death of behavioral norms or social conventions, which are rules that naturally emerge out of social systems without any kind of global planning and that prevail in dictating individual agents' behaviours by means of punishment/reward mechanisms that apply to agents that decide (or not) to follow the norm. After Axelrod's [6, 5] influential work, the dynamics of norms in distributed agents networks are seen as an evolutionary process, which ACE an especially relevant testbed for economists and game theorists interested in the role social conventions play in the pursuit of maximum outcome.

From the economics standpoint, of considerable influence is the work of Schelling [26] who studied how norms can arise out of the iteration of localized simple-minded interactions.

The work by Delgado [9] should be mentioned too, for being especially relevant to the work that will be pursued under the scope of the DREAM project. The rise of social norms is studied with respect to the topology of the network that connect agents with their neighbours.

Bottom-up Modeling of Market Processes Using ACE systems to model the behaviour of specific actual markets in terms of self-organization to understand them and eventually take advantage of them is probably the most active field of ACE research. The focus has been mainly placed on electricity markets and, overall, financial markets. For the latter ones, one can see argues that the reason that makes financial markets especially well suited for study from an ACE standpoint is twofold. First of all, LeBaron [18] argues that, besides their providing the researcher with high quality, high frequency data on trades that are well-organized, centralized and efficient, financial markets are the most acute example of an economic system that organizes itself through multiple unconcerted interactions of individuals around a central coordinating mechanism (the price fixing process) that both acts as a referee and as an information provider. But more importantly, as underlined by Tesfatsion [28], it happens that the traditional analysis of financial markets, based on the rational expectations equilibria [23] theory is unable explain numerous empirical features observed in

financial markets such as fat-tailed asset return, high trading volumes, persistence and clustering in asset return volatility and cross correlations between asset returns, trading volume and volatility. And it turns out that a number of ACE oriented studies managed to satisfactorily reproduce some of those phenomena. As for electricity markets, they seem to be an especially well suited test bed for competitive double auctions settings where buyers and sellers compete and learn from each other pursuing their own interest and incidentally yielding high market efficiency.

Formation of Economic Networks Another possible topic of interest for ACE researchers is the way in which interconnection networks between traders form and evolve over time. Two kinds of connection networks are to be distinguished : those who link trading partners and those who represent information routes. As for the former ones, a particular kind of graph architecture, known as *small world networks* has drawn much attention from the field. These networks are defined by the fact that every node is both strongly connected to its nearest neighbours and has a number of shortcuts to other strongly connected neighbourhoods located far away. Besides their strong scaling properties, these networks have proved to yield surprisingly high market efficiency and to form naturally as they give incentives in terms of individual outcomes. Concerning information networks, much attention is placed on how information transmission and the connection topology that underlies it can influence the birth and rise of *information cascades* structures that in turn can trigger off panicky phenomena such as market crashes or speculative bubbles. An *information cascade* is said to happen when individual agents discard private rational information in favour of simple mimicking of what their neighbours do, thereby facilitating the spread of irrational behaviour.

Modelling of Organizations In economic terms, an organization is a group of people working together toward an objective that transcends those of its individual members. Agent-based computational modelling, together with Object Oriented Programming, which seems to have played a great role in organization modelling research [22], is well adapted to observe how, for instance, firms' organizational structures can influence the behaviour of the market they are involved in. This direction of ACE research currently does not call much attention from the economics community but this should change in the near future.

Design of Computational Agents for Automated Markets The use of automatic computational traders is a matter of potentially consequent profit as well as of considerable controversy. The design of efficient trading agents for asset trading or internet auctions is an extremely active area as experiments tend to show that they are able to achieve higher strategic efficiency and faster opportunity seizing than their human counterparts. The controversy [14] is in that information markets could be the first example of a setting where human beings would be completely replaced by artificially intelligent agents.

Parallel Experiments with Real and Computational Agents Driving similar experiments with both artificial agents and human subjects can be useful in several ways and can, in general, provide researchers with three things. *Validation* first of all, on the basis of similar observations, of the assumptions that were made for artificial agents and market structure design. *Improvement of the model*: observing human behaviour in clearly

defined market settings could allow for identification of key learning mechanisms that could be used to design more realistic artificial agents. *Understanding*, finally, could be brought as to why human societies behave in such or such way when confronted to such or such conditions. Provided indeed that these behaviours can be reproduced with artificial agent systems and as elementary reactions and learning mechanisms of agents are clearly defined, a clever researcher could possibly trace back the origin of whatever behavioural feature he is interested in by finding out what zones of the parameter space are favorable to the rise and persistence of these features.

Building ACE Computational Laboratories Studying all of the previously mentioned subjects requires the design and implementation of Agent-Based Computational Laboratories that allow for observation and parameter tweaking in adequate conditions.

Besides the fact that it might be healthy for economic theories to be calibrated and eventually validated to some extent by such simulation tools, it appears that the energetic development of these Computational Laboratories could be crucial for the development of the ACE field itself as it lives on the edge of two (to make it simple) usually unmixed fields, Artificial Intelligence and pure Economics and as it requires such a common playground to attract people from both sides, computer scientists being attracted by sophisticated models or appealing computational projects and economists being attracted by convenient tools that could free them from the imperious need of strong programming skills.

Remains however and unfortunately the issue that it seems quite doubtful that such a Computational Laboratory can be designed to be general enough to be of interest to other researchers than those who originated its implementation to answer their particular needs.

3 On financial markets: suggested introductory readings by LeBaron [17]

As we intend to simulate financial phenomena within the DREAM framework, this survey by LeBaron [17] appears to be especially relevant, being intended as a tutorial. It details six fundamental early works widely referenced in more recent papers. These early artificial financial market design are the following.

3.1 Lettau's adaptive agents and mutual fund flows

In his 1997 work [20], Lettau focuses on artificial traders' behaviours in a very simple setting. Agents have to decide how many shares of a risky (i.e. paying a stochastic dividend) asset they should hold. To do so, they are trying to maximize the utility function $U(\omega) = E(-e^{\gamma\omega})$ where $\omega = s(d - p)$. s is the number of shares being held, d the dividend and p the price, which is set exogenously. It is a well known result that the optimal solution is linear in terms of p and $\mu(d)$ the mean dividend: $s^* = \alpha^*(\mu(d) - p)$. α^* is therefore the only unknown on the way to optimality and is the only decision variable agents are trying to optimize.

Lettau uses a traditional GA to search for α^* : each agent is represented by a bitstring that encodes a candidate alpha value as a real number. α^* can also

be derived analytically and experimental runs show that the GA consistently converges to the appropriate value. Two other findings besides arise from these experiments. First, there appears to be a bias toward risky attitudes and the convergence value tends to be greater than α^* . This can be explained by too weak sampling that triggers rewarding of risky but lucky strategies. Second, agents, in Lettau’s own words, “exhibit an asymmetric response after positive and negative returns where the portfolio adjustment is more pronounced after negative returns.” This means that agents tend to withdraw their money faster after bad days than they invest new money after good ones. Interestingly enough, as Lettau shows with real financial data representing cash flows in a variety of mutual funds, this behaviour is also exhibited by real traders.

This artificial market setting is controversial in two ways. First of all, it uses a somewhat rusty bitstring representation that yields disruptive mutations in the amounts of holdings. Secondly, the price is set exogenously, which appears as very counterintuitive when endogenous price fixing is the central coordinating mechanism of financial markets. This latter simplification however is meant to allow for focus on agents’ behaviour.

3.2 Zero-Intelligence traders

This other very influential study by Gode and Sunder [10] aims at questioning how rationality distribution influences market behaviour and efficiency. More precisely, they want to find out what is the minimum amount of rationality that is needed by artificial agents to behave in way that reasonably close to real to what is observed in experiments involving human subjects.

They base their observations on a double auction setting where two communities, buyers and sellers make successive bids and offers until a deal is concluded. Agents behave randomly except for one only restriction: buyers cannot buy an asset for a price that would be higher than what the asset is worth to them (its redemption value) and, reciprocally, a seller cannot sell an asset at a price that is lower than what he paid for it in the first place.

Experiments show that with this only constraint, and as opposed to strictly random ones, artificial agents based markets are able to achieve to close to perfect efficiency and to exhibit a behaviour, in terms of price series, that is very close to that observed with humans. The conclusions to be drawn out of this work is that ACE researchers should be quite cautious on what to focus on when it comes to explaining market behaviour for this study shows that market structure can play a great role and take over individual’s rationality to impose a seemingly robust behaviour.

3.3 Arifovic and the behaviour of exchange rates

Arifovic’s experiments [2] are based on what is called an overlapping generations economy. It means that individual agents live a constant alternation of two periods. At period t , N “young” agents, said to be of generation t , come to life. At period $t+1$, those are said to be “old”. The economy involves a single consumption good. Agents are endowed with w^1 units of it when they in their young period and with ω^2 when they are old. They consume part of these goods: $c(t)$ when young, $c(t+1)$ when old. When they are young, they also have the possibility to sell those goods to constitute savings. Our agents live in two countries, 1 and 2, with two different currencies and therefore two different prices for the consumption good. Agents are trying to maximize their utility function:

$$U(t) = \ln(c(t)) + \ln(c(t+1))$$

subject to the following constraints:

$$c(t) \leq w^1 - \frac{m_1(t)}{p_1(t)} - \frac{m_2(t)}{p_2(t)}$$

and

$$c(t+1) \leq w^2 + \frac{m_1(t)}{p_1(t+1)} + \frac{m_2(t)}{p_2(t+1)}$$

m_i represent holdings in the different currencies acquired by agents when young. An agent's savings is equal to the sum of holdings in both currencies:

$$s(t) = \frac{m_1(t)}{p_1(t+1)} + \frac{m_2(t)}{p_2(t+1)}$$

The exchange rate between the two currencies, center of this study, is defined as:

$$e(t) = \frac{p_1(t)}{p_2(t)}$$

An important particularity is that, differently from Lettau [20], these equilibrium considerations are made within a framework where the price is fixed endogenously.

Arifovic uses a genetic algorithm to evolve her population of agents. Each agent is represented by a bitstring that encodes three real numbers: first period consumption and savings in both currencies. She then follow standard GA techniques except for what she calls the *election operator* that prevents an offspring to enter the population if it is said to perform worse than its parents.

She conducts agent based and human subjects experiments in parallel and shows that the same behaviour is observed in both cases. More precisely, and contrarily to what the rational expectations theory would predict, the first period consumption is shown to settle down at a constant value while the exchange rate keeps on oscillating, underlining the fact that there are no incentive to place savings in one currency rather than in the other.

3.4 Routledge on costly information acquisition

Routledge [24] places his agent in a context that corresponds to the canonical framework described by Grossman and Stiglitz [11]. Traders have the possibility to purchase information at some non zero cost that helps them guessing what the dividend of the risky asset (as opposed to a risk free but non profitable one) will be. They can be said to purchase a noisy signal of the dividend. Agents aim at maximizing their utility functions by finding out the optimal amount of shares of the risky asset they should hold. This learning process is based on the expectations agents can form about the value of the dividend. Informed traders base theirs on a linear ponderation of the signal y they purchased:

$$E^i(d|y) = \beta_0^i + \beta_1^i y$$

Uninformed traders have to base their expectations on the sole basis of the public price:

$$E^u(d|p) = \beta_0^u + \beta_1^u p$$

. A genetic algorithm is used to evolve these forecasting parameters. Each agent is represented by a vector of the form

$$(\theta, \beta_0^i, \beta_1^i, v^i, \beta_0^u, \beta_1^u, v^u)$$

where $v^{u/i}$ are the conditional variances for uninformed and informed traders and θ is a bit that is set to 0 for uninformed traders and to 1 for informed ones.

A rational expectations equilibrium in such a case is known to exist as proven by Grossman and Stiglitz [11] and Routledge shows that his learning framework does converge to it eventually, provided adequate parameter settings are made. It is to be noted however about that equilibrium that its stability is strongly sensitive to parameters setting the amount of noise. This might an encouraging clue on the ability of ACE experiments to identify the aforementioned boundary between settings where rational expectations theories are valid and settings where general collective instability prevails.

3.5 The Santa Fe artificial stock market

This is probably the most famous artificial stock market experiment [4]. It again borrows from canonical frameworks such Grossman and Stiglitz' [11] and is original in the sense that it uses classifier systems as a learning mechanism. The market design is simple and analytically tractable so that plausibility of the observed behaviour can be assessed.

Agents have to decide between a risky asset bringing a stochastic dividend and a risk free bond in infinite supply paying at a constant rate r . The risky stock's dividend is as follows:

$$d_t = \mu(d) + \rho(d_{t-1} - \mu(d)) + \epsilon_t$$

where ϵ_t is a gaussian noise and ρ is a constant fixed to 0.95. It is well known that under these conditions, the demand for shares of the risky asset is given, for agent i at time t by:

$$s_{t,i} = \frac{E_{t,i}(p_{t+1} + d_{t+1}) - p_t(1+r)}{\gamma\sigma_{t,i,p+d}^2}$$

where p is the price for a share of the risky asset and γ the risk aversion coefficient. Agents will derive forecasts for the value of p and use this equation to calculate the amount of risky shares they should hold. Forecasts are modelled as follows:

$$E_{t,i,j}(p_{t+1} + d_{t+1}) = a_{i,j}(p_t + d_t) + b_{i,j}$$

Agents are therefore to come up with accurate values of a and b along with an estimation of $\sigma_{i,j}^2$ to be able to generate their demand function.

A standard classifier system is used as a learning mechanism. Each agent is given a set of rules, encoded as bitstrings (making use of the # generalization symbol), that map a particular state of the economy to a particular (a, b, σ) triplet. The condition part of the rules uses 13 bits representing the following conditions:

$$\mathbf{1-7} \quad \frac{p * Interest}{d} > \frac{1}{2}, \frac{3}{4}, \frac{7}{8}, 1, \frac{9}{8}, \frac{5}{4}, \frac{3}{2}$$

8 $p > 5\text{-period Moving Average (MA)}$

9 $p > 10\text{-period MA}$

- 10 $p > 100$ -period MA
- 11 $p > 500$ -period MA
- 12 control bit, always on
- 13 control bit, always off

Agents pick a rule among their personal sets according to how well it has performed in the past. This performance is measured in terms of squared forecast error, the same conditional variance that is used in the calculation of the demand function:

$$\sigma_{t,i,j}^2 = \beta \sigma_{t-1,i,j}^2 + (1 - \beta)((p_{t+1} + d_{t+1}) - E_{t,i,j}(p_{t+1} + d_{t+1}))^2$$

Every K periods, worst performing rules are replaced through the standard course of genetic recombination and mutation. The system appears to be very sensitive to the value of K : high values (around 1000) yield a behaviour of the market that is close to what is expected in terms of Rational Expectations Equilibrium which is not the case with low values (around 250) where agents start to make use of technical trading (which should not make sense in a REE context). These results too might shed light on the rationality boundary there is between theoretical and actual trading. The interested reader can refer to [16, 19] for a variety of other findings and insights into the design of the SFI artificial stock market.

3.6 Using Artificial Neural Networks

Beltratti and Margarita [7] work with a population of agents that evolve forecasts of future prices using neural networks. Three kinds of agents exist: Dumb (D), Naive (N) and Smart (S) ones that differ in both the data they use (for instance how far back in the history of prices they go to derive their predictions) and the structure of the ANN they use to compute the forecast, namely the number of neurons in the hidden layer. Trade happens in a decentralized way: agents are matched randomly and trade at an intermediate price situated between their two expectations. Additionally, traders are allowed to move from one category to another by purchasing neural complexity at a given cost. This particular feature places this study right in the the debate of costly information acquisition set up by Grossman and Stiglitz [11] and also tackled by the aforementioned study by Routledge [24].

Experiments study the impact of information cost on the composition of the population at the end of run. It is observed that if certain intermediate cost levels allow for an ecology of the different types, most values with high magnitudes lead to the complete domination by the corresponding type (smart agents tend to take over the population if cost is low and naive ones do so if costs are prohibitive). A more refined analysis of the dynamics of the population's composition however shows that smart agents tend to dominate early in the run when prices are highly volatile and are then taken over by naive agents that progressively invade the whole population as prices get stabilized.

4 A builder's guide to Agent Based Financial Markets [18]

4.1 Context

To take a more pragmatical turn, this section will follow a paper by LeBaron [18] that aims at helping researchers interested in building their own artificial stock markets. To do so, he quickly positions the problem in terms of research orientations and then enumerates crucial design issues to which the newcomer will necessarily end up being confronted. These crucial issues are briefly outlined hereunder.

4.2 Design issues

Agents The very heart of an agent-based artificial market is course the way in which agents themselves are designed. This means the designer has to decide on the amount of rationality each agent is going to have and what heuristics it is going to use to map perceptions of its environment (e.g. series of prices) to relevant actions and decisions. Three main categories are identified with increasing levels of complexity:

1. *Zero Intelligence (ZI) traders* [10] are the easiest one to implement: they have no rationality whatsoever and behave in a random manner. These are used to isolate the influence of market structures and rules on the behaviour of the market as well as to identify the impact of local irrationality.
2. *Rule-Based traders* use a dynamic set of deterministic logic rules that maps states of the world to subsequent decisions taken by the agent. Implementation is straightforward as rules are described in explicit terms. Critiques of this approach, typical of classical AI, are mainly focused on its hard-wired character that inadequately stand for the fuzzy, dynamic and coevolutionary aspects of collective human behaviour.
3. *Adaptive traders* use modern fuzzy heuristics and representations such as Artificial Neural Networks or Evolutionary Algorithms. This allow for a continually changing set of innovative strategies. The absence of need for explicit assumptions is also supposed to be a more realistic model of the natural emergence of the interactions patterns of traders with each other and the market. Critiques mainly have to do with potential explosion in computational and strategy complexity, with the poor mathematical understanding we have of the dynamics of evolutionary algorithms and with the fact that building an evolutionary model that comes up with relevant indicators and strategies from scratch is impossible *stricto sensu*.

Trading mechanisms As illustrated by works with ZI traders [10], the market structure, i.e. the rules according to which financial assets are exchanged, prices are fixed or information is spread can have a great influence on the system's behaviour. Design decisions of course are to be made accordingly to what one aims at questioning. Three ways are identified to design the central price fixing mechanism:

1. *Use direct response to Supply/Demand discrepancies.* The price simple goes up when there is excess demand and down when there is

excess supply. This method, that adequately represents a market in systematic disequilibrium is found to be somewhat too sensitive to parameters tuning as can be seen in [3].

2. *Compute temporary equilibrium prices as in [2, 4].* This method requires well defined demand functions and fixes, analytically or computationally, prices to allow for a temporary balance. Although these methods get rid of the parameter sensitivity issues and allow for varying market depth, they rely on strong assumptions regarding demand behaviours that might make them irrelevant models of real high frequency trading.
3. *Model actual mechanisms of continuous trading.* Mimicking complex rules used in actual markets allows for refined analysis but is only doable in a straightforward way if the said rules are simple and deterministic enough. Otherwise, in the case the market is cleared by human intervention, yet another modelling and learning is introduced along with its cortege of design issues.

Securities The financial assets exchanged by agents also come with several characteristics that need to be modelled. These characteristics are the pillars on which traders are going to base their strategic choices and position the security in the *Profitability/Risk/Liquidity* triangle. Securities are usually modelled in a very simple way to allow for tractability and focus on other features of the market but two features appear to be profitable ways of design complexification:

- To correspond to real world situations, information about the security's fundamental should not be revealed without noise.
- Most artificial markets deal with one or two types of securities. A great deal of classic financial phenomena or strategies such as portfolio diversification cannot be studied in such frameworks.

Evolution Designers have to be especially cautious about how, when, and where to implement evolution. Especially because evolution is the key competitive mechanism that drives both the individual agent's learning and the overall market's structuration and also because distributed agents systems are to be analysed in intricate coevolutionary terms, which makes it very important to understand what evolutionary mechanisms come into play and in what relation to each other.

Often delicate in socio-economic settings is the question of the choice of the fitness function and even if financial markets usually provide with a wealth of potentially relevant indicators (utility, forecast accuracy, average wealth, etc.) care should be taken to ensure that fitness measures are robust enough to resist misleading noise and to encompass the variety of criteria that can justify an agent's survival.

Benchmarks The crucial issue of validating agent-based models by comparing them to real markets remains a difficult one and stands as an open problem. It is important however to tackle it as it is the only way to show that these toy problems ACE researchers spend so much time and energy on are of any interest to real world practitioner. LeBaron [18] makes three suggestions:

1. Study cases where the behaviour of the market is well defined and analytically tractable so one can check that the artificial market indeed

behaves as it is expected to. Such an observation gives credit to the considered market even outside of the well understood cases.

2. It is important to go beyond the simple finding of parameter settings that allow for the reproduction of actual market behaviour, researchers should try to gain a more refined knowledge of the parameter space and be able, in particular, to find out what the boundaries are that separate simple, rational behaviours and more complex or chaotic ones.
3. A final possibility consists in using parameters values estimated from actual economic data. This traditional approach could help theoretical markets get a little closer to reality but it does not go without usual controversies about how well numerical values can be carried on from one context to another.

Time Time is an important aspect of dynamic evolutionary systems. LeBaron [18] outlines three directions of concern:

1. *How to deal with the past ?* Any agent that does forecasting or learning needs to be concerned with how long its memory should be and with how ambitious its historical data sampling should be.
2. *How fast should individuals adapt to changes ?* This parameter could have a huge impact on the behaviour one can observe and conditions the possibility to settle to states of equilibrium.
3. *What about synchronicity ?* This particular issue is ignored by most of the current ACE research where trading is cleared in a clean synchronous fashion. In the real world on the contrary, markets are made of asynchronous, high frequency exchanges. This difference issue needs to be addressed in the future by market designers concerned with realism and the implementation of artificial financial markets on large scale computing platforms such as the DREAM seems to be especially relevant in this context.

5 Conclusion

Agent-based Computational Economics and Finance is an extremely active area of research that has the particularity to be consistently sound from both the theoretical and applicative standpoints. It moreover comes with an avenue of open questions that leaves much room for exploration, especially by researchers equipped with refined and modern Artificial Intelligence and/or Computer Science tools. As such, ACE research could greatly benefit both from a competent practice of advanced Evolutionary Computation and from a sophisticated distributed computing framework, making it an especially well suited subject of interest for the DREAM project.

References

- [1] Y. S. Abu-Mostafa, B. LeBaron, A. W. Lo, and A. S. Weigend, editors. *Computational Finance 1999*. The MIT Press, Cambridge, MA, 1999.
- [2] J. Arifovic. The behavior of the exchange rate in the genetic algorithm and experimental economies. *Journal of Political Economy*, 104(3):510–541, 1996.

- [3] W. B. Arthur. Inductive reasoning and bounded rationality (the el farol problem). *American Economic Review*, 84:406, 1994.
- [4] W. B. Arthur, J. Holland, B. LeBaron, R. Palmer, and P. Tayler. *The Economy as an Evolving Complex System II*, chapter Asset pricing under endogenous expectations in an artificial stock market, pages 15–44. Addison-Wesley, Reading, MA, 1997.
- [5] R. Axelrod. *The complexity of cooperation: Agent-based models of conflict and cooperation*. Princeton, NJ: The Princeton University Press.
- [6] R. Axelrod. *The Evolution of Cooperation*. New York, NY: Basic Books, 1984.
- [7] A. Beltratti and S. Margarita. *From Animals to Animats 2*, chapter Evolution of trading strategies among heterogeneous artificial economic agents. The MIT Press, Cambridge, MA, 1993.
- [8] H. Dawid. *Adaptive learning by genetic algorithms: Analytical results and applications to economic models*. Berlin: Springer Verlag, 2nd edition, 1999.
- [9] J. Delgado. Emergence of social conventions in complex networks. *Elsevier Artificial Intelligence*, 141:171–185, 2002.
- [10] D. K. Gode and Sunder S. Allocative efficiency of markets with zero-intelligence traders: Markets as a partial substitute for individual rationality. *Journal of Political Economics*, 101(1):119–137, 1993.
- [11] S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *The American Economic Review*, 70:393–408, June 1980.
- [12] F. A. Hayek. The use of knowledge in society. *The American Economic Review*, XXXV(4):519–530, September 1945.
- [13] H. Joyce. Adam smith and the invisible hand. *+plus Journal, Issue 14* - <http://plus.math.org>.
- [14] J. O. Kephart. Software agents and the route to the information economy. *Proceedings of the National Academy of Sciences*, 99:7207–7213, 2002.
- [15] T. Krebs. Rational expectations equilibrium and the strategic choice of costly information. Brown University, working paper 2001-44, September 2001.
- [16] B. LeBaron. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23:1487–1516, 1999.
- [17] B. LeBaron. Agent based computational finance: Suggested readings and early research. *Journal of Economic Dynamics and Control*, 24:679–702, 2000.
- [18] B. LeBaron. A builder’s guide to agent based financial markets. *Quantitative Finance*, 1(2):254–261, February 2001.
- [19] B. LeBaron. Building the santa fe artificial stock market. Working Paper, Graduate School of International Economics and Finance, Brandeis University, Waltham, MA., June 2002.

- [20] M. Lettau. Explaining the facts with adaptive agents: The case of mutual fund flows. *Journal of Economics Dynamics and Control*, 21:1117–1147, 1997.
- [21] B.G. Malkiel. *A Random Walk Down Wall Street*. W. W. Norton and Company, 8th edition, April 2003.
- [22] M. J. Prietula, Carley K. M., and L. Glasser. *Simulating organizations: Computational models of institutions and groups*. The MIT Press - Cambridge, MA, 1998.
- [23] R. Radner. Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica*, 47(3):655–678, May 1979.
- [24] B. R. Routledge. Genetic algorithm learning to choose and use information. *Macroeconomic Dynamics*, 5(2), April 2001.
- [25] B. Ruffieux. Les marchés testés en laboratoire. *Pour la science*, 307:85–92, May 2003.
- [26] T. C. Schelling. *Micromotives and macrobehavior*. New York, NY: W. W. Norton and Co., 1978.
- [27] A. Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations*. 1776.
- [28] L. Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2001.
- [29] L. Tesfatsion. *Special Issue on Agent-Based Computational Economics*, volume 18. Computational Economics, October 2001.
- [30] L. Tesfatsion. *Special Issue on Agent-Based Computational Economics*, volume 25. Journal of Economic Dynamics and Control, March 2001.
- [31] L. Tesfatsion. *Special Issue on the Agent-Based Modeling of Evolutionary Economic Systems*, volume 5. IEEE Transactions on Evolutionary Computation, October 2001.
- [32] N. J. Vriend. Was hayek an ace ? *Southern Economic Journal*, 68(4):811–840, 2002.